

Learning, Behavior, and Evolution

Domenico Parisi Stefano Nolfi Federico Cecconi
 Institute of Psychology
 National Research Council - Rome
 e-mail: domenico@irmkant.Bitnet stiva@irmkant.Bitnet
 cecconi@irmkant.Bitnet

Abstract

We present simulations of evolutionary processes operating on populations of neural networks to show how learning and behavior can influence evolution within a strictly Darwinian framework. Learning can accelerate the evolutionary process both when learning tasks correlated with the fitness criterion and when random learning tasks are used. Furthermore, an ability to learn a task can emerge and be transmitted evolutionarily for both correlated and uncorrelated tasks. Finally, behavior that allows the individual to self-select the incoming stimuli can influence evolution by becoming one of the factors that determine the observed phenotypic fitness on which selective reproduction is based. For all the effects demonstrated, we advance a consistent explanation in terms of a multidimensional weight space for neural networks, a fitness surface for the evolutionary task, and a performance surface for the learning task.

1. Introduction

Are behavior and learning among the causes of evolution? Do they influence the course or the rate of evolution? With reference to behavior Plotkin has written: "Whether behavior is also a cause and not just a consequence of evolution is a significant theoretical issue that has not received the attention it deserves from evolutionary biology" (Plotkin, 1988, pag. 1). He notes that the subject index of Mayr's *The Growth of Biological Thought* contains only three entries for "behavior". One might add that there is no single entry for "learning".

There are several reasons why the question whether learning and behavior are among the causes of evolution tends to be ignored by evolutionary biologists, and by biologists generally. One reason is that the orthodox view represented by the Modern Synthesis tends to be reductionist, which implies that the causes and basic mechanisms of evolution are only to be found at the level of genetics. Behavior and learning are too wholistic to be considered as important to understand the intimate nature of evolutionary processes. Another reason is that the idea that such phenotypic processes as behavior and learning might be among the causes of evolution

sounds too Lamarckian, and the rejection of inheritance of lifetime changes in the phenotype is one of the foundations of the Modern Synthesis. A third reason is that behavior and learning are the province not of biology but of psychology and ethology, and biologists would hardly admit that among the causes of the central phenomenon studied by their discipline and the theory of which is the organizing framework for biology (evolution and the theory of evolution), there are processes on which are competent such "soft" disciplines as psychology and ethology.

But there is still another reason that can explain why behavior and learning are not seriously considered as possible causes of evolution within biology. The various claims that have been advanced in the course of the present century in support of the idea that behavior and learning can influence evolution have had a limited empirical basis and, what is even worse, have been generally rather vague conceptually. Therefore, it has been easy for evolutionary biologists to dismiss these claims as irrelevant and to consider the whole issue as marginal at best.

The purpose of this paper is to examine some aspects of the problem "Do learning and behavior influence evolution?" within the framework of neural networks

(Rumelhart and McClelland, 1986) and genetic algorithms (Holland, 1975; Goldberg, 1989). The justification for such an enterprise is that one can hypothesize that the theoretical apparatus of the theory of complex dynamic systems (of which neural networks and genetic algorithms are considered here as special applications) and the methodology of computer simulation typically used in research on neural networks and genetic algorithms can help make more precise and testable claims on the role played by learning and behavior in evolution. If we succeed in convincing students of evolution that this hypothesis is a reasonable one, then we will have contributed to eliminating at least the last of the four reasons listed above for ignoring this potentially very important issue.

2. How can learning help evolution

The orthodox view of evolution is that changes due to learning during life are not inherited and, more generally, that learning does not influence evolution. The basis for such a view is the physical separation between the germ cell line and the somatic cell line. Changes due to learning concern somatic cells whereas evolution is restricted to the germinal cells. Since the two types of cells are physically separated, it is impossible that whatever happens to the somatic cells can have an influence on evolution. On the other hand, Baldwin (1896), Waddington (1942), and several others have claimed that there is an interaction between learning and evolution and, more specifically, that learning can have an influence on evolution.

Computer simulations that apply evolutionary methods to populations of neural networks have recently shown that changes during the 'life' of individual neural networks which are not inherited can still have an influence on the course of the evolutionary process. Hinton and Nowlan (1987) have demonstrated how modifying at random the connection weights during life allows the simulated evolutionary process to select networks that are more adapted to the given task. Belew (1989) indicates how the beneficial effect of learning on evolution can increase if the weight changes are not random but are correlated with the task for which the networks are being selected. Ackley and Littman (1991) and Nolfi and Parisi (1991) have shown that when evolution is free to select what networks will learn during their life, useful learning tasks are evolved yielding an increase in performance with respect to the situation in which lifetime learning is not allowed. Nolfi and Parisi (1991) also demonstrate how an evolved learning capacity might emerge and then get

extinguished if it is no longer useful for the evolutionary process.

In order to analyze how learning can help evolution in simulated organisms let us consider the simple artificial organisms described in Nolfi, Elman, and Parisi (1990). Each organism (O) lives in a bidimensional environment containing randomly distributed pieces of food. The organisms are modelled by a feedforward neural network (Figure 1) which basically receives sensory input from the environment concerning the position of the nearest food element and generates as output motor actions that allow the organism to displace itself in the environment.




Figure 1. O's architecture. Sensory input is encoded by the 2 input units representing the angle and the distance of the nearest food element (both values are scaled from 0.0 to 1.0). Movement is encoded in the 2 output units that codify four possible actions: go ahead, turn left, turn right, and stay still. Outputs on the motor output units are feedback as additional input at time $t+1$.

In each activity cycle activation spreads from the input units through the hidden units to the output units. The kind of movement of the organism which is generated by the network given a certain sensory input depends on the quantitative weights on the connections of the network. Initially these weights are assigned at random and therefore the organism wanders randomly in the environment.

In order to obtain Os which are able to reach food elements in an efficient manner an evolutionary method based on selection and mutation is used. The process starts with 100 Os, each having the same architecture and a different random assignment of

connection weights. This is Generation 0 (G0). G0 networks are allowed to "live" for 20 epochs, where an epoch consists of 250 actions in 5 different environments (50 actions in each), for a total of 5000 actions. The environment is a grid of cells with 10 randomly distributed pieces of food. Os are placed in individual copies of the environment, i.e. they live in isolation.

At the end of their life Os are allowed to reproduce. However, only the 20 Os which have accumulated the most food in the course of their random movements are allowed to reproduce by generating 5 copies of their weight matrix. These best ranking individuals have been assigned for purely random reasons weight matrices that cause them to sometimes respond to food elements by approaching them. The $20 \times 5 = 100$ new Os constitute the next generation (G1). Mutations are introduced in the copying process by selecting at random 5 weights and adding a random value between +1.0 and -1.0 to these weights.

After the Os of G1 are created they also are allowed to live for 5000 cycles. The behavior of these Os differs slightly from that of preceding generation (G0) as a result of two factors. First, the 100 Os of G1 are the offspring (copies) of a subset of the Os of G0. Second, the offspring themselves differ slightly from their parents because of the mutations in their weights. These differences lead to small differences in average food eaten by the Os in G1 with respect to those of G0. At the end of their life the 20 best individuals are allowed to reproduce 5 times, forming G2. The process continues for 50 generations.

Mutations can result in offspring that are better than their parent or offspring that are less good. However, selective reproduction will insure that the former individuals will be more likely to reproduce than the latter. The net result is a progressive increase in food approaching ability due to selective reproduction and random mutations.

Nothing changes in the neural networks during their life up to this point. We then run another set of simulations in which we added a learning task during life. Os learned to predict the sensory consequences of their own actions, i.e. how the sensory information from a food element was going to change when a planned action was actually executed. Using backpropagation (Rumelhart, Hinton, and Williams, 1986) the networks were taught to specify at time T in two additional output units (prediction units) the sensory input that the network will receive at time T+1 (see Figure 2).




Figure 2. O's architecture. The two additional output units codify O's prediction.

The addition of the learning task during life increases the power of the evolutionary process. Even if none of the changes which occur in the parent's weight matrix as a consequence of learning are transmitted to offspring, populations of Os which evolve with such a learning are able to reach a larger number of food elements than Os evolved without learning (see Figure 3).




Figure 3. Average number of food elements eaten by successive generations for organisms evolved with and without life-learning. Each of the two curves represents the average performance of 10 different simulations with different random assignment of weights.

It is important to notice that the performance increase is not due to the fact that networks, having the possibility to learn, increase their fitness during their life. Such an increase exists (in other words, learning to predict how food position changes with the organism's actions leads to a better food approaching performance during life) but it is not enough to explain the difference between the two curves. In fact, life learning allows the evolutionary process to evolve networks that perform better at birth, i.e. before learning takes place. This means that learning, in addition to have a life-time adaptive function, has an evolutionary function that results in an increase of offspring's fitness.

How can learning during life have an influence on evolution if inheritance is strictly Darwinian (or better, Waismannian) and not Lamarckian? In other words, if the learned changes in connection weights are not inherited and a reproducing individual transmits its intact inherited weight matrix to its offspring? To try to answer this question we must look at the fitness of genotypes in a more abstract way.

We begin by identifying the inherited genotype of an individual with a specification of the weight matrix of the individual's neural network at birth. It is this specification that is transmitted to the individual's offspring if the individual is among those that reproduce. By definition the genotype does not change except for mutations. We can view one particular genotype (weight matrix) as a single point in a multidimensional abstract space. Each dimension of this space corresponds to one particular network connection. Hence, if our networks have N connections, the corresponding space will be an N -dimensional space. A particular network will occupy the position on a dimension that indicates the quantitative weight of the corresponding connection. Therefore, each point in the N -dimensional space represents one particular weight matrix, and all possible weight matrices are represented in the space.

We now assume that each possible genotype (point in the N -dimensional space of weight matrices) has a certain fitness. That is, if an individual with that genotype (weight matrix) is allowed to live in some environment for a certain lifespan, it will generate a behavior which will result in a certain fitness value given a certain fitness criterion¹. As a consequence, we can conceive of fitness as an additional dimension of the space, with weight matrices with a higher fitness located higher on that dimension than matrices with a lower fitness. Since the additional dimension of fitness is added to a space that already has N dimensions, we will talk of a (multidimensional) fitness "surface". If the fitness surface is "smooth" this means that weight matrices that are near to each other in weight space will have similar fitnesses, whereas if the fitness surface is "rugged" one cannot predict what the fitness of a particular weight matrix will be given the fitness of a nearby matrix.

A first generation of randomly generated weight matrices is a collection of randomly distributed points in weight space. Selective reproduction means that the reproducing individuals have weight matrices that correspond to higher points on the fitness surface than those of non-reproducing individuals. Consider now

mutations. Mutations mean that a reproducing matrix (parent) is replaced by one or more matrices (offspring) corresponding to points on the fitness surface located in a region of that surface just around the point to which the parent's matrix corresponds². Since mutations are random, offspring matrices are a random sampling of points in that region.

Consider now two very different individuals (weight matrices) corresponding to distant points in weight space, which have the same fitness (hence, the two points are the same height on the fitness surface). Since the two individuals have the same fitness, selective reproduction has no way to choose between them. In fact, if there were no mutations it would be irrelevant from an evolutionary point of view to choose between them. In such a case the offspring of one individual would be exact copies of their parent and no consequence for the next generation's average fitness would result from choosing one of the two individuals rather than the other. However, since there are mutations, an individual will be replaced by individuals which are similar to their parent but not exact copies. More specifically, the offspring's weight matrices will sample the region surrounding the parent's point on the fitness surface.




Figure 4. 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 as represented as a single dimension.

On consequence of this is that, although the two individuals correspond to points that are the same height on the fitness surface, the average fitness of the next generation will depend on the nature of the two regions surrounding the two points. The two points (parents) may be equally high but the region surrounding one of them can include points (offspring) that are on the average higher than the points in the surrounding region of the other (see Figure 4). It would then be appropriate for the selection process to select the individual with a better surrounding region rather than the other since the offspring of the former individual will be on average better than the offspring of the latter. More generally, it would be useful for the

selection process to know the nature of the surrounding regions of candidates for reproduction since the fitness of offspring is more important than the fitness of their parents from the point of view of the next generation's average fitness. Selective reproduction per se has no means to know that. It sees the heights of candidates for reproduction on the fitness surface but it does not see their surrounding regions. Our hypothesis is this is exactly what learning does: learning illuminates the regions on the fitness surface surrounding the points on that surface corresponding to candidates for reproduction, and makes what it sees available to the selection mechanism. The result is that selection is improved and there may be a positive influence of learning on evolution³.

Learning involves weight changes and therefore it implies a movement in weight space of the point representing an individual matrix. If an individual X has a better surrounding region than another individual Y, even if both start from the same height on the fitness surface, by moving in weight space (learning) X is more likely to end up on a higher point on the fitness surface than Y, and therefore more likely to be selected for reproduction.

Given this sort of analysis, even random changes in the weight matrix during life-time, are enough to ensure the selection, on average, of the better between two individuals that have the same fitness but correspond to points on the fitness surface with different surrounding regions. As a consequence, random changes in weights during life should result in a positive effect on evolution and this is exactly what Hinton and Nowlan (1987) have found. We have run some simulations in which individual networks are taught by using randomly generated teaching inputs on the two additional output units (see Figure 2), and in this case too there is a positive influence of learning on evolution, even if the improvement is less great than in the case of prediction learning (see Figure 5).

It remains to be explained why learning a task such as predicting the consequences of one's own actions, which is correlated with the task for which organisms are selected, results in a larger beneficial effect on evolution.




Figure 5. Average number of food elements eaten by successive generations for organisms evolved without life learning, with random learning, and with prediction learning. Each of the three curves represents the average performance of 10 different simulations with different random assignments of weights.

Our hypothesis is that some types of explorations of surrounding regions can be more "intelligent" than others, in the sense that they preferentially explore the sections of the surrounding region with good fitness points. If such points exist, i.e. if there are points in the surrounding region of an individual that have a higher fitness than the individual's fitness, an intelligent exploration of the surrounding region would increase the reproduction chances of such an individual. Individuals located in regions with such high fitness points should be preferred by the selection process because they will have some probability of generating at least some offspring better than themselves. Exploring an individual's surrounding region in such a way that sections with higher fitness points can be detected allows exactly that. It allows the selection process to prefer an individual located in a region with higher fitness points than itself to another individual even if the two regions globally include points with the same average fitness (see Figure 6). In the situation depicted in Figure 6, random learning (i.e. random exploration of the surrounding region) would not confer to point B more reproductive chances than to point A. If we want point B to be selected rather than A, since it is more likely to have better offspring than itself, we need an intelligent exploration of the surrounding regions, i.e. a movement of the point on the fitness surface such that the point is more likely to end up on a higher fitness level if there is such level.




Figure 6. Fitness of all possible weights matrices. Point B has A better surrounding region than point A, because in its surrounding region there are points with higher fitness than points in B's surrounding region, even if the two points have identical fitnesses and identical surrounding regions on average.

This we believe is what takes place in the simulations with prediction learning. Non-random learning tasks have a performance surface which is analogous to the fitness surface of evolutionary tasks. Each specific weight matrix corresponds to a particular point (height) on the performance surface for the given learning task. If the performance surface of a certain learning task and the fitness surface of a certain evolutionary task (as defined by a certain fitness criterion) are correlated, i.e. a matrix of weight which is good at the learning task is also good at the evolutionary task, and viceversa, then learning the task during life will have a stronger effect on evolution than just random learning. The reason is that learning the task implies weight changes that are also useful from the point of view of evolutionary fitness. Hence, by moving during learning to higher positions on the performance surface of the learning task, point B in Figure 6 will be simultaneously pushed toward higher positions on the fitness surface - which is impossible for point A. The net result is that B will have more chances of reproduction - which is useful from the point of view of evolution.

It is interesting to note that if one allows the evolutionary emergence of the life-time learning task rather than arbitrarily deciding what is the task at the outset, as in Ackley and Littman (1991) and Nolfi and Parisi (1991), evolution can select the learning tasks which is most appropriate for the displacement of points in weight space. In other words, it is plausible to expect that the learning task will change, at the evolutionary time-scale, in order to obtain the most intelligent exploration of the surrounding regions of individual points. This actually happens in the computer simulations described in Nolfi and Parisi (1991).

3. Indirect inheritance of acquired characters

In the preceding section we showed that, in our simulated organisms, learning can help the evolution of adaptive behavior and we discussed how this effect can be explained without postulating the inheritance of acquired characters. The existence of a phenomenon of this type has been postulated by Waddington (1942). This is one aspect of the possible interaction between evolution and learning. In the present section we will examine another related aspect of this interaction, that is, we will try to figure out if learning some ability during life can facilitate the acquisition of that ability in successive generations. In other words, we want to demonstrate that a learning ability can be indirectly transmitted to descendants even if the inheritance mechanism remains strictly Darwinian and individual are not selected for that learning ability.

If we analyze the results of the simulations described in the previous section and in Nolfi, Elman, and Parisi (1990), we find some evidence of inheritance of acquired characteristics. Figure 7 graphs the average error curve for the prediction task learned during life by Os belonging to the first and to the last generation. Although both groups of Os start from an identical error level at the beginning of their life (i.e. there is no inheritance of the capacity to predict at birth), the Os of the last generation learn more of the prediction task than the Os of the first generation. In other words, in a population of Os that learn to predict the sensory consequences of their own actions during their life there is an observed increase in the ability to learn the task generation after generation. Hence, we can conclude that there is inheritance of the ability to learn the particular task, although not directly of the ability to perform the task.




Figure 7. Global error on the prediction task as a function of epochs of training for naive Os (i.e. Os of generation 0) and evolved Os (i.e. Os of generation 49).

How such inheritance of a capacity to learn to perform

a task might be explained? The answer could be that, as we have already suggested, the evolutionary task of approaching food elements and the learning task of predicting how the position of a food element changes with the organism's actions, are correlated. In other words, a weight matrix which is good for the first task is also good for the second task. In fact, as detailed in Nolfi, Elman, and Parisi (1990), it can be empirically demonstrated that in many cases input stimuli must be classified in the same manner for both the prediction and the approaching tasks. Hence, the same set of weights may be appropriate for both tasks. This kind of explanation does not require to postulate any inheritance of acquired characteristics. The Os of later generations learn to predict better than random Os simply because they have been selected for approaching food elements. The presence of prediction learning during the life of previous generations does not have any role in the increased ability to learn to predict of the Os of the later generations. The only role of prediction learning is in determining an increased capacity to reach food elements, as we have shown in the previous section.

This explanation assumes that there is inheritance of the ability to learn a particular task only if the learning task is correlated with the evolutionary task, that is, with the task that dictates who will reproduce and who won't. To test this hypothesis we run a new set of simulations in which Os had to learn during their life a task which presumably is not correlated⁴ with the task for which Os are selected. Our choice has been the XOR task. At each time step Os, in addition to generating an useful output on the motor output units, are taught by backpropagation to generate on an additional output unit a value of 0 if both input units have an activation value which is greater or less than .5, and a value of 1 otherwise.

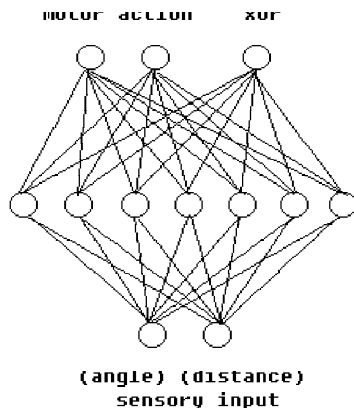


Figure 8. Network architecture for Os that learn the XOR task in addition to being selected for generating

the appropriate motor actions in response to sensory stimuli. Sensory input is encoded by the 2 input units representing the angle and the distance of the nearest food element. Movement is encoded in the 2 output units that specify the amount and direction of turn and the length of the step forward. (At each time step Os can turn from 90 degrees left to 90 degrees right and then move from 0 to 5 cells forward.) The third output unit is the response unit for the XOR task.

If the previous explanation is correct we should expect that learning the XOR task during life should not influence how the task is learned by successive generations, that is, that there is no inheritance of the ability to learn the XOR task, since this task is not correlated with the evolutionary task of reaching food. Contrary to this expectation, we found that Os of successive generations are able to learn the XOR task better and faster than the Os of previous generations (see Figure 9). In other words, a capacity to learn the XOR task is genetically transmitted even if the weight changes that result from learning are not transmitted and the ability to perform the XOR task does not correlate with the capacity for which Os are selected.




Figure 9. Global error on the XOR task as a function of epochs of life. The error curves of several generations are represented. Each curve is the average result of 10 simulations each starting from different initial random assignment of connection weights. Global error at epoch 0 is calculated by testing Os for an epoch of life without letting back-propagation operate.

Although the correlation between the learning task and the evolutionary task can be a (partial) explanation of the inheritance of the ability to learn the learning task when such correlation exists, it is clear that there can be inheritance of a learning ability even in the absence of this correlation and therefore we need another explanation for the case in which the learning task is not correlated with the evolutionary task.

A possible explanation which is consistent with our

previous explanation of the influence of learning on evolution could be the following. We defined two tasks as correlated if it is probable that, given an arbitrary point in weight space, the performance of this weight matrix on both tasks will be equally good or equally bad. However, we should expect that, even when two tasks are not globally correlated, there may exist some sub-regions of weight space in which the two tasks are more correlated than in other regions. Individuals located in these more correlated regions will be more likely to reproduce because learning (moving to a higher position on the performance surface of the learning task) would involve moving to a higher position of the fitness surface (see Figure 9). But since these reproducing individuals are located in correlated regions, an increase in the evolutionary ability (approaching food) across generations will be accompanied by a parallel increase in the ability to learn the life-time task (doing the XOR task).




Figure 10. Fitness surface for the evolutionary task and performance surface for the life-learning task for all possible weights matrices. Life-time movements due to learning are represented as arrows. Point A is in a region in which the two surfaces are correlated. As a consequence, A has more probability to be selected than B even if A and B have the same fitness on the evolutionary surface at birth, since A will be more likely to increase its performance during life than B.

We conclude that even the ability to learn arbitrary tasks can be genetically transmitted because evolution will progressively select individuals that lie in those sub-regions of weight space that correspond to correlated segments of the learning task surface and of the fitness surface.

4. The influence of behavior on evolution: Self-selection of input stimuli

We have assumed so far a notion of genetic fitness in terms of which, given a certain fitness criterion, each genetically transmitted weight matrix is assigned a fixed fitness value. In our case for each given weight

matrix there is a corresponding single value on the fitness surface which is the number of food elements eaten during life. However, this is a simplification in that what is actually observed and used by the selective reproduction mechanism is not this hypothetical genetic fitness value but a specific phenotypic fitness value. For each assumed genetic fitness value there may be various actually observed phenotypic values depending on a number of additional factors including (a) the environment in which an individual happens to live, (b) the experiences the individual happens to have in that environment, and (c), in so far as the behavior of the individual determines these experiences and in some cases changes the environment itself, the behavior of the organism.

We won't consider the role in evolution of all these additional factors that determine the phenotypic fitness of individuals on which the evolutionary process is based, but we will restrict ourselves to a particular aspect of the role of the behavior of the organism in determining the course of evolution.

Our starting point is that our organisms are ecological neural networks, i.e. neural networks that live and learn in an environment (Parisi, Cecconi, and Nolfi, 1990). The input to a network at each time step is not arbitrarily decided by the researcher but is a function of the structure of the initially defined environment and of the behavior of the organism in that environment. More specifically, the sensory input to an O (angle and distance of the nearest food element) depends on the local distribution of food but also on what has been the motor output of the network in the previous cycle. In other words, in ecological networks Os can control input stimuli with their behavior.

Now there are at least two different strategies that can be followed to maximize the number of food elements eaten. One strategy is to increase one's capacity to respond in an efficient way to all kinds of input stimuli. The other strategy is to develop a capacity to respond efficiently to a subset of input stimuli and then behave in such a way that one is more likely to encounter this subset of stimuli rather than the remaining ones. We have analyzed the data of our previous simulations to test the hypothesis that Os are able to follow the second strategy, which implies an important role of behavior in determining the phenotypic fitness of individuals.

We divided the input stimuli Os can receive during their life into 10 classes that correspond to different amplitudes of the angle of the currently perceived

food element, and we calculated the frequency with which stimuli belonging to each of the 10 classes are perceived by a particular O.




Figure 11. Percentage of occurrence for each of 10 classes of stimuli during 5000 actions of a particular O.

As Figure 11 shows, different classes of stimuli have very different frequencies of occurrence. For the particular individual that we have examined (other individuals may have different frequency distributions), stimuli with a very small angle (i.e. stimuli just on the right of O's facing direction) have very high frequency while stimuli with a very large angle (i.e. stimuli just on the left of O's facing direction) have very low frequency. At this point we can look at the O's performance for each class of stimuli. Since our Os are being selected for their ability to approach food, we defined the goodness of the performance in response to each particular stimulus as the amount of decrease in the distance between O and the stimulus (food element) after O's action.




Figure 12. Average performance of the same O for each class of stimuli.

As Figure 12 shows, O reacts in a more efficient way to stimuli with small angles than to stimuli with large angles with respect to O's facing direction. This implies that O has developed a behavior which allows

it to be exposed, most of the time, to stimuli to which it is able to react in a efficient way⁵.

We can also measure how much of the O's performance can be explained as an ability to select the most appropriate stimuli and how much as an ability to correctly react to input stimuli. The results of this analysis are shown in Figure 13. The average performance of the O in the standard situation (i.e. the situation in which the stimulus at time t depends on O's action at time t-1) is plotted against the average performance of the same O obtained by generating each time a new stimulus in a random position with respect to O.




Figure 13. Average performance of an O which can indirectly select the incoming input stimuli compared with the performance of the same O when it is positioned, at each time step, in a new arbitrary situation.

The large loss in performance obtained when O is deprived of the possibility to indirectly select the input stimuli shows how this ability can be important in explaining O's evolved behavior. This can also explain why an ability to react equally efficiently to all classes of stimuli does not emerge evolutionarily. Os will still benefit from acquiring a capacity to react efficiently to all classes of stimuli, because infrequent stimuli to which Os do not react efficiently may still appear. On the other hand, the beneficial effect of such a generalized capacity would be relatively small when compared with the more specialized capacity to react efficiently to self-selected stimuli, so that there would not be enough accumulated evolutionary pressure for the generalized capacity to emerge.

5. Conclusions

In the last few years, thanks to the large increases in available computational power, the "artificial life" experimental approach to the study of natural evolutionary phenomena has spread in the scientific community. Within this approach, neural networks

and genetic algorithms have been the most common tools used to simulate, respectively, the individual organisms and the natural evolutionary process. (In addition to the work already cited see: Miller, and Todd, 1990; Belew, McInerney, and Schraudolph, 1990).

This simulative approach has already produced many interesting results that have contributed to clarify important arguments discussed in the evolutionary biology literature. This despite the fact that the simulative models currently implemented are extremely simplified with respect to the real phenomena.

In this paper we have offered new results on the interaction between learning, behavior, and evolution and we offered a general and consistent explanation of the different findings. We have shown that life-time learned changes can have an influence on evolution although changes that are correlated with the criterion used for selective reproduction are larger than random changes. We also demonstrated how an ability to learn some task can emerge and be transmitted evolutionarily for both tasks that are correlated with the reproduction criteria and for uncorrelated tasks. Finally, we have indicated how behavior - more specifically self-selection of input stimuli - can influence evolution in that behavior is one factor that determines the observed phenotypic fitness on which selective reproduction is based.

References

- Ackley, D.E. and Littman, M.L. (1991). Proceedings of the Second Conference on Artificial Life. Addison-Wesley: Reading, MA.
- Baldwin, J.M. (1896). A new factor in evolution. American Naturalist, **30**, 441-451.
- Belew, R.K. (1989). Evolution, learning, and culture: computational metaphors for adaptive algorithms. CSE Technical Report CS89-156. University of California, San Diego.
- Belew, R.K., McInerney, J., Schraudolph, N. (1990). Evolving networks: using the genetic algorithm with connectionist learning. CSE Technical Report CS89-174. University of California, San Diego.
- Goldberg, D.E. (1989). Genetic Algorithms in Search, Optimization, and Machine Learning. New York: Addison-Wesley.
- Hinton, G.E., Nowlan S.J. (1987). How Learning Guides Evolution. Complex System, **1**, 495-502.
- Holland, J.J. (1975). Adaptation in Natural and Artificial Systems. Ann Arbor, Michigan: University of Michigan Press.
- Miller, G.F. and Todd, P.M. (1990) Exploring adaptive agency I: theory and methods for simulating the evolution of learning. In D.S. Touretzky, J.L. Elman, T.J. Sejnowski and G.E. Hinton (eds.), Proceedings of the 1990 Connectionist Models Summer School. San Matteo, CA: Morgan Kaufmann.
- Menczer, F. and Parisi, D. (1990). "Sexual" reproduction in neural networks. Technical Report PCIA-90-06. Institute of Psychology, C.N.R. Rome.
- Menczer, F. and Parisi, D. (1991). Evidence of hyperplanes in the genetic learning of neural networks. Technical Report PCIA-91-08. Institute of Psychology, C.N.R. Rome.
- Nolfi, S., Elman, J, and Parisi, D. (1990). Learning and evolution in neural networks. CRL Technical Report 9019. University of California, San Diego.
- Nolfi, S., and Parisi, D. (1991). Auto-teaching: neural networks that develop their own teaching input. Technical Report PCIA-91-03. Institute of Psychology, C.N.R. Rome.
- Parisi, D., Cecconi, F., Nolfi, S. (1990). Econets: Neural Networks that Learn in an Environment. Network, **1**, 149-168.
- Plotkin, H.C. (1988). The role of behavior in evolution. Cambridge, Mass.: MIT Press.
- Rumelhart, D.E., Hinton G.E., and Williams, R.J. (1986). Learning internal representations by error propagation. In D.E. Rumelhart, and J.L. McClelland, (eds.), Parallel Distributed Processing. Vol.1: Foundations. Cambridge, Mass.: MIT Press.
- Rumelhart, D.E. and McClelland, J.L. (1986). Parallel Distributed Processing. Cambridge, Mass.: MIT Press.
- Todd, P.M. and Miller, G.F. (1991) Exploring adaptive agency II: simulating the evolution of associative learning. In: J.A. Meyer and S.W. Wilson (eds), From Animals to Animats. Cambridge, MA: MIT Press/Bradford Books.

Waddington, C.H. (1942). Canalization of development and the inheritance of acquired characters. *Nature*, **150**, 563-565.

¹ We will see later in the paper that talking of the fitness of genotypes is inaccurate and it ignores an important aspect of the processes involved in evolution. However, for the moment we can be satisfied with what has been said.

² This is not true in the case of sexual reproduction and genetic recombination because in this case offsprings can have a very different weight matrix from that of each of their parents. On sexual reproduction in neural networks, see Menczer and Parisi 1990;1991).

³ A similar explanation has been given by Hinton and Nowlan (1987).

⁴ To verify if two tasks are correlated one should calculate the performance on the two tasks for each point in weight space. Instead, we have used our intuitive judgment.

⁵ Other Os develop different preferences in stimulus selection (for example left stimuli can be preferred to right ones) but all react better to stimuli they indirectly select more often.