

The ability to select favourable stimuli and the ability to react to input stimuli efficiently, at least in our simple approaching task, interact in a competitive fashion during the development of the overall ability to perform the task. The first ability plays a very important role at the beginning of the developmental phase while the second ability emerges later in development. In other words, normally Ss first acquire the capacity to behave in such a way that they encounter more frequently a certain restricted class of stimuli to which they know how to react and only later on they extend their ability to react efficiently to other less frequent stimuli.

Finally, it seems that the particular interaction between these two abilities which is actually observed can predict the probability that a particular behavior strategy will represent a stable strategy or a local minimum from a developmental point of view.

Acknowledgements

This research was supported by P.F. "ROBOTICA", C.N.R., Italy.

References

Ackley, D.H., and Littman, M.,S. (1990). Learning from natural selection in an artificial environment. In *Proceedings of the International Joint Conference on Neural Networks* (pp. 189-193). Hillsdale,NJ:Erlbaum.

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.

Holland, J.J. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor, Michigan: University of Michigan Press.

Hopfield, J.J. (1982). Neural Networks and Physical Systems with Emergent Collective Computational Abilities. *Proceedings of the National Academy of Sciences, U.S.A.*, **79**, 2554-2558.

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). Growing Neural Networks. *Technical Report*, PCIA-91-18 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.

increase in overall reacting ability.




Figure 16: Frequency of input classes across generations in Simulation 3




Figure 17: Performance on input classes across generations in Simulation 3

4. Conclusion

A system which behaves in an environment can increase its performance level in two different ways. It can improve its ability to react efficiently to any stimulus that may come from the environment or it can acquire an ability to expose itself only to a sub-class of stimuli to which it knows how to respond efficiently. The possibility that a system can solve a task by selecting favourable stimuli is rarely considered in designing intelligent systems. In this paper we have shown that this type of ability can play a very powerful role in explaining a system's performance. Moreover, artificial systems which are left free to develop their own way of solving a problem, may exhibit a strong preference to found their behavior strategy on this stimulus selection ability.




Figure 14: Frequency of input classes across generations in Simulation 2




Figure 15: Performance on input classes across generations in Simulation 2

At generation 60 Ss' behavior reaches a stable state. The small angle stimuli (classes 1 to 4) are the only one to which Ss are able to react efficiently. On the other hand, Ss are very efficient in response to these stimuli, which are very frequent. One consequence of this is that Ss are unable to find a new behavior strategy which would allow them to improve their ability to react to other input stimuli efficiently and at the same time to preserve their self-selection ability. In other words, if we examine the particular behavior strategy which has evolved at around generation 60, we can understand why changes resulting in additional increases in performance do not arise in successive generations.

Simulation 3 shows a more complex pattern (see Figures 16 and 17). After a few generations both class 10 stimuli (that correspond to stimuli just on the left of an S's facing direction) and, to a less degree, class 1 stimuli (that correspond to stimuli just on the right of an S's facing direction) become frequent. This is a reasonable solution because both classes of stimuli are stimuli approximately in front of an S. However, performance with respect to these self-selected stimuli remains rather poor. As a consequence, at around generation 120 a new global rearrangement occurs leading to a different distribution of class frequencies and to a large

frequencies. However, performance, even performance in reaction to class 1 stimuli, is not particularly good at this stage of the evolutionary process. It is necessary to reduce the excessive frequency of class 1 stimuli to obtain better performance levels to this class of stimuli and to stimuli of nearby classes. Apparently, an excessive initial specialization leads to poor performance; hence, more balanced organisms tend to evolve.




Figure 12: Frequency of input classes across generations in Simulation 1




Figure 13: Performance on input classes across generations in Simulation 1

Simulation 2 shows a quite similar pattern (see Figures 14 and 15). In the first generations class 1 stimuli become very frequent in comparison with all other stimuli and performance increases only when the frequency of class 1 stimuli is drastically reduced. The capacity to select favourable stimuli and the ability to react to all stimuli tend to compete, i.e. an increase in one of them corresponds to a decrease in the other. Changes of behavior resulting in an increase of one ability without perturbing the other are very rare although they sometimes do occur. One such change occurs around generation 50 when an increase of the performance in reaction to many low frequency stimuli is obtained without modifying the frequency distribution of the various classes.




Figure 10: Performance increase across generations in Simulation 3




Figure 11: Performance due to the ability to select favourable input stimuli and to the ability to react to input stimuli in Simulation 3.

Another analysis that can be conducted on our results is to examine the frequency (percentage of occurrence) of the various classes of stimuli (angle of food) and to determine how well Ss perform, in terms of reducing their distance from food, for each of these different classes. Figure 12 and 13 show the result of this analysis for Simulation 1. (Label 1 identifies stimuli with an angle between 0 and 36 degrees, label 2 stimuli with an angle between 36 and 72 degrees, etc.) As can be seen, very soon class 1 stimuli, i.e. food elements a little on the right with respect to an S's facing direction, become the most frequent class with all other stimuli kept at very low




Figure 9: Performance due to the ability to select favourable input stimuli and to the ability to react to input stimuli in Simulation 2.

As in Simulation 1, the increase in overall performance in the first generations is mainly due to an increase of the capacity to select favourable stimuli while the increase from generation 50 to around generation 100 is mainly due to an increase in the ability to react efficiently to all stimuli. In addition, we can see again that the two abilities compete until a stable state is reached around generation 75.

In the third and last simulation a more complex pattern was obtained (see Figures 10 and 11). The shape of the curve for the overall performance (Figure 10) indicates a slower increase than in Simulation 1 and 2 for most of the evolutionary process followed by a terminal phase of rapid increase leading to a final level which is better than that of the other two simulations. Figure 11 tells us that, as in Simulations 1 and 2, the increase in performance in the very first generations is caused by a development of the self-selection capacity but very soon (around generation 15) a stable state is reached in which both components of the overall performance have a comparable role in determining the performance. In this phase of the evolutionary process (from around generation 10 to generations 115-120) there appears to be no competition between the two components and we observe a very slow increase in overall performance. After generation 120, however, competition is restored and, interestingly, in this simulation the component "ability to approach any food" ends up being a more important factor in explaining overall performance than the other component of selecting a restricted set of more favourable stimuli. This terminal phase of renewed competition between the two components in the use of Ss' computational resources might be the cause of the final high level of overall performance in this simulation.

A comparison between the evolution across generations of the approaching ability and the selecting ability indicates that the role of selection becomes soon very significant while approaching takes more time (generations) to reach an importance comparable to the first component. In other words, Ss first rely very much on their capacity to behave in a way that selects a restricted class of stimuli as input to which they know how to react and then they concentrate on acquiring a more general ability to approach food stimulating them from any angle and distance. Another observation evoked by Figure 7 is that there appears to be some sort of competitive interaction between the two components of performance. Very often to an increase of the self-selection capacity corresponds a decrease in the ability of Ss to approach food whatever the input angle and distance. This process continues until generation 100 is reached after which both components of performance tend to remain in a stable state.

We have done the same analysis for Simulation 2 which is identical to Simulation 1 except that a different random seed for generating the initial set of weight matrices and architectures is used. Figure 8 presents the increase in performance across the 150 generations and Figure 9 shows how much of this increase is due to the capacity to self-select the input and how much to the ability to reach food whatever the input.




Figure 8: Performance increase across generations in Simulation 2

subtracting the "approaching" curve from the "overall performance" curve of Figure 6. This second curve is labeled "selecting". As the comparison between these two curves shows, our Ss had a better performance, in terms of approaching food, when they were allowed to self-select their input (as it was the case in the actual evolution) than when they were tested with externally imposed input which sampled all the possible stimuli. Hence, by subtracting from the overall performance the performance which was due to their ability to approach all kinds of stimuli we thought we had a measure of the contribution of their capacity to self-select stimuli to the overall performance.




Figure 6: Performance increase across generations in Simulation 1




Figure 7: Performance due to the ability to select favourable input stimuli and to the ability to react to input stimuli in Simulation 1.

At this point we can look at S's performance for each class of stimuli. We defined the performance for each particular stimulus as the amount of decrease (scaled between 0 and 1) of the distance between S and the stimulus after S's action. If the system shows better performance on the more frequent stimuli we can hypothesize that the frequency of the stimuli is manipulated by the system itself in order to increase its performance.

As Figure 5 shows, this is what happens. The figure describes the performance of the same individual of Figure 4 for each class of stimuli. S reacts in a more efficient way to stimuli with small angle than to stimuli with large angle with respect to S's facing direction. And, as shown in Figure 4, stimuli with small angle are much more frequent than stimuli with large angle as a consequence of S's behavior. This implies that S 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 more efficient way .




Figure 5: Performance with respect to different classes of input stimuli

3. An analysis of performance during the course of evolution

A system which behaves in an environment, like our Ss, has two different ways of improving its performance: (a) by acquiring a capacity to react better to input stimuli; and (b) by acquiring a capacity to behave in a way that increases the frequency of stimuli to which it is able to react efficiently. If we look at how our Ss actually improve their performance we can conclude that both ways are pursued by them. In this Section we will examine how the overall performance of our Ss, and these two components separately, evolve across generations.

Figure 6 shows how the overall performance of the best individual of each generation, measured as the average reduction of the distance between the individual and food after the individual's action, improves in the 150 generations for Simulation 1. In order to have a separate measure of the capacity to react to input stimuli, whatever the stimuli, we positioned each of these individuals in a different random position for 5000 times and we calculated the average reduction of the distance between S and food after S's action. In other words, in this test an S was not allowed to select its own input stimuli with its actions but was exposed to the complete range of possible stimuli (with the restriction of a sample of 5000 stimuli) and its performance with respect to all possible stimuli was measured. The results are shown in Figure 7 in the curve labeled "approaching". The figure also shows a second curve which is obtained by simply




Figure 3: Fitness increase across generations in Simulation 1

At the end of the training process we can examine Ss' behavior and analyze the particular solution found by each S. We are interested in testing the hypothesis that Ss increase their performance by learning to select the input stimuli in a useful way in addition to learning to react to stimuli in an appropriate way. In order to verify this hypothesis we divided the input stimuli 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 class were perceived by S.

The results for the best S in the last generation of Simulation 1 are shown in Figure 4 which gives the percentage of occurrence of each class of stimuli during 5000 actions. (Class 1 identifies stimuli with angle between 0 and 36, class 2 stimuli with angle between 36 and 72, etc.).




Figure 4: Frequency of different classes of input stimuli

As the figure shows, different classes of stimuli have very different frequencies of occurrence. In particular, stimuli with a very small angle (i.e. stimuli just on the right of S's facing direction) have very high frequency while stimuli with a very large angle (i.e. stimuli just on the left of S's facing direction) have very low frequency.

When S is placed in the environment, a sequence of events will occur. Sensory input is received on the input units. Activation flows up through the hidden units to the output units. The values on the output units are then used to move S in the manner specified by these values, thereby changing the sensory input for the next cycle.

To train the network (i.e. to optimize the free parameters of S) we used a kind of genetic algorithm (Holland, 1975) applied to Ss' "genotypes". Each genotype represents a set of instructions for building the corresponding S's neural system (for more details see Nolfi and Parisi, 1991). We begin with 100 different random genotypes each yielding a network with a different architecture and a different assignment of connection weights. This is Generation 0 (G0). G0 networks are allowed to "live" for 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 40x40 cells with 10 pieces of food randomly distributed in it. The Ss are placed in individual copies of these environments, i.e. they live in isolation.

At the end of their life (5000 actions) Ss are allowed to reproduce. However only the 20 Ss 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 20x5=100 new Ss constitute the next generation (G1). Mutations are introduced in the copying process resulting in possible changes of the architecture or of the weights' value.

After Ss of G1 are created they are allowed to live for 5000 cycles. The behavior of these Ss differs slightly from that of preceding generation (G0) as a result of two factors. First, the 100 Ss of G1 are the offspring (copies) of a subset of the Ss of G0. Second, the offspring themselves differ slightly from their parents because of the mutations. These differences lead to small differences in mean food eaten by the Ss in G1. At the end of their life the 20 best individuals are allowed to reproduce 5 times, forming G2. This process continues for 150 generations.

This method, like other unsupervised learning algorithms (Hopfield, 1982), allows the researcher to avoid specifying the correct solution for the task. The solution emerges through an optimization process during which Ss are evaluated simply on the basis of the total number of food elements eaten.

Simulations

We ran 3 simulations starting with different initial populations of S. If we look at Ss' fitness (i.e. number of food elements eaten) across generations we can see that Ss more able to approach food elements tend to evolve. Figure 3 shows the fitness value for the best individual of each generation for Simulation 1.




Figure 1: System and its Environment

At any particular moment *S* occupies one of the cells. A number of food elements are randomly distributed in the environment with each food element occupying a single cell. *S* has a facing direction and has a rudimentary sensory system that allows it to receive as input a coding of the angle (relative to where *S* is currently facing) and distance of the nearest food element. *S* is also equipped with a simple motor system that provides it with the possibility, in a single action, to turn any angle from 90 degrees left to 90 degrees right and then move from 0 to 5 cells forward. Finally, when *S* happens to step on a food cell, it eats the food element which disappears.

Ss are implemented as feed-forward neural network such as that of Figure 2. Sensory input is encoded by 2 input units representing the angle and the distance of the nearest food element (both values are scaled from 0.0 to 1.0). Motor action is encoded in 2 output units that specify the amount and direction of turn and the length of the step forward (these two values are also scaled from 0.0 to 1.0).




Figure 2: Each system is implemented as a neural network

Self-selection of Input Stimuli for Improving Performance

Stefano Nolfi **Domenico Parisi**

Institute of Psychology, CNR

V.le Marx 15, 00137 Rome - Italy

E-mail: stiva@irmkant.Bitnet domenico@irmkant.Bitnet

Abstract

A system which behaves in an environment can increase its performance level in two different ways. It can improve its ability to react efficiently to any stimulus that may come from the environment or it can acquire an ability to expose itself only to a sub-class of stimuli to which it knows how to respond efficiently. The possibility that a system can solve a task by selecting favourable stimuli is rarely considered in designing intelligent systems. In this paper we show that this type of ability can play a very powerful role in explaining a system's performance.

Introduction

Cognitive science and robotics have been dominated for years by a static conception of intelligence. More recently there has been a growing awareness of how important active motor behavior, for example in object manipulation, is for understanding natural intelligence and for developing really useful artificial intelligence.

In this paper we analyze the capacity of a system to self-select with its behavior its own stimuli. Systems that behave in an environment have the possibility to determine, on the basis of their behavior, the stimuli they receive as input. This possibility can be used to learn to predict the sensory consequences of the system's own behavior and to develop a map of the environment which will facilitate the system's ability to attain goals in the environment (Nolfi, Elman and Parisi, 1990; Parisi, Cecconi, and Nolfi, 1990). Another use of the possibility to determine the stimulus input is to select more favourable stimuli. In this paper we will examine this second aspect and, more particularly, we will show how by self-selecting more favourable stimuli a system can increase its goal attaining ability.

The problem

We developed a system (S) that must perform a very simple navigation task in order to approach "objects" randomly distributed in a "world". One can think of our system as a very simple animal that must find and eat food elements distributed in its environment. The S's environment is a two-dimensional square divided up into cells (see Figure 1).