

Evolving Robots Able to Visually Discriminate Between Objects with Different Size

S Nolfi^{*}, D Marocco^{**}

^{*}Institute of Cognitive Science and Technology, National Research Council (CNR)
Viale Marx, 15, 00137, Rome, Italy
nolfi@ip.rm.cnr.it

^{**}University of Calabria, "Centro Interdipartimentale della Comunicazione"
87036 Arcavacata di Rende, Cosenza, Italy
davidem@ip.rm.cnr.it

Abstract

In this paper we present a set of experiments in which mobile robots able to visually discriminate between two different objects were developed through artificial evolution. The results obtained by evolving robots in stationary environment showed that evolving individuals, by exploiting sensory-motor coordination, develop an extremely simple and effective solution. The analysis of the results obtained by evolving robots in varying environmental conditions showed that evolving individuals develop a slightly more complex solutions that also require an ability to modify the behavior on the fly by adapting to the current environmental circumstances. Overall, the obtained results show that artificial evolution might lead to the development of simple and effective solutions also in the case of problem that has to be solved in varying environmental conditions.

Keywords: Evolutionary Robotics, Sensory-Motor Coordination, Active Vision

1. Introduction

Evolutionary Robotics is a new technique for the automatic creation of autonomous robots that is inspired by the Darwinian principle of selective reproduction of the fittest. It is a new approach which looks at robots as autonomous artificial organisms that develop their own skills in close interaction with the environment without human intervention. The basic idea behind evolutionary robotics is very simple. An initial population of different artificial genotypes, each encoding the control system and sometimes the morphology of a robot, are randomly created and put in the environment. Each robot is then let free to act (move, look around, manipulate) according to a genetically specified controller while its performance on a given task is automatically evaluated. The fittest robots are allowed to reproduce by generating copies of their genotypes with the addition of changes introduced by some genetic operators (e.g., mutations, crossover, duplication). This process is repeated for a number of generations until an individual is born which satisfies the performance criterion (fitness function) set by the experimenter. Robots might also change during lifetime and therefore might adapt to their environment both phylogenetically and ontogenetically.

There is a growing body of literature describing successful applications of evolutionary methods [1-3] to the development of mobile robots [4]. The impact of these researches on our understanding of adaptive behavior in natural agents and on our ability to develop effective artificial agents (robots), however, depends on the extent to which this method can scale-up to tackle richer robot/environmental interactions.

A large number of animal species have exploited visual receptors. On the other hand, the majority of the research projects in evolutionary robotics involve robots provided only with short distance sensors (e.g. infrared sensors or bumpers) that provide information of the immediate surrounding of the robot (for a notable exception see [5-6]).

In this paper, we present a set of experiments in which evolving robots are asked to visually discriminate between objects with the same shape but with different sizes by navigating toward one of the two. In section 2 and 3 we will present the result of a set of experiments conducted by evolving robots in stationary and varying environmental conditions, respectively. In section 4 we will report our conclusions.

2. The basic experiment

The mobile robot Khepera [7] equipped with a vision module was placed in the environment shown in Fig. 1. The environment is a rectangular arena 58 x 37 cm surrounded by white walls with two black stripes with a width of 4 and 16 cm respectively. Walls are made of white painted wood and are 30 cm in height. A neon lamp (24 W output) provided illumination over the arena. The robot has to approach the bigger (B) but not the smaller (S) stripe. In other words, it has to visually discriminate the two objects.

One interesting aspect of this task is that often the robot receives the same sensory patterns while it is facing both B and S. In particular B at a certain distance looks exactly like S at a closer distance (this is a problem that tends to arise in any real world situation). Therefore, to solve its task, the robot has to overcome the problem that often sensory states are ambiguous.




Figure 1. The environment and the Khepera robot with the linear vision turret. The environment consist of an arena of 58 x 37 cm that contains two visual targets that consist of a bigger (B) and a smaller (S) stripe of 16 and 4 cm respectively.

The robot is supported by two lateral wheels that can rotate in both directions and two rigid pivots in the front and in the back. The sensory system employs eight infrared sensors distributed around one side of the body and a linear camera. The infrared sensors can detect the walls at a maximum distance of approximately 4 cm. The linear camera consists of one array of 64 photoreceptors providing a linear image composed of 64 pixels of 256 gray-levels each and subtending a total view-angle of 36°. The optics are designed to bring into focus objects at distances between 5cm and 50cm while an additional sensor of light intensity provides information for automatic adjustment of photoreceptor sensitivity. The robot was attached to the host computer by means of a lightweight aerial cable and specially designed rotating contacts. This configuration makes it possible to trace and record all important variables by exploiting the storage capabilities of the host computer, and at the same time

provides electrical power without using time-consuming homing algorithms or large heavy-duty batteries.

Each individual of the population was tested on the robot for 30 epochs, each epoch consisting of a maximum of 300 sensory-motor cycles (each cycle last 100 ms). An epoch ended either when the robot reached a distance lower than 10 cm from the center of B or S or after 300 sensory-motor cycles (corresponding to 30 seconds). At the beginning of each epoch the robot is positioned at a random position and orientation. Individuals' fitness was increased of 1 point each time the robot approached B (i.e. each time it reached a distance lower than 10 cm from the center of B) and decreased of 1 point each time the robot approached S.

The controller (Fig. 2) consisted of a simple feed-forward neural network with 14 sensory neurons encoding the state of 6 frontal infrared sensors and 8 photoreceptors of the camera (i.e. only the 6 frontal infrared sensors and only 8 evenly spaced photoreceptors out of the 64 provided by the camera were used) and two motor neurons encoding the speed of the two wheels. The activation state of the infrared sensors was normalized between 0.0 and 1.0. The activation of the two motor neurons was computed by using the sigmoid function and then normalized between -10 and +10 to set the desired speed of the two corresponding motors.




Figure 2. The neural controller.

In order to keep things as simple as possible and given the small size of the parameter set, we used a direct genetic encoding [8] and a fixed neural architecture. Each parameter (i.e. the connection weights and threshold values of the two motor neurons) was encoded using 8 bits and normalized between -10.0 and +10.0. Therefore, the genotype of each individual robot was $8 \times (28 \text{ synapses} + 2 \text{ thresholds})$ bits long. Population size was 100. The best 20 individuals of each generation were allowed to reproduce by generating 5 copies of their genotype with 3% of their bits replaced with a new randomly selected value.

To save time, the evolutionary process was carried out in simulation by using an extended version [9] of the software Evorobot described in [10]. The results of the experiment described in this section have also been validated by downloading the evolved controllers into the real robot (see below).

We ran 10 replication of the evolutionary process in simulation by starting with different randomly generated genotypes. Each replication of the experiment took about 1 hour by using a standard PC and would have taken up to 52 days on the physical robot. As can be seen from Fig. 3, evolving individuals reached optimal performance after about 10 generations.

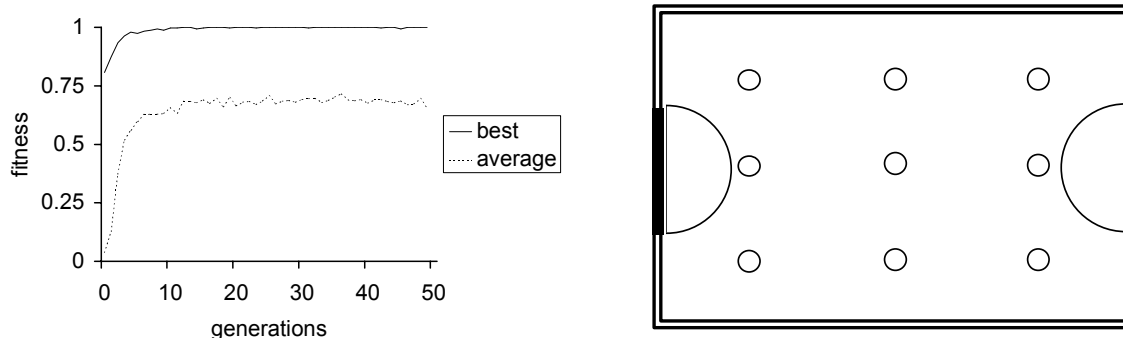


Figure 3. **Left.** Fitness through out generations normalized between -1.0 and +1.0 (i.e. 1.0 means that the robot was able to reach the correct visual target in 30 out of 30 trials). The dashed and full lines represent the average performance of the population and the performance of the best individual respectively. Average results of 10 replications. **Right.** A schematic representation of the environment. The double lines represent the walls painted in white. The two thick lines represent the portion of the walls painted in black. The left and right semi-circles represent the area in which the robot was rewarded or punished respectively. The circles indicate the 9 positions in which evolved robots were initially placed to evaluate their performance (see text).

By initially placing the best evolved individuals of the last generation in the 9 starting positions indicated in the right side of Fig. 3 with four different orientations (north, south, east, and west) we observed that they navigate toward the right landmark most of the times. In some cases, however, they navigate toward the wrong landmark or do not reach one of the two targets within the allotted time (300 sensory-motor cycles corresponding to 30 seconds). Performance, on the average, are shown in Fig. 4. Similar but slightly worse results (see Fig. 4) have been observed in the real environment by downloading the evolved controllers on the physical robot¹ and by testing the individuals in the environment shown in Fig. 1. Also in this case, individuals were placed initially in the same $9 \times 4 = 36$ initial positions and orientations described above. One of the evolved controllers, out of 10 replications, displayed optimal performance on the physical robot by navigating toward the right landmark in 36 out of 36 trials.

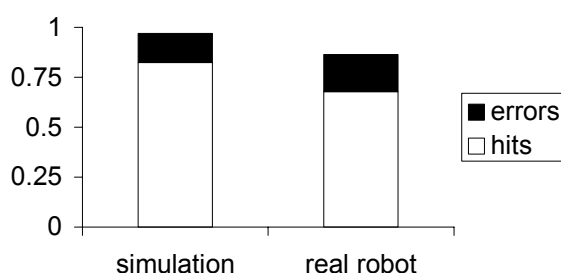


Figure 4. Percentage of hits (i.e. number of navigations toward the right visual landmark) and errors (i.e. number of navigations toward the wrong visual landmark) for the best individuals of generation 49. Performance in the real environment have been obtained by downloading and testing the controllers evolved in simulation into the real robot. Average results of 10 replications.

Interestingly, evolved individuals solve the problem by exploiting sensory-motor coordination [11-12]. It should be noted that to navigate toward B, robots should not only take into account the perceived size of the visual target but also its relative distance. Given that robots are not provided with binocular vision the only way to compute the actual distance of an object is to observe how the perceived image vary while the robot is moving. Indeed evolving robots act

¹ Given that the refresh of the camera requires more than 100 ms we stopped the robot for 500 ms after each sensory-motor cycle (i.e. after each 100 ms) during the test.

so to lose visual contact after a period of time that is proportional with the distance and the perceived size of the object. This in turn allows evolved robots to move more toward B than toward S independently of the relative size of the two objects. In other words, robots determine which is the target object by exploiting sensory-motor coordination.

This can be illustrated by analyzing the behavior of the individual shown in Fig. 5 that represents the typical strategy adopted by evolved individuals. The robot rotates clockwise on the spot until it starts to perceive one of the two objects with the camera (in this case, the robot happens to face B first). At this point it starts to move forward by slightly turning clockwise. This allows the robot to lose visual contact with the object after some time. At this point the robot starts to rotate clockwise again and then to move toward the other object for a while by slightly turning clockwise. The robot thus moves back and forth between the two objects. However, given that it loses visual contact more slowly with the larger and more distant than with smaller and closer objects, it tends to move more toward B than toward S (see the sub-parts of the trajectory indicated with empty and full arrows respectively). This in turn allows the robot to finally reach B after a certain number of back and forth segments (where each segment consists in a certain number of forward movements toward one of the two objects).

To explain why the robot loses the visual contact more slowly with larger than smaller objects (from the robot position point of view) one should consider that the former persists longer than the latter in the view angle of the camera while the robot move forward by slightly turning clockwise. To explain why the robot loses the visual contact more slowly with more distant than closer object (from the robot point of view) one should consider that, while approaching an object, its angular projection varies more slowly for more distant than for closer objects. In other words, the robot uses motion not only to approach the desired target but also to implicitly determine which is the object with the larger size (independently from the robot's point of view).




Figure 5. The behavior of a typical evolved individual. The empty circle represents the robot and the curve line represents its trajectory during a trial. In this example the robot is initially placed in the north-east side of the arena facing south and navigate toward B that is located on the west side of the arena.

In a control experiment in which we asked evolving robots to approach S and avoid B (i.e. approach the smaller and avoid the bigger stripe) we obtained similar results and we observed the same type of behavioral strategy. Also in this case evolved individuals moves back and forth between the two objects by losing visual contact more slowly with larger and more distant than with smaller and closer objects. The only difference is that in this case evolved individuals moved backward.

3. Evolving robots in varying environmental conditions

The strategy described in the previous section is simple and effective. However, one might argue that such simple strategy is effective in the particular environmental conditions in which the robots has been evolved but might not generalize well in varying environmental conditions. To ascertain this point and to verify if evolving robot might develop an ability to discriminate the two objects even in varying environmental conditions we ran a new set of experiments.

The robot was placed in a rectangular arena of 160 x 140 cm surrounded by white walls with two black cylinders located in the east and west side of the arena (Fig. 6). As in the previous experiment, the robot has to approach the bigger (B) but not the smaller (S) cylinder. In this case, however, the radius of the small and large cylinders were randomly selected each trial between 5 and 125 cm and between 15 and 185 cm respectively. Lifetime consisted of 40 epochs each consisting of a maximum of 200 sensory-motor cycles (each cycle lasted 100 ms). To stress the problem of disambiguating identical sensory pattern corresponding to the two different objects, in each trial, the robots were initially placed with a random selected orientation along the points from which the two visual objects have the exactly the same perceived size. All other parameters were identical to those described in the previous section.




Figure 6. The environment. The lines represent the white walls. The full circles represent two black cylinders of varying size (i.e. size is randomly selected each trial). The plain and dashed circles represent the areas in which the robot is rewarded and punished respectively.

We ran two set of experiments. In the former individuals were provided with a simple reactive controllers consisting of 14 sensory neurons and 2 motor neurons shown in Fig. 2. By testing the evolved individuals for 40 trials in changing environmental conditions (i.e. environments in which the size of the two cylinders were randomly varied each trial) we observed that the best individuals produced up to 77% of correct navigations in the case of the best replication and 59% on the average (see left of Fig. 7, thin and thick lines respectively).

In the latter experiment neural controllers were provided with two additional internal neurons receiving connections from themselves and from the sensory neurons and projecting connections to the two motor neurons (see top-left of Fig. 9). The activation state of sensory and internal neurons were updated according to the following equations (motor neurons were updated according to the logistic function):

$$A_j = t_j + \sum w_{ij} O_i \quad O_j = \tau_j O_j^{(t-1)} + (1 - \tau_j) \left(1 + e^{-A_j}\right)^{-1} \quad 0 \leq \tau_j \leq 1 \quad (1)$$

With A_j being the activity of the j th neuron (or the state of the corresponding sensor in the case of sensory neurons), t_j the bias of the j th neuron, W_{ij} the weight from the i th to the j th

neuron, O_i the output of the i th neuron. O_j is the output of the j th neuron, τ_j the time constant of the j th neuron. The time constant parameter of each neuron is genetically encoded with 8 bits and normalized between 0.0 and 1.0. All other parameters were the same of the experiment described above. By testing the evolved individuals for 40 trials in randomly varying environmental conditions we observed that in this case the best individuals produced up to 97% of correct navigations in the case of the best replication and 67% on the average (see right of Fig. 7, thin and thick lines respectively).




Figure 7. Fitness through out generations normalized between -1.0 and +1.0 (i.e. 1.0 means that the robots were able to reach the correct visual target in 40 out of 40 trials). Performance of the best individual of each generation tested for 40 trials from new 40 randomly selected positions. Thick and thin lines represent the average performance of 10 replications of the experiment and the performance in the case of the best replication, respectively. **Left:** experiment in which robots are provided with the reactive neural controllers shown in Fig. 2. **Right:** experiment in which robots are provided with the neural controllers shown in the top-left of Fig. 9.

By analyzing the behavior displayed by the individuals of the former experiment (i.e. the individuals which were provided with the neural controller shown in Fig. 2) we observed the emergence of the same strategy displayed by robots evolved in the stationary environment described in the previous section. This can be illustrated by analyzing the behavior of the individual shown in Fig. 8 that shows the behavior of a typical successful individual. The robot rotates clockwise on the spot until it starts to perceive one of the two objects with the camera (in this case, the robot happens to face S first) and then move forward by slightly turning clockwise. This allows the robot to loose the visual contact with the object after some time. At this point the robot starts to rotate clockwise again and then to move toward the other object for a while by slightly turning clockwise. The final picture is that, also in this case, the robot tends to reach B after a certain number of back and fourth movements in which it moves more toward B than toward S. As in the case of the previous experiments, evolved individuals exploit sensory-motor coordination to identify the target object and to navigate toward it.

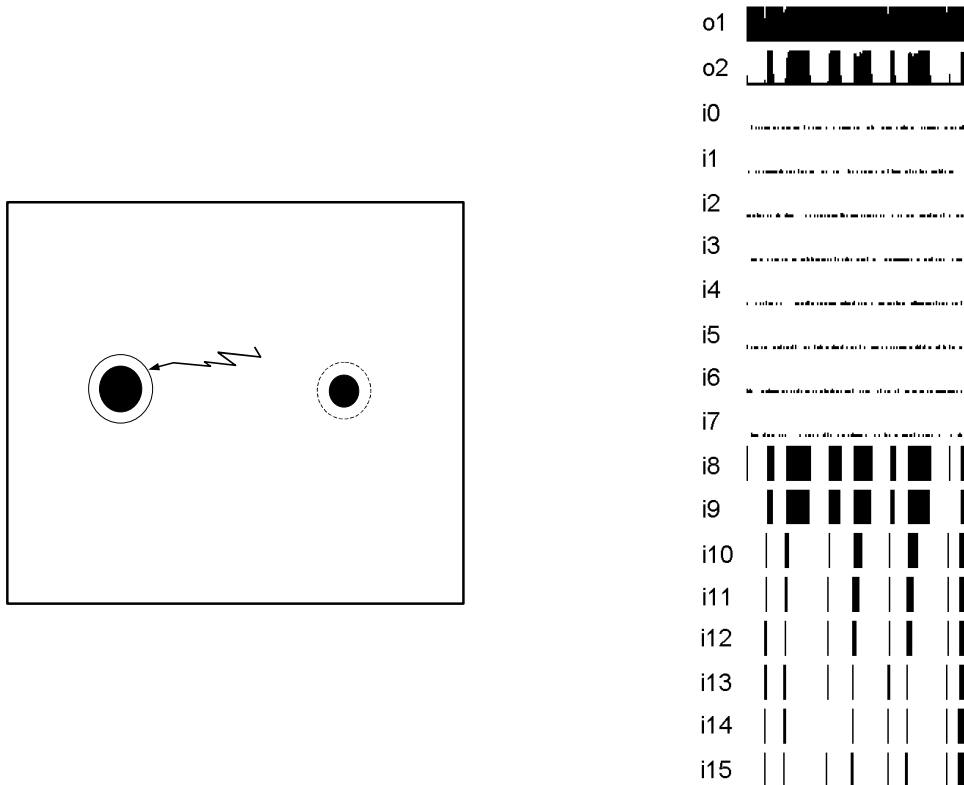


Figure 8. **Left:** typical behavior displayed by an evolved robot during a trial. The sizes of the two objects have been randomly generated within the range described in the text. **Right:** state of the two motor neurons (o1 and o2) and of 16 sensory neurons (from i0 to i15) throughout time of the individual displayed on the left side of the Fig..

The problem of this strategy that explains while it does not allow robots to reach optimal performance when the sizes of the two objects vary is that is only effective if the range of variation of the objects is not too large. Indeed, if the robot loses visual contact with the objects too quickly (i.e., if it turns clockwise significantly while approaching an object), it will move more toward B than S overall but it might be unable to reach B within the 200 lifecycles given that it will cover a small distance in each segment (i.e. in each sequence of straight movements toward the current object). On the contrary, if the robot loses visual contact too slowly, it might risk to reach the wrong object during the first approaching segment without the possibility to "measure" the relative size of the two objects by moving toward both objects at least one time. Given that how quickly the robot loses visual contact depends on the absolute size of the objects, the propensity to loose visual contact should be carefully selected by taking into account the size of the objects themselves. Individuals display a good generalization ability. However, when the range of variation is too high (as in the case of this experiment) evolved individuals might be unable to navigate toward B when objects are very large and/or very small.

By analyzing the behavior displayed by the individuals of the latter experiment (i.e. the individuals which were provided with the neural controllers shown on the top-left of Fig. 9) we observed the emergence of a new strategy that is effective also in varying environmental conditions. Indeed evolved individuals, in this case, do not keep moving back and fourth between the two objects but, after a while, tend to move consistently toward B (see the trajectory indicated with a full line, bottom-left of Fig. 9, that displays the behavior exhibited by a typical evolved individual). As in the case of the reactive robot described in Fig. 8, this robot rotates clockwise until one of the two objects (S in this example) enter in the view angle

of the camera. At this point it starts to move toward the current object by slightly turning clockwise so to loose the visual contact with the object after some time. Later on, however, when the robot starts to perceive (B), it consistently move toward it by finally reaching the right object. This individuals moves toward S only once before moving and reaching B during the second segment. Others individuals might move back and fourth a couple of times. But all successful individuals consistently move toward (B) after some time.

At a first sight, the analysis of the behavior displayed by evolved individuals (e.g. the individual shown in Fig. 9) might suggest that in this case evolved robots determine which is (B) during the first part of the trial and then move consistently toward it. In other words, one might conclude that in this case evolved individuals develop a completely different strategy that extract the relevant information from the environment and then use such information to navigate toward the target. However, a deeper analysis demonstrated that evolved individuals exploit the same sensory-motor strategy described above. The difference is that, in this case, evolved individuals are able to modulate such a strategy on the basis of the current environmental conditions.

The fact that the sensory-motor strategy described above still play an important role in these individuals can be demonstrated by lesioning the internal neurons. In the case of the individual described in Fig. 9, for instance, by freezing the activation state of the two internal neurons to 0.0 (see the right part of Fig. 9), the individual displays exactly the same strategy displayed by the reactive robots described above. In this condition, in fact, the individual approach B by moving back and forth between B and S and by loosing visual contact more quickly with S than with B (see the dashed line in the bottom-left part of Fig. 9).

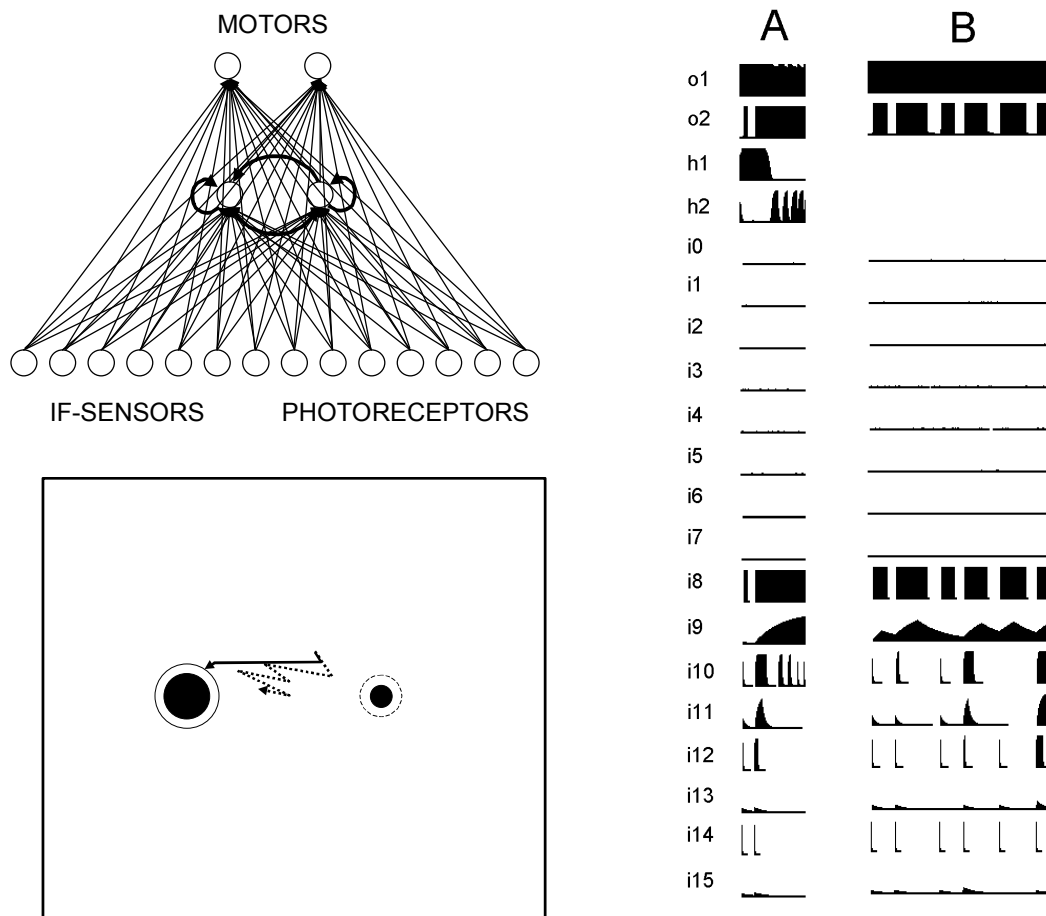


Figure 9. **Top, left:** the architecture of the neural controller. **Bottom, left:** typical behavior displayed by an evolved robot (the full line represents the trajectory of the robot in normal conditions, the dashed line represents the trajectory produced by the robot when the activation state of the two internal units is set to 0.0). **Right:** (A) displays the state of the two motor neurons (o1 and o2), two internal neurons (h1 and h2), and of the 16 sensory neurons (from i0 to i15) throughout time of the individual displayed on the left side of the Fig. (B) displays the state of the neurons when the activation state of the two internal neurons is set to 0.

This basic strategy is modulated by the internal neurons so to avoid the problems described above. In particular, in the case of the individual described in the Fig. 9, the second internal neuron modulates the robot behavior so to not lose the visual contact with the current object and consistently move toward it. The first internal neuron instead, by inhibiting the second internal neuron, forces the robot to quickly lose visual contact with the current object. During the very first cycles the first internal neurons get fully activated and therefore the robot starts to move toward the current object by quickly losing the visual contact with the object. This allows the robot to move more toward B than toward S, on the average according to the same strategy that we described above. Later on, however, the first internal neuron is inhibited and, when this happens, the robot starts to consistently move toward one of the two targets without going back and forth anymore thus avoiding the risk to not reach B in time.

There are several factors that contribute to inhibit the first and activate the second internal neurons. A first factor is simply time (i.e. more is the time that the robot spent negotiating with the two objects higher is the probability that it is now closer to B and higher is the risk to not reach B in time). A second factor is how much time the robot has spent moving toward the object that is currently in the view angle of the camera (higher is the distance traveled during a single segment higher is the probability that the current object is B). A third factor is the average size of the two landmarks. When both B and S are very small the robot, by losing the visual contact with the objects too quickly, might risk to not reach B in time. In these situations the second internal neuron becomes slightly activated so to allow the robot to lose the visual contact with the two objects more slowly.

To summarize individuals in the former experiment still rely on the same sensory-motor strategy adopted by reactive individuals and do not discover a completely different strategy. The additional internal states are used to modulate the basic sensory-motor strategy by taking into account the information gathered during the previous interaction with the environment.

4. Conclusion

In this paper we presented a set of experiments in which mobile robots able to visually discriminate between objects of different size were developed through artificial evolution. Evolved individuals were able to visually navigate toward the target objects and to solve the problem that sensory patterns are ambiguous (i.e. that the smaller object from a certain distance looks exactly like the bigger object from a higher distance) both in stationary and in varying environmental conditions.

The analysis of the results obtained by evolving robots in stationary environment showed that evolving individuals, by exploiting sensory-motor coordination, develop an extremely simple and effective strategy. Such form of behavior can be produced by a simple reactive neural controller in which sensory neurons are directly connected with motor neurons. Evolved individuals use motion to implicitly compute the distance and relative size (with respect to the robot point of view) of the two objects and to navigate toward the target.

The analysis of the results obtained by evolving robots in varying environmental conditions showed that, in this case, individuals develop a more complex strategy that also requires an ability to adapt on the fly to the current environmental circumstances. Indeed,

evolved individuals rely on the same sensory-motor strategy described above but are also able to modulate such behavioral strategy on the basis of the current environmental conditions by taking into account what they did before, that implicitly provides information on the external environment, and some aspects of the perceived environment such as the average perceived size of the two objects. The development of such adaptive strategy requires slightly more complex neural controllers provided with additional internal neurons with recurrent connections and time constant parameters that determine the amount of change in time of their activation state. In this case, evolved individuals looks organized in a hierarchical fashion where the higher layer, consisting of the internal neurons and their incoming and outgoing connections, keeps trace of information that is used to modulate the lower layer, consisting of the direct connections between sensors and motors, that is responsible for producing the same behavioral strategy displayed by the reactive individuals described above.

Overall these results demonstrate that artificial evolution might lead to the development of simple and effective solutions also in the case of problems that has to be solved in varying environmental conditions. The emergence of mixed strategies that rely both on simple reactive components and on more complex components that also take into account the previously experienced sensory states is particularly important from an evolutionary point of view because it suggest that artificial evolution might lead to an open-ended process in which progressively more complex solution are synthesized.

References

- [1] J K Holland, *Adaptation in Natural and Artificial Systems* (Ann Arbor, MI, University of Michigan Press, 1975)
- [2] H P Schwefel, *Evolution and Optimum Seeking* (New York, Wiley, 1995)
- [3] J R Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection* (Cambridge, MA, MIT Press, 1992)
- [4] S Nolfi, & D Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (Cambridge, MA, MIT Press/Bradford Books, 2000)
- [5] D Cliff, I. Harvey, & P. Husbands, Explorations in evolutionary robotics, *Adaptive Behavior*, 2, 1993, 73-110.
- [6] I Harvey, P Husband, & D Cliff, Seeing the light: artificial evolution, real vision, in D Cliff, P Husband, J-A Meyer, & S. Wilson (Eds.), *From Animals to Animats 3, Proceedings of the Third International Conference on Simulation of Adaptive Behavior* (Cambridge, MA: MIT Press/Bradford Books, 1994)
- [7] F Mondada, E. Franzi, & P. Ienne, Mobile robot miniaturization: A tool for investigation in control algorithms, in T Y Yoshikawa, & F Miyazaki (Eds.), *Proceedings of the Third International Symposium on Experimental Robots* (Berlin, Springer-Verlag, 1993)
- [8] X Yao, A review of evolutionary artificial neural networks, *International Journal of Intelligent Systems*, 4, 1993, 203-222.
- [9] S Nolfi, & D Marocco (2000). Evolving visually-guided robots able to discriminate between different landmarks, in J-A Meyer, A Berthoz, D Floreano, H L Roitblat, & S W Wilson (Eds.), *From Animals to Animats 6. Proceedings of the VI International Conference on Simulation of Adaptive Behavior* (Cambridge, MA, MIT Press, 2000)
- [10] S Nolfi, Evorobot 1.1 User Manual, *Technical Report* (Rome, Italy, Institute of Psychology, CNR, 2000)
- [11] S Nolfi , Power and limits of reactive agents, *Neurocomputing*, 42, 2002, 119-145.

- [12] S Nolfi, & D Marocco, Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences*, 120, 2001, 287-310.