

Detour Behavior in Evolving Robots: Are Internal Representations Necessary?

Orazio Miglino¹, Daniele Denaro², Guido Tascini³, and Domenico Parisi⁴

¹ II University of Naples, Via G.Paolo I, 81055 S.Maria.C.V., Italy
orazio@caio.irmkant.rm.cnr.it

² Department of Computer Science, University of Ancona, Via Breccie Bianche, 60131
Ancona, Italy
denaro@caio.irmkant.rm.cnr.it

³ Department of Computer Science, University of Ancona, Via Breccie Bianche, 60131
Ancona, Italy
tascini@inform.unian.it

⁴ Institute of Psychology, National Research Council, Via Marx 15, 00137 Rome, Italy
domenico@kant.irmkant.rm.cnr.it

Abstract. Internal representations of the environment are often invoked to explain performance in tasks in which an organism must make a detour around an obstacle to reach a target and the organism can lose sight of the target along the path to the target. By simulating a detour task in evolving populations of robots (Khepera) we show that neural networks with memory units perform better than networks without memory units in this task. However, the content of the memory units need not be interpreted as an internal representation of the position of target. The memory units send a time-varying internally generated input to the network's hidden units that allows the network to generate the appropriate behavior even when there is no external input. Networks without memory units do not have this internal input and this explains their inferior performance.

1 Detour behavior

Imagine a robot (an artificial organism) that must reach a target object located somewhere in the environment. The robot has a camera mounted on its body with which the robot can see the target. Since the camera has a restricted visual field the robot turns until the target happens to be within its visual field and it then can move in the direction of the target and reach it.

If there is no obstacle in the environment the task is easy enough (cf. e.g., [5]). Consider however an environment that includes an obstacle in addition to the target object. The obstacle is a low wall that allows the robot to see the target even if the target is located on the other side of the wall (provided the robot is oriented in the direction of the target) but prevents the robot from proceeding directly toward the target. The only solution for the robot is to make a detour around the wall in order to reach the target. But this creates a serious problem for our robot. When the robot is

making a detour around the wall the target is likely to exit the robot's visual field. The problem is particularly serious if the wall has the shape shown in Figure 3b. In order to negotiate the obstacle the robot may be forced to go in the opposite direction with respect to the target and therefore to lose sight of the target. (The type of task we have described is called "detour task" and is a standard procedure for testing real organisms' spatial abilities. [3, 9])

An organism which is able to make a detour around an obstacle is sometimes said to possess an "internal representation" of the spatial environment that mediates between stimulus and response (cf., e.g., [8]). When the organism sees the target and it can solve the task by simply approaching the target (i.e., there is no obstacle), the organism can directly map stimulus into response. However, when the organism must make a detour around an obstacle and therefore it is forced to lose sight of the target, one can think the organism is in possession of some internal representation of the environment that tells the organism where is the target even if the target is currently not visually accessible.

We do not find the notion of an internal representation particularly perspicuous, although this notion is often invoked in explaining the spatial behavior of organisms [6, 10]. In the present paper we analyze the detour behavior of (artificial) organisms by using a neural network model of the nervous system controlling the robot's behavior. We train a population of networks using a genetic algorithms so that evolved networks come to possess connection weights allowing them to exhibit various degrees of the ability to make a detour around an obstacle in order to reach a target object. We compare two conditions. In one condition the robot's behavior is controlled by a simple feedforward neural network. The network has two distinct sets of input units encoding the input, respectively, from a camera and from a set of infrared sensors. The first set of input units (camera) encode the position of the target with respect to the robot provided the robot's body is turned in the direction of the target and therefore the target falls inside the robot's visual field; otherwise, the target is not perceived by the robot. The second set of input units (infrared sensors) encode the position of the robot's body with respect to either the obstacle or the peripheral wall enclosing the environment provided the body is sufficiently close to either the obstacle or the peripheral wall; otherwise, the robot does not perceive the obstacle or the enclosing wall. Both sets of input units send their activations to a layer of hidden units that in turn send their activations to two output units. The two output units encode the rotation speed of the two wheels that allow the robot to turn and move in space.

In the second condition the neural network includes an additional layer of special units called memory units [4]. In any particular cycle the activation level of each hidden unit is copied in the activation level of a corresponding memory unit. Each memory unit is linked with normal excitatory or inhibitory connections to all hidden units. Therefore in the next cycle the activation level of the hidden units - and therefore the network's output - depends both on the current activation of the input units (and the input-to-hidden weights) and on their own activation in the preceding cycle stored in the memory units (and the memory-to-hidden weights). The memory units maintain a cumulative trace of the past and make this trace available to the hidden units so that the network in each cycle generates an output influenced by both the current input from the environment and this trace of the past.

In choosing to compare these two conditions our reasoning is the following. (For other simulations exploring the role of memory in robot's or animat's behavior, see

[1, 2]). It is tempting to equate the existence of memory units in a network architecture with the possession of an internal representation of the environment that makes it possible for an organism to respond to features of the environment even when these features are not currently perceived. When a robot without memory units must make a detour around an obstacle and it loses sight of the target, it has no way of knowing where is the target. Therefore, it will have trouble orienting itself and regaining visual access to the target in order to approach and reach the target. On the contrary, if the neural network controlling the robot's behavior includes an additional layer of memory units, the problem can be more easily solved. If in cycle N the robot sees the obstacle and in cycle $N+1$ it loses sight of the obstacle, the trace of the location of the obstacle will remain somehow stored in the memory units. Therefore, even if the robot has no direct visual access to the target it can respond to the current input using this stored trace of the position of the target. This trace will allow the robot to orient itself and to turn towards the target, thereby regaining visual access to it.

What we would like to determine is whether a robot without memory units can solve the detour problem or, in order to solve the problem, it is necessary to have memory units that allow the robot to maintain a trace of the past. In the latter case the content of the memory units would constitute some sort of representation of the environment which is necessary to solve the detour problem. Another possibility is that both robots with memory units and robots without memory units can solve the detour problem but the robots with memory units perform better than those without memory units. As we will see, this second result is obtained in our simulations. However, we don't think that this means that the content of memory units constitutes an internal representation of the environment. As we will show, the role of the memory units in enhancing performance in a detour task can be analyzed in a simpler and more basic way that does not require internal representations

2. Simulations

In our simulations we use a small robot called Khepera which has a cylindrical body with a diameter of 5 cm and a height of 8 cm (Figure 1). Around the body there are 8 proximity (infrared) sensors with a nonlinear range of 3 cm. However, we use only 3 of these 8 sensors: left, front, and right. Mounted on the top of the robot there is a camera with a linear array of 64 pixels. The camera has a horizontal visual field of 36 degrees at 8 cm from the ground. Each pixel can have one of 256 grey levels. For our purposes the 64 pixels were reduced to 8 (sampled each 8 pixels on the entire array of 64 pixels) with only 2 grey levels (black and white). The robot turns and moves by regulating the rotation speed of two wheels.




Fig. 1. The Khepera robot

The environment is a rectangular walled arena of 55x40 cm. It contains a low obstacle that can have various shapes and a cylindrical target. The proximity sensors of the robot can sense the obstacle or the target or the peripheral wall provided the robot is sufficiently close to them. The camera can see the target (but not the obstacle or the wall) at any distance (provided the target falls within its restricted visual field) even when the target is on the other side of the low obstacle with respect to the robot.

We developed a simulator for the robot and the environment because we wanted to use a genetic algorithm to train the robot and using a genetic algorithm with real robots is too expensive in terms of time and management. The simulator is based on a software library in C language for managing bidimensional environments and mobile units developed at our Institute. After a population of robots has evolved the required capacity we test the control systems (neural networks) in a real (physical) environment with the physical Khepera. Generally, very similar behaviors were exhibited by the simulated and the physical robots.

The control system of our robots is constituted by a neural network. We used two types of networks in two different populations: with and without memory. The network without memory is a normal feedforward neural network with a single layer of hidden units. The network has 11 input units (3 for the infrared sensors and 8 for the camera), 2 output units (for the robot's 2 wheels), and 5 hidden units. The network with memory has an additional layer of memory units, one for each hidden unit. In any given cycle each of the memory units is assigned the same activation level of the corresponding hidden unit. Since all the memory units are linked with normal connections to all the hidden units, in any given input/output cycle the activation pattern appearing on the hidden units (and therefore the network's output) depends on both the current activation pattern of the input units and the cumulative trace of the past stored in the memory units. (Notice that in the networks with memory the hidden units are 4, that is, one unit less than in the networks without memory. This guarantees a comparable total number of connections for the two types of networks (83 connections for the networks without memory units and 89 connections for the networks with memory units) and therefore a connection space of comparable size to be searched by the genetic algorithm).




Fig. 2. Neural network controlling the robot' s behavior: (a) without memory; (b) with memory

The simulation starts with a population of 100 organisms all with the same network architecture and randomly assigned connection weights. The weights are randomly selected from a rectangular distribution ranging from -5 to +5. Each individual lives alone in an environment that contains a target and an obstacle. The individual must make a detour around the obstacle to reach the target. The individuals that perform better are selected for reproduction. In order to facilitate the evolution of the detour behavior the population evolves in a simpler environment in the first 200 generations, and then for 200 additional generations in a slightly more complex environment. For the first 200 generations the environment contains a rectangular obstacle (a bar of 20 cm of length and 3 cm of height) and a target (a cylinder of 2 cm of diameter and 10 cm of height) (Figure 3a). The life of all individuals lasts a total of 240 input/output cycles divided into two epochs of 120 cycles each. At the beginning of each of the two epochs the individual is positioned in the environment below the obstacle facing North. The target is positioned near the upper left corner of the environment in one epoch and near the upper right corner in the other epoch in such a way that the target falls outside the organism' s restricted visual field. The organism has available 120 cycles to make a detour around the obstacle and reach the target.

After the first 200 generations the environment changes and it becomes more complex. The obstacle now has a U-inverted shape (Figure 3b) and is moved somewhat more North. The lifetime of an organism lasts 6 epochs of 120 cycles each. In three of these 6 epochs the organism is initially placed below the obstacle facing North and the target is located once in the upper left corner, once in the upper right corner, and once midway between the two corners, i.e., in the middle of the upper portion of the environment (cf. Figure 3b). In the remaining 3 epochs there is no obstacle and the organism and the target are placed as in the 3 epochs with the obstacle. These 3 epochs with no obstacle were included to allow the organisms to evolve an ability to choose between a direct route to the target when no obstacle is present and an indirect detour route when there is an obstacle.

The fitness on the basis of which individuals are selected for reproduction is based on two components. An individual is assigned one fitness unit for each cycle the individual spends near the target, where nearness to the target means to be located 12 cm or less from the target. This is the first component. The second component is

calculated according to the following formula (used in standard evolutionary robotics experiments [7]):

$$F = \text{abs}(V) * (1 - \text{sqrt}(DV)) * (1 - I) \quad (1)$$

Where V is velocity (absolute value of the sum of activations of the two wheels), $\text{sqrt}(DV)$ is the square root of the difference between the activations of the two wheels, and I is the average activation of the 3 infrared sensors. In other words, the second fitness component rewards the individual if it moves fast and straight (without turning) and it punishes it for being too close to the obstacle or wall (i.e., when the infrared sensors are activated; notice that the second fitness component is ignored when the organism is near the target). An individual's performance is evaluated in each cycle using the above formula. An individual that exhibits the best possible performance is assigned one fitness unit. Individuals that are not so good are assigned less than one unit of fitness according to their performance.




Fig. 3. First environment (a); second environment (b). The dotted areas identify alternative target positions.

At the end of life each individual is assigned a fitness which is the average of the fitnesses obtained in each epoch of life (2 epochs in the first block of 200 generations and 6 epochs in the second block of 200 generations). The 100 individuals of each generation are ranked on the basis of their fitness and the 20 best individuals are selected for reproduction. Each of these 20 individuals generates 5 offspring, i.e., that is, 5 neural networks with the same connection weights of their (single) parent except that 20% of the weights are randomly modified by adding a quantity randomly selected in the interval between -1 and +1 to their current value.

3 Results

Figure 4 shows how fitness changes across the 200+200=400 generations for (a) the best individual, (b) the 20 best individuals (average), and (c) the total population of 100 individuals (average) in each generation, for both the population with memory and the population without memory. (The data are the average of 5 replications of the

simulation with different "seeds" for constructing the initial population.) The fitness increases across the 400 generations, which implies that our organisms learn (in an evolutionary sense) the task, i.e., they acquire an ability to reach the target after making a detour around the obstacle. In most of the curves there is a rapid increase in fitness after generation 200 which appears to be due to the fact that in the second environment our organisms live half of their life (3 out of 6 epochs) in an environment without the obstacle, which allows them to reach the target more easily. Aside from that, there is a clear and significant advantage of the individuals with memory units in their network architecture over the individuals without memory units.

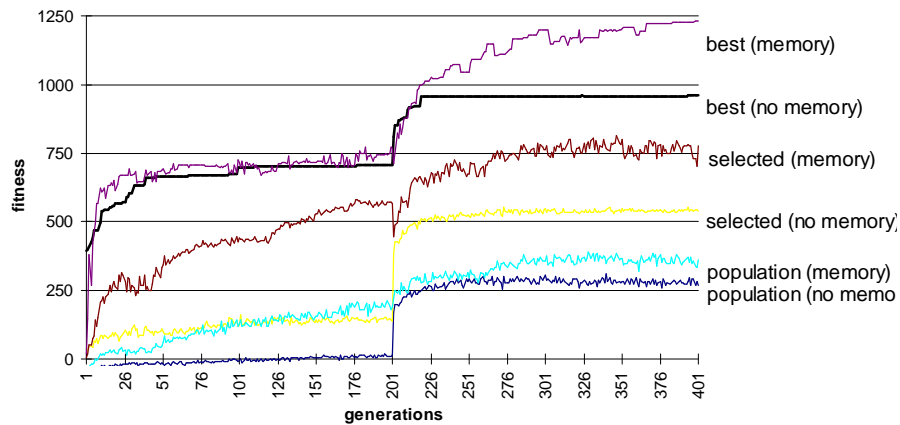


Fig. 4. Change in fitness across 400 generations for the best individual, the 20 best individuals, and the entire population of 100 individuals in each generation, for both the population without memory and the population with memory.

At the end of evolution we tested for generalization the 5 best individuals of the last generation in each of the 5 replications of the simulation for both the population with memory and the population without memory, for a total of $25+25=50$ individuals. Each of these 50 individuals is allowed to live for 100 epochs (120 cycles each) and in each epoch (a) the initial position of the organism is the same as during evolution but its orientation is chosen at random (whereas it was always North during evolution), (b) the location of the target in the upper portion of the environment is chosen at random (whereas it was restricted during evolution, cf. Figure 3b), and (c) the obstacle in different epochs is either (1) the bar of the first environment, (2) the inverted-U obstacle of the second environment, (3) an L-shaped obstacle, or (4) is absent. The performance in this generalization test (generalization test A) was measured as the percentage of epochs in which an organism was able to reach the target. An organism was considered as having reached the target if it was able to reach a location at less than 3 cm from the target.

The results indicate that when tested in an environment which is different from the

environment they have experienced during evolution the organisms with memory perform more efficiently than the organisms without memory. They are able to reach the target in 78% of the epochs compared with 65% of the organism without memory. Using the t Student test we found a statistically significant difference ($t(48)=2.89$, $p<0.01$).

These results establish that keeping a neural trace of the past in such a way that in any given cycle the organism responds to both the current input and this trace of the past results in a better performance in the detour task both in the environment in which the task has been learned (has evolved) and in new environments never experienced before. Both the organisms with memory and the organisms without memory evolve an ability to make a detour around an obstacle to reach a target but the organisms with memory perform better than those without memory.

Then our next questions are: How is the behavior of the organisms with memory different from the behavior of the organisms without memory? How can the memory units help an organism to solve the task more efficiently?

To answer these questions we divide the entire behavioral sequence that takes the organism from its starting position to the target into two successive segments. In the first segment the organism must negotiate the obstacle and turn around it. In the second segment the organism leaves the obstacle behind itself and approaches the target. Operationally, the first behavioral segment ends when the organism finds itself North of the obstacle for the first time and the second behavioral segment includes everything which happens after that event. We will show that the organisms with memory and the organisms without memory behave differently with respect to both the first and second segments of the behavioral sequence.

As the fitness curves of Figure 4 show, an interesting result of our simulations is that even the organisms without memory are able to turn around the obstacle and head towards the target even if in order to do so they must necessarily lose sight of the target. These organisms can solve the problem by adopting the strategy of following the obstacle. They might approach the obstacle and respond to the sensory (infrared) input from the obstacle to generate movements that allow them to walk along the obstacle until they have turned around the obstacle and can regain visual access to the target. At this point they can simply approach the target. Or they can go away from the obstacle, approach the peripheral wall, and use the sensory (infrared) input from the wall to reach the upper portion of the environment where the target is located.

On the other hand, the organisms with memory have an alternative option open to them. When there is no visual input from the target because the organism is facing away from the target and there is no input from the infrared sensors because the organism is too distant from either the obstacle or the wall, the organism is completely deprived of sensory input since it is exposed to a constant zero input. The organisms without memory can only escape from this difficult situation by approaching either the obstacle or the wall in order to get some useful sensory input. The organisms with memory, on the contrary, can use the constantly changing internal input from their memory units to generate appropriate behaviors.

This analysis predicts that the organisms without memory will evolve behavioral trajectories that tend to bring them near either the obstacle or the peripheral wall more often than the organisms with memory. These trajectories allow the organisms without memory to be exposed to sensory input from either the obstacle or the wall more often than the organisms with memory. The organisms with memory do not need to do so because they have a time-varying input which is internally generated by

the memory units.

To test this prediction we have determined the percentage of cycles in generalization test A (in the epochs with an obstacle) in which the organism receives sensory input from its infrared sensors, for both the organisms with memory and the organisms without memory. This percentage is 42% for the organisms with memory and 63% for the organisms without memory (the difference is statistically significant, $t(48)=3.19$, $p<0.01$). We conclude that the two types of organisms behave differently in the first portion of the behavioral sequence that allows them to reach the target. They both must turn around the obstacle and therefore lose sight of the target. However, their behavior of negotiating the obstacle is different. The organisms without memory tend to stay close to the obstacle or to approach the peripheral wall in order to receive sensory input from the obstacle or the wall that they can use to go around the obstacle and reach the upper portion of the environment where they will find the target. In fact, this input is the only input that can guide their behavior in the first segment of the behavioral sequence that will take them to the target. In contrast, the organisms with memory are able to turn around the obstacle and reach the upper portion of the environment with less input from their infrared sensors because they can count on the internal input from their memory units to guide their behavior.

We now turn to the second segment. The second segment begins when the organism has already made a detour around the obstacle and has reached the upper portion of the environment where the target is located. What is demanded of the organism at this stage is to turn toward the target in order to gain visual access to the target and to use this visual input to head toward the target. This is a simple stimulus-response task that can be solved without memory units and, therefore, at least in principle the behavior of the organisms with memory and the behavior of the organisms without memory should not be different at this stage.

To test how the two kinds of organisms behave in the second segment of the behavioral sequence we compare the behavior of the 25 individuals with memory and of the 25 individuals without memory in 100 epochs in which the environment contains no obstacle (generalization test B). At the beginning of each epoch (120 cycles) the organism is placed in the standard starting position at the center of the environment with a randomly selected orientation and the target is placed in a randomly selected position in the upper portion of the environment. We measure (a) the percentage of epochs in which an organism is able to reach the target (i.e., to reach a location at 3 or fewer centimeters from the target), (b) the average number of cycles necessary for an organism to reach the target, and (c) the percentage of cycles an organism receives sensory input from the wall.

The results show that there are no differences between the organisms with and without memory with respect to measure (a). Both types of organisms reach the target equally often, that is, in most of the epochs (92% for the organisms with memory; 90% for the organisms without memory). (Remember that there is no obstacle in the environment in this test.) However, the average time taken by the organisms with memory to reach the target is 20 cycles while it is 32 cycles for the organisms without memory (statistically significant difference; $t(48)=4.32$, $p<0.01$). Therefore, memory allows our organisms to reach the target more quickly. Furthermore, the organisms without memory are less likely than those without memory to end up near the peripheral wall in their search for the target. This is suggested by the fact that the average number of cycles in which the organisms with memory receive input from their infrared sensors is 40 (in each epoch) while it is 67 for the organisms without

memory ($t(48)=2.97, p<0.01$).

Table 1 summarizes the results of the generalization tests. Typical trajectories of organisms with and without memory in an environment with an obstacle and in an environment with no obstacle are shown in Figure 5.

Table 1. Summary of generalization test results.

Type of measure	Without memory	With memory
% of target reachings in generalization test A	65	78
% of target reachings in generalization test B	90	92
# of cycles to reach the target in generalization test B	32	20
# of cycles to reach the target from first detection (test B)	21	10
# of cycles with infrared sensory input activated (test A)	63	42
# of cycles with infrared sensory input activated (test B)	67	40




Fig. 5. Typical trajectories of organisms without memory (top) and with memory (bottom) in an environment with an obstacle (left column) and in an environment with no obstacle (right column).

4 Discussion

One important result of our simulations is that even simple input-output networks without memory units can to some extent solve a detour task. Although in some cycles the robot is necessarily unable to see the target because it has to turn away from the target in order to make a detour around the obstacle, the robot does not need to wander randomly in the environment until it happens to have gone past the obstacle and to regain visual access to the target. The robot can approach the obstacle and/or the peripheral wall and using the sensory input from its infrared sensors it can generate a trajectory that makes a detour around the obstacle and takes the robot to the upper portion of the environment where the target can be directly reached.

However, solving a detour problem with a simple feedforward neural network without memory units generates an inferior performance compared with solving the problem with memory units. The organisms with memory units perform better than those without memory units both during evolution and in generalization tests with environments different from those experienced during evolution. They have higher fitness during evolution and their capacity to make a detour around the obstacle and reach the target is more robust and generalizable.

However, it is not clear that the superior performance of the organisms with memory can be attributed to the fact that they store a trace of the position of the target and use this trace to guide their behavior when they do not have direct visual access to the target. The memory mechanism of our neural network is a cumulative one, which implies that the memory trace of each successive events is superimposed on the trace of the previous events. Therefore, the memory trace of an event becomes weaker and weaker as the event recedes back in time. Since our organisms can lose sight of the target for a long sequence of input/output cycles, the memory trace of the position of the target inevitably fades away and is unlikely to be used to guide the robot's behavior.

Our more detailed analysis of the behavior of our robots seems to indicate that the advantage of possessing a memory mechanism for solving the detour task lies elsewhere. This analysis has shown that the robots without memory tend to approach the obstacle and/or the peripheral wall more often than the robots with memory during the trajectory that brings them to the target. In fact, their proximity sensors are more activated than those of the robots with memory both during the early portion of the trajectory when the robots have to negotiate the obstacle and in the later portion when the robots can directly approach the target. The reliance on sensory input from the obstacle and the wall allows the robots without memory to solve the task but causes them to exhibit a less efficient performance than the robots with memory. They need to go near the obstacle or the wall to get access to sensory input from the obstacle or the wall to guide their behavior, and this inevitably slows down their performance because they take more time to reach the target.

The organisms with memory are in a different situation. Even if the external input from the environment does not change from one cycle to the next cycle when they have no visual access to the target and no sensory input from the obstacle or the peripheral wall, and therefore the external input cannot guide their behavior, these robots have access to an internal input from their memory units which does change from cycle to cycle. Therefore, in the «blind» cycles (i.e., the cycles without external

input), the robots with memory can use the varying internal input to generate the appropriate behavior. The result is that they do not need to waste time approaching the obstacle and/or the peripheral wall and they can generate shorter trajectories to the target. Hence, their superior performance both during evolution and in the generalization tests.

Returning to the notion of an «internal representation» as a necessary prerequisite for solving a detour task, we conclude that this notion is not necessary, at least in the conditions examined in our simulations. Simple feedforward neural networks can solve the task, although not very efficiently. But even the more efficient performance in the detour task of neural networks with memory units does not require a notion of internal representation. One might think that neural networks with memory units store an internal representation of the environment (more precisely, of the position of the target relatively to the organism) so that the network can use this internal representation when it has no direct visual access to the target. But our analyses have shown that the superior performance of the networks with memory units is probably determined not by a internal representation of the position of the target but by a simpler and more basic mechanism according to which networks with memory units can rely on the varying internal input from the memory units when the external input is not available - an option not open to networks without memory units.

References

1. Donnar, J. Y., Meyer, J. A. Learning reactive and planning rules in a motivationally autonomous animat. *IEEE Transactions on Systems, Man, and Cybernetics-Part B* **26**(3) (1996) 381-395.
2. Dorigo, M., Colombetti, M. Robot shaping: Developing autonomous agents through learning. *Artificial Intelligence* **71**(2) (1994) 321-370
3. Ellen, P., Thinus-Blanc, C. (eds.): *Cognitive processes and spatial orientation in animal and man*. Dordrecht, Martinus Nijhoff (1987)
4. Elman, J.L.: Finding structure in time. *Cognitive Science* **14** (1990) 179-211
5. Floreano, D., Mondada, F.: Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics-Part B* **26**(3) (1996) 396-407.
6. Gallistel, C. R.: *The Organization of Learning*. Mit Press, Cambridge (1990)
7. Nolfi, S., Floreano, D., Miglino, O., Mondada, F.: How to evolve autonomous robots: Different approaches in evolutionary robotics. In Brooks R. A., Maes P. (eds.): *Artificial Life IV. Proceedings of the Fourth International Workshop on The Synthesis and Simulation of Living Systems*. MIT press, Cambridge (1994) 190-197
8. Regolin, L., Vallortigara, G., Zanforlin M.: Object and Spatial representations in detour problems by chicks. *Animal Behavior* **48** (1994) 1-5
9. Thinus-Blanc, C.: *Animal Spatial Cognition*. World Scientific, Singapore (1996)
10. Tolman, E. C.: Cognitive maps in rats and men. *Psychological Review* **36** (1948) 13-24