

# Formation of spatial representations in evolving autonomous robots

O. Gigliotta<sup>1,2</sup>, S. Nolfi<sup>1</sup>

<sup>1</sup>Institute of Cognitive Sciences and Technologies, CNR, Viale San Martino della Battaglia 44, 00185, Roma, Italy

<sup>2</sup>Department of Psychology, University of Palermo, Viale delle Scienze, Ed.15, 90100, Palermo

**In this paper we describe how a population of evolving robots can autonomously develop forms of spatial representation that allow them to discriminate different locations of their environment. Developed representation forms consist of patterns of activations of internal neurons that are generated by integrating sequences of sensory-motor states while the robots interact autonomously with their environment. Moreover, these representation forms are allocentric, i.e. they allow evolved robots to identify a given spatial location independently from the robots' initial position and orientation in the environment. The analysis on the robots' representation system indicate that it can be characterized as a limit cycle resulting from the transient dynamic between fixed attractor points that alternate while the robot move in the environment. We also demonstrate how the evolved representation systems display remarkable generalization properties. We conclude the paper by discussing the characteristics of the representation system developed by the robots and its relation with other models described in the literature.**

## I. INTRODUCTION

The objective of this research work is to investigate whether evolving robots provided with simple recurrent neural controllers can solve a problem that requires to self-localize and to discriminate different spatial locations of the environment. The environment and the task are constructed in a way that require the categorization and the use of abstract features of the robot/environment interaction that cannot be inferred from any single sensory pattern and that require to integrate sensory-motor information through time [2, 10, 11, 7]. The goal is that to identify the simplest experimental setting in which this type of problem can be solved and to analyse the characteristics of the solution found.

The problem of discriminating different spatial locations is a key feature of animal navigation abilities. Indeed, although the ability to discriminate different spatial locations in the environment might not be necessary to exhibit simple navigation forms, it is certainly crucial for the ability to display more complex forms of navigation [12]. However, this research also aims to study, more generally, how robots can discover abstract categories, extract these categories online, while interacting with the external environment, and use them to identify different locations of the environment.

The reason for using a self-organizing technique such as artificial evolution is due to the fact that we are interested in studying whether robots can develop and use an effective categorization ability autonomously rather than developing

robots that can solve their task on the basis of hand-crafted solutions. We chose to use artificial evolution, rather than other learning techniques, since it allows leaving robots free to determine the control parameters and the behaviours that satisfy their adaptive goal by reducing the constraints on how the problem should be solved to the minimum. However, we expect that similar results can be obtained by using other learning techniques providing that variations occurring during learning modify the fine-grained interactions between the robot and the environment and variations are retained or discarded on the basis of their effects on the overall behaviour exhibited by the agent [8].

## II. RELATED LITERATURE

In this section, we review the related literature by restricting our analysis to research works in which the way in which robots encode spatial information is not pre-determined by the experimenter and is developed by robots through an adaptive process while they interact with the environment.

In a recent work, Vickerstaff and Di Paolo [13, 14] evolved simulated robots (provided with a compass sensor, speed sensors, and two light sensors) for exhibiting a homing behaviour. The robots, which are initially placed in their home location, are selected for the ability to return in the same location after having reached a variable number of light beacons placed randomly. Beacons are placed one at a time (i.e. a new beacon is displaced as soon as the robot reaches the previous beacon). Robots were provided with a CTRNN [1]. The free parameters, which were encoded in evolving genotypes, included the connection weights and the architecture of the neural controller. The analysis of the obtained behaviour indicates that evolved robots solve their problem by using a path integration mechanism, i.e. by encoding the homing location in the activation state of a vector of neurons (or a vector of modifiable weights, in a variation of the experiment). This internal vector is continually updated by integrating sensor information through time while the robot travels in the environment so to always encode an updated estimation of the current position of the robot relative to the home location. Evolved robots also display a searching behaviour, once they have reached their estimate of the nest location, similar to that exhibited by ants.

As in the case of Vickerstaff and Di Paolo work, we develop robots able to display navigation ability by leaving them free to determine the way in which they internally

encode spatial information. The main difference with the work described in this paper lays in the nature of the problem. The work of Vickerstaff and Di Paolo involves a homing problem that consists in returning to a recently visited environmental location and that can be solved through a path integration method --- a navigation method similar to that employed by sailors under the name of dead reckoning that consists in continuously updating the estimation of the relative location of the reference point (the nest) on the basis of the estimated direction and speed of the movements performed after abandoning the reference point. The work described in this paper instead involve a place recognition problem in which robots should be able to discriminate their current location by identifying regularities in their sensory-motor flow while they move in the environment.

In an earlier related work, Floreano and Mondada [3] evolved a Khepera robot [5], provided with infrared, light, and battery level sensors, for the ability to move as straight and as fast as possible by periodically returning to a recharging area located close to a light beacon to recharge its battery. Evolving robots display an ability to integrate spatial and energy consumption information so to allow the robot to periodically return to the charging station just before the energy level goes below a "survival" threshold. Since the light gradient provided a straightforward indication of the recharging area and it is visible from any environmental location, evolving robots does not need to evolve a path integration mechanisms to identify the relative location of the target. Moreover, the relative location of the target can be identified in any single time step on the basis of the current state of the light sensors. The main difference between Floreano and Mondada and our experimental scenario is that in our case robots cannot infer their location in the environment on the basis of a single sensory pattern only since robots experience identical sensory patterns in different environmental locations.

Nolfi [6] evolved a simulated Khepera robot, provided with infrared sensors only, to navigate into a two room's environment by discriminating the room in which it is currently located. As in the case of the experiment described in this paper, robots cannot rely on a path integration method since robots' initial position in the environment is randomly initialised. Moreover, as in the case of the work described in this paper, robots experience identical sensory states in different environmental locations, and therefore should be able to discriminate their current environmental location by integrating sensory-motor patterns through time. The work described in this paper is related to this previous work but present new features that allow agents to display a greater expressive power and remarkable generalization abilities.

### III. THE EXPERIMENTAL SETUP

Consider the case of an E-puck robot (Fig. 1, left) placed in a double T-maze environment (Fig. 1, right) that should be able to explore the environment, memorize the location of the target area, and recognize that location after being placed in a

randomly selected position in the environment. The location of the target area is marked by a black disk placed on the ground and can be easily detected by the robot through an infrared sensor pointing toward the ground (Fig. 2, left). The black disk, however, is removed after the target area is reached by the robot for the first time (Fig. 2, right). To recognize the location of the target area after the marker has been removed, therefore, the robot should keep in memory an indication of the estimated position of the target area and recognizes when it reaches the same location again.

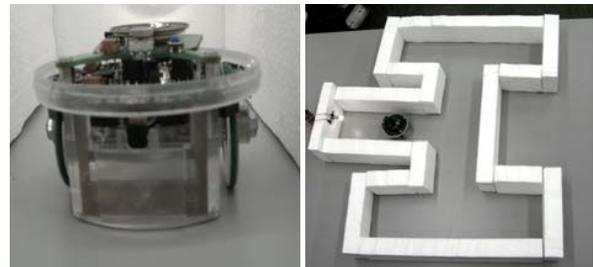


Fig. 1. Left: The e-puck robot developed at EPFL, Switzerland (<http://www.e-puck.org/>). Right: The double T-maze environment including the light bulb and the robot.

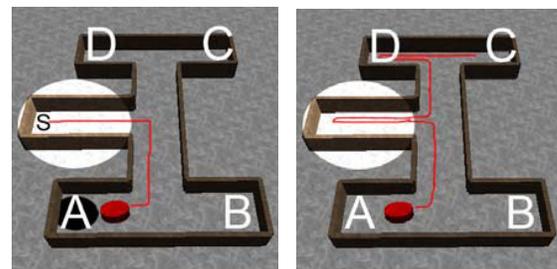


Fig. 2. The robot and the environment in simulation. The white circles indicate the amplitude of the light gradient. The A, B, C, and D labels indicate the position in which the target area is located in different trials of the experiment. The black line indicates the trajectory of an evolved robot navigating in the environment. **Left:** during the first phase of each trial the robot explore the environment until it encounter the target area that is marked with a black circle. **Right:** during the second phase of each trial the black circle is removed, and the robot is asked to return and to stop on the target area location after having moved to a new randomly selected location.

The T-maze arena covers an area of 52cm by 60cm. The robot has eight infrared sensors (that provide information about nearby obstacles up to a distance of about 5cm), eight light sensors (that provide the light gradient information up to a distance of about 40 cm from the light), one ground sensor (that detects the colour of the ground), and two motors (that control the desired speed of the two corresponding wheels).

Each robot is controlled by a neural network with a fixed architecture including 18 sensory neurons (that encode the state of the eight infrared sensors, of the eight light sensors, and of the two location sensors), two internal neurons with recurrent connections, and three motor neurons (that encode the desired speed of the two wheels and whether or not the location of the target area has been reached). During the first part of each trial in which the robot has not yet reached the

target area, the location sensors are set to the maximum value (1.0). During the second part of each trial, the state of the two location sensors is set to the absolute difference between the current activation state of the internal neurons and the activation state of the internal neurons observed when the robot detected the black disk indicating the location of the target area for the first time.

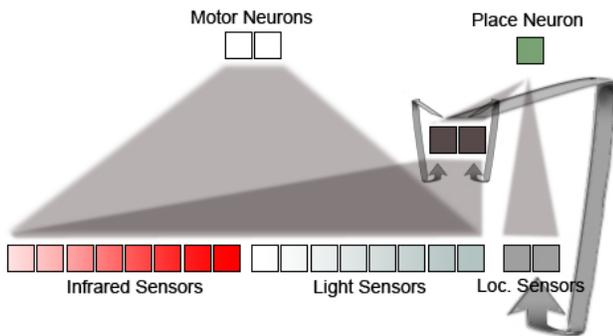


Fig. 3. The architecture of robot's neural controller. The grey areas indicate the connections between blocks of neurons. The arrow on the left indicates that the activation state of the two internal neurons is copied into the location sensors when the robot detects the black disk placed on the target area. The place output neuron should be turned on when the robot returns on the location of the target area.

The free parameters of the robots' neural controllers are evolved [9]. The initial population consisted of 100 randomly generated genotypes that encoded the connection weights, the biases, and the time constants [6] of sensory and internal neurons of 100 corresponding neural controllers. Each parameter is encoded with 8 bits and normalized in the range  $[-5.0, +5.0]$ , in the case of connection weights and biases, and in the range  $[0.0, 1.0]$ , in the case of time constants. The 20 best genotypes of each generation were allowed to reproduce by generating five copies each, with 2% of their bits replaced with a new randomly selected value. The evolutionary process lasted 800 generations (i.e. the process of testing, selecting and reproducing robots is iterated 800 times). The experiment was replicated 10 times.

Robots are tested for 100 trials. At the beginning of each trial the robot is placed in location S with a randomly selected orientation and the target area is randomly set within one of the four possible locations indicated with letters A, B, C, and D (Fig. 2). The robot is allowed to interact with the environment for 3300 cycles (each cycle last 100ms). After 950 cycles a black disk is placed in the location of the target area and is then removed as soon as the robot reaches it. The fitness of evolving robots is computed according to a function with two components that reward respectively: (1) robot's ability to navigate in the environment so to visit all its sub-areas, and (2) robot's ability to return on the target location and to stop there by turning its' place neuron on (i.e. by exciting a threshold of 0.75).

#### IV. RESULTS

By analysing the behaviour displayed by the best individual of the last generation we observed that, in most of the replications of the experiments, robots display a navigation behaviour that allow them to periodically visit all the locations of the environment by moving forward in corridors, turning right on T-junctions, and turning back at corridors' end points (see Fig. 2). Moreover, evolved robots display an optimal ability to return and to recognize the location of the target area after having visited the target area once and after the black disk indicating the location of the target area has been removed (Fig. 4, white histograms). Evolved robots also display an ability to recognize the location of the target area after being placed in a random selected location of the environment at the beginning of the trial (Fig. 4, grey histograms). Finally, by downloading the evolved neural controller into a real e-puck robot and by testing the robot in the real environment shown in Fig. 1, we observed that controllers evolved in simulation display an ability to solve the task also when embodied in the real robot and situated in a real environment (Fig. 4, black histograms).

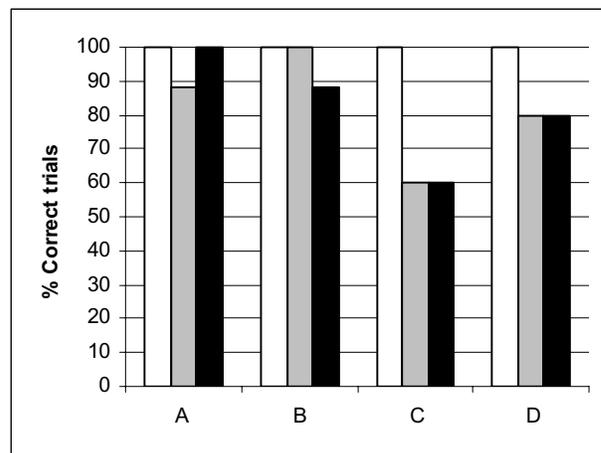


Fig. 4. Percentage of trials in which the robot is able to return and to stop on the target area after the robot experienced the black disk indicating the location of the area and after the disk was removed. Data for one of the best evolved individuals. The four groups of histograms indicate the performance for the four possible location of the target area. White histograms indicate the performance obtained in simulation by initially placing the robot close to the light. The grey histograms indicate the performance obtained in simulation by initially placing the robot in a randomly selected location and orientation of the environment. The black histograms indicate the performance observed on the real robot in the real environment by placing the robot in a randomly selected location and orientation in the environment at the beginning of each trial.

To solve this problem, evolved robots develop an ability to integrate over time the state of the sensory neurons experienced while the robot navigates in the environment into internal states that are different when the robot is located in different part of the environment, and that are similar when the robot visit and re-visit the same environmental location. Indeed, by plotting the state of the robots' internal neurons

recorded when the robot is located in the four possible locations of the target area (Fig. 5) we can see how the states corresponding to the same location are clustered together and the states corresponding to different locations are separated in the two dimensional space that includes all possible values of the internal neurons. Data have been obtained in the condition in which robots start from a randomly selected position and orientation in the environment.

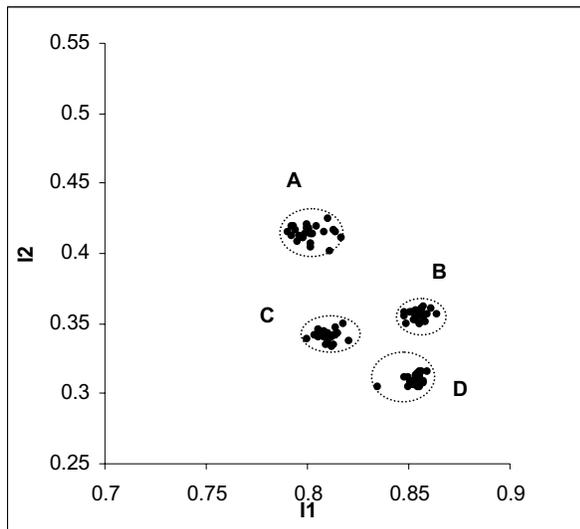


Fig. 5. The black dots indicate the activation state of the internal neurons recorded while the best evolved robot is located in the four possible locations of the target area in different trials in which the robot start from randomly selected locations and orientations in the environment. Data recorded during 100 trials in which the robot start from different randomly assigned location and orientation in the environment after the robot moved in the environment for 3300 cycles corresponding to 330 s. I1 and I2 indicate the activation state of the two internal neurons. The letter A, B, C, and D, and the circles indicate how activation states corresponding to similar locations in space are clustered in different areas of the state space corresponding to the two internal neurons.

#### V. ON ROBOTS' ABILITY TO REPRESENT ENVIRONMENTAL LOCATIONS

In this section, we analyse how robots represent different locations of the environment and how are able to generalize their skills when they start from different randomly assigned locations.

By analysing the dynamics of the two internal neurons while the robot moves in the environment we can see how after 95 s (i.e. after the robot makes about 2 laps of the environment) the state of the internal neurons converges toward a limit cycle attractor (Fig. 6). By analysing the dynamics of the internal neurons when the robot is situated in the environment in this and in other replications of the experiment (results not shown), we observed the following characteristics: (1) *periodicity*: the dynamic of the internal neurons tends to converge toward a periodic limit cycle that repeats itself for every lap of the environment, and (2) *expressiveness*: different points along the trajectory of the limit cycle correspond to different locations of the robot in the

environment (beside few points in which the trajectory of the limit cycle crosses points visited previously while the robot was located in a different part of the environment). Periodicity and expressiveness can be demonstrated by ascertaining that the variation of the state of the internal neurons in a given environmental location for different laps of the environment is lower than the variation of the state of the internal neurons for different locations of the environment (result not shown).

It should be noted that the shape of the limit cycle significantly vary in different replications of the experiments. Indeed, several different shapes might preserve the periodicity and expressiveness characteristics that are necessary to solve the problem.

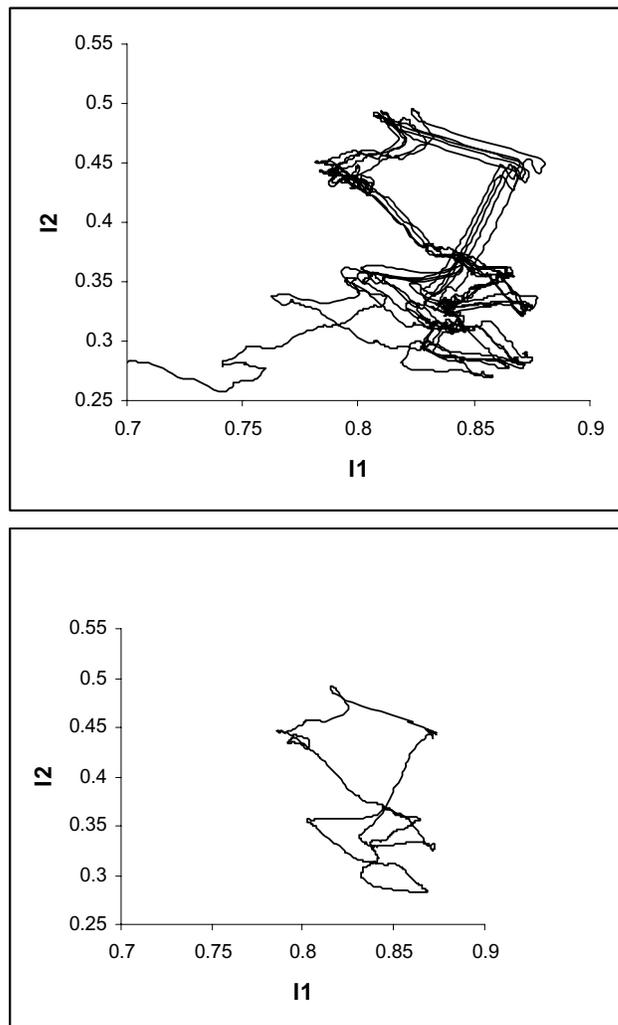


Fig. 6. Top: the activation state of the internal neurons recorded while the robot moves in the environment for 330 s. Bottom: limit cycle attractor obtained by averaging the activation state of the internal neurons in different laps of the environment after the robot performed the first lap.

The periodicity of the limit cycle explains how an evolved robot generalizes its skill with respect to its initial position. In fact, after a certain amount of time in which the robot moves

in the environment, the state of its internal neurons correlates only with the current position of the robot in the environment and it is not affected anymore by the initial sensory experiences of the robot that in turn were affected by the initial location and orientation of the robot in the environment. The expressiveness of the limit cycle ensures that different locations of the environment correspond to different internal states. The combination of these two properties ensures that the robot can reliably represent its position in the environment so to distinguish, for example, between location A, B, C, and D and can self-localize after being placed in a randomly selected location and orientation.

The periodic limit cycle described above results from a transient dynamics in which the state of the internal neurons moves toward fixed attractor points corresponding to the sensory states that alternate while the robot navigates in the environment (see also [2]). The existence of fixed attractor points is demonstrated by the fact that by fixating the state of sensory neurons on any given state experienced by the robot while it moves in the environment and by keep updating the internal neurons for a sufficient time, the state of the internal neurons converge toward a fixed attractor point independently from the initial neuron state (Fig. 7).

For example, as shown in Fig. 8, when the robot ends one laps of the environment, the state of the internal neurons occupies the central position of the limit cycle. Then, while

the robot travels toward the light, the state of the internal neurons moves toward the TL average attractor situated in the top-right corner of the state space (see Fig.8 “s”). Then, when the robot turns back and moves away from the light, its internal state tends to move toward the AL average attractor located in the top-left corner of the state space. Later on, when the robot turns right and then moves along the corridor, its internal state moves toward the average attractors TR and CO located in the left and right side of the state space respectively. These latter movements bring the state of the internal neurons in the area indicated with the letter A that is experienced when the robot reaches the corresponding area of the double T-maze environment (see Fig. 2). The movements toward the attractor points corresponding to the sensory states experienced later on produce the following path of the limit cycle until a new lap of the environment and a new path in the limit cycle starts.

The fact that the fixed point attractors are never fully reached by the state of the two internal neurons while the robot navigates in the environment, since the corresponding sensory states are experienced only for a limited amount of time and since internal neurons have time constants that cause that they change their state relatively slowly, ensures that the internal states tend to preserve information about previously experienced sensory states.

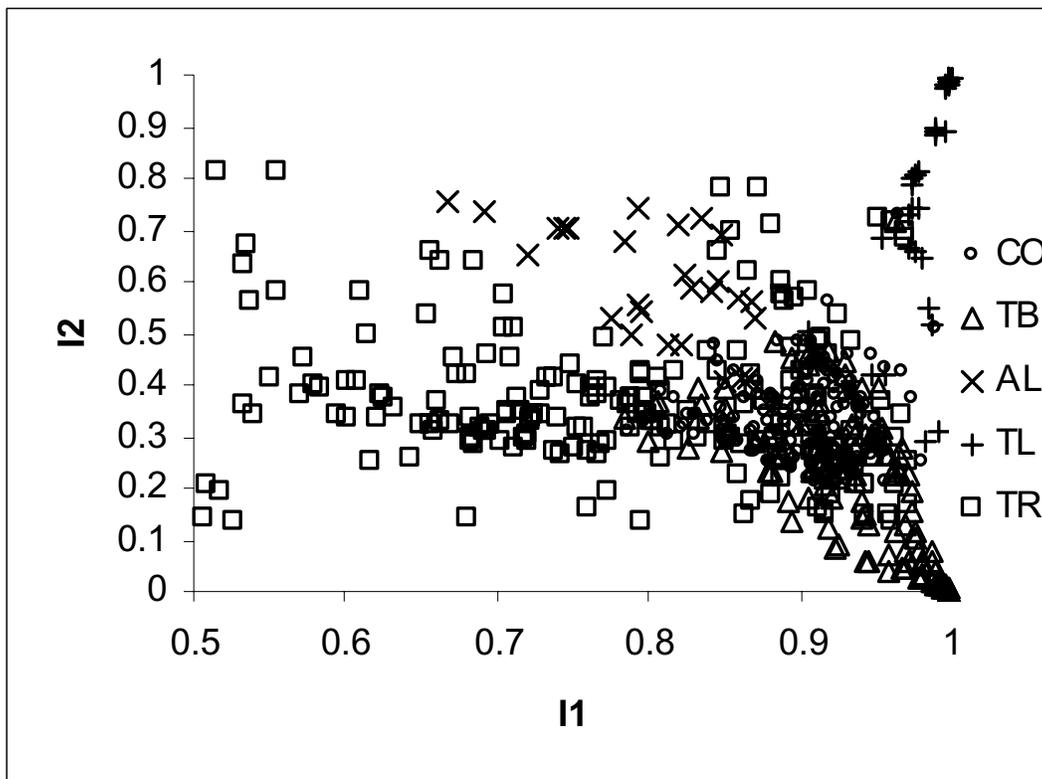


Fig. 7. Fixed point attractors in the internal neurons state space corresponding to all sensory states experienced by the robot during one lap of the environment. The fixed point attractors have been generated by fixing the state of the sensory neurons and by updating the state of the internal neurons for 150s. Fixed point attractors for different classes of sensory patterns are visualized with different symbols. Classes include the sensory status experienced by the robot while it: (TL) move toward the light, (AL) moves away from the light, (TR) turns right, (TB) turns backward, (CO) moves along a corridor that is not illuminated.

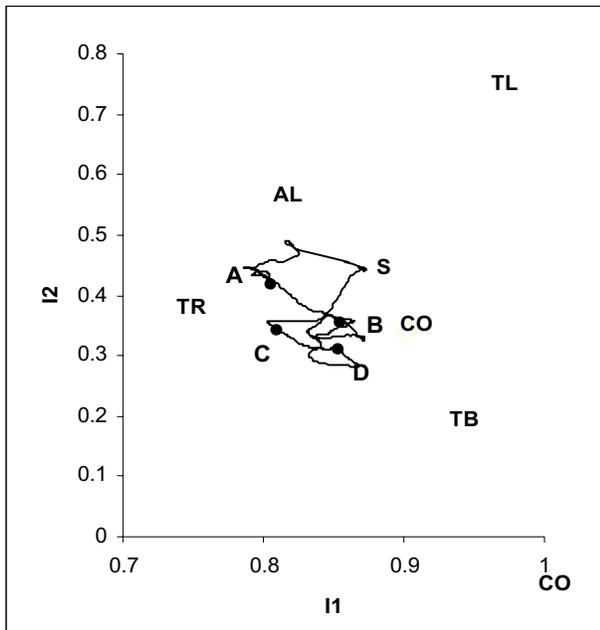


Fig. 8. Internal neurons state space including the averaged limit cycle (some data of Fig. 6), the indication of the state experienced by the robot in location A, B, C, D, and S (Fig. 2), and the average fixed point attractors corresponding to the five categories of sensory states (based on the same data shown on Fig. 7).

#### VI. ON ROBOTS' ABILITY TO REPRESENT DIFFERENT SPATIAL STRUCTURES

In this section we analyse how robots generalize their abilities when tested in environments that have different spatial organizations from the environment in which they have been evolved. As we will see, the dynamics of the internal neurons lead to different limit cycles in different environments. However, the fact that the limit cycles preserve the characteristics of periodicity and expressiveness described above ensures that the robots are able to generalize their skills in different environmental conditions.

To verify robots generalization abilities and to analyse the dynamics of the internal neurons in different environmental conditions we tested the same individual shown in Fig. 5, 6, 7, and 8, that has been evolved in the double T-Maze environment, in a simple T-maze environment and in a single/double T-Maze environment (Fig. 9). As shown in Fig. 10, robots evolved in the double T-Maze environment show rather good performance also when tested in the two new environments.

As shown in Fig. 11, the shape of the limit cycles changes in different environments. However, the fact that the limit cycles obtained in the new environments also display periodicity and expressiveness allows the robots to generalize their ability to self-localize and to recognize previously visited target areas also in environments never experienced before.

The shape of the limit cycle observed in a particular environment is the result of two factors: (1) the fixed point

attractors that determine the way in which the state of the internal neurons changes for any given sensory state, and (2) the sequence of sensory state experienced by the robot while it moves in the environment. The first factor is a function of the characteristics of the neural controller and therefore does not change when the robot is situated in a new environment. The second factor, however, obviously depends on the characteristics of the environment in which the robot is placed.

This implies that robots can generalize their abilities in a wide range of environments with different spatial organizations providing that the local characteristics of the new environment are similar to that of the environment in which they have been evolved (i.e. in this particular case, providing that the new environment has light gradients, corridors, and dead-ends that produce a variation in robots internal states) and providing that the size of the environment is comparable to that of the environment in which they have been evolved. This latter point can be explained by considering that the time constants that determine the rate at which the state of the internal neurons vary toward the corresponding fixed point attractors should be sufficiently small to avoid reaching the attractor points and sufficiently high to produce enough variation in the state of the internal neurons while the robot moves in different locations of the environment.

Interestingly, while the state of the internal neurons can be described as a representation of robot's relative position in the environment, the limit cycles consisting of the sequences of internal states experienced by the robot while it navigates in the environment, can be described as a representation of the spatial organization of the environment. In fact, different environments correspond to different limit cycles. Moreover, the level of similarity between two environments or between two sub-parts of two environments are reflected in the level of similarities of the corresponding limit cycles or of the corresponding sub-sections of the limit cycles (see Fig. 6 and 11).

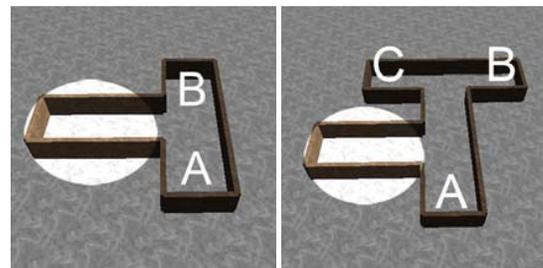


Fig. 9. Left: The T-Maze environment. Right: The single/double T-maze environment.

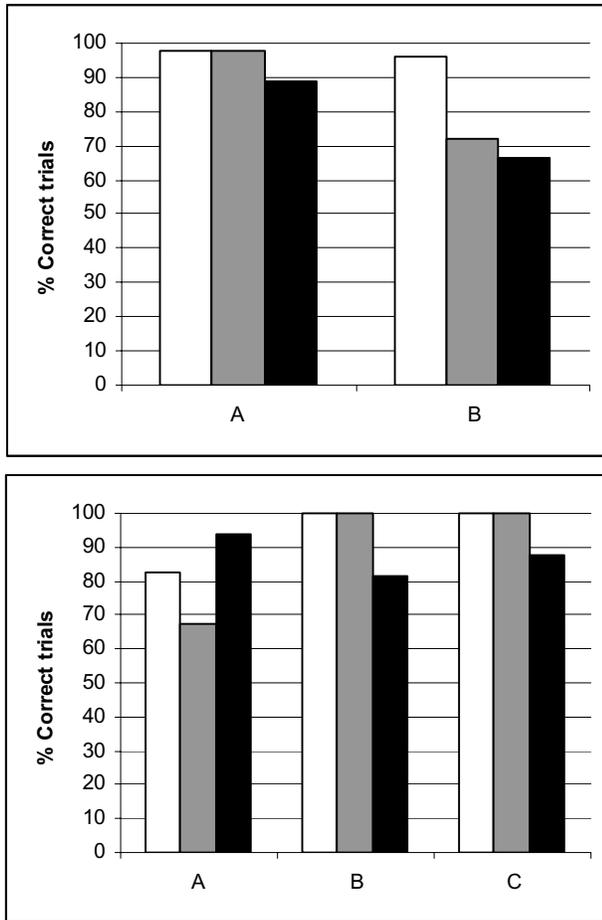


Fig. 10. Percentage of trials in which the robot is able to return and to stop on the target area after the black disk indicating the location of the area was removed. The groups of histograms indicate the performance for the corresponding location of the target area (see Fig. 8). White histograms indicate the performance obtained in simulation by initially placing the robot close to the light. Grey histograms indicate the performance obtained in simulation by placing the robot in a randomly selected location and orientation in the environment. Black histograms indicate the performance observed on the real robot in the real environment by placing the robot in a randomly selected location and orientation in the environment. Top: data for the T-Maze environment. Bottom: data for the single/double T-Maze environment.

## VII. DISCUSSION

In this paper, we described how evolved robots provided with a form of continuous time recurrent neural controller [1, 6] solve a problem that requires to identify their own relative location in the environment and the location of a target area so to be able to recognize the target location later on.

The analysis of the obtained results indicates that the coupling between the robot's internal dynamics and the robot/environmental dynamics leads to limit cycles in the state space of the internal neurons. These limit cycles are characterized by expressiveness (i.e. by the fact that different

states correspond to different relative location of the robot in the environment) and periodicity (i.e. by the fact the sequence of states of the internal neurons corresponding to a limit cycle correspond to a complete lap of the robot in the environment).

More precisely the observed limit cycle dynamics is the result of a transient dynamical process based on the alternation of fixed point attractors, with different time durations, that alternate in time while the robot moves in the environment. The characteristics of the transient dynamics and the characteristics of the observed limit cycles allow evolved robots to generalize their capacity in a wide range of environmental conditions.

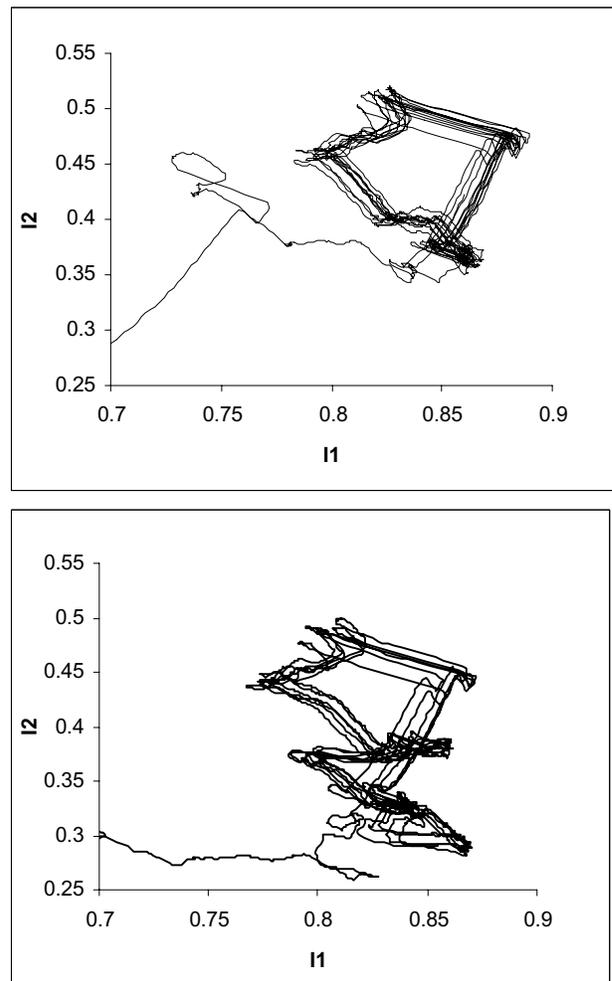


Fig. 11. The activation state of the two internal neurons recorded while the robot moves in the is environment for 330 s. Top: data for the T-Maze environment. Bottom: data for the single/double T-Maze environment.

More generally, the analysis of the obtained results indicates how evolved robots display an ability to represent different locations in the environment and to represent different environments (i.e. environment with different spatial organizations). The characteristics and the mechanisms that lead to these representation capabilities, however, have a

specificity with respect to other models described in the literature that is worth noting. We are using the term representation to indicate a robot's internal state which covaries with the robot's location in the environment or with the characteristics of the environment in which the robot is situated. Such specificity consists in the fact that the way in which a location of the environment and the overall structure of the environment are represented in our robot is implicit rather than explicit as in other models like metric or topological maps approaches in which the geometrical features of the environment or the spatial relations between relevant locations of the environment are represented in 2D maps or in topological graphs [4]. In our case, instead, both the representation of spatial locations and the representation of the relations between different spatial locations are implicit in the sense that are generated while the robot navigates in the environment and do not exist before that.

One important implication of the difference between implicit and explicit representation forms is that, as we have shown in this paper, implicit representation forms can allow robots to generalize their ability in new environments immediately without the need to acquire an explicit representation of the new environment. In the case of implicit representation forms, in fact, the new representation corresponding to the new environment is generated immediately as soon as the robot moves in the new environment without the need to change the free parameters of the system.

In future work we plan to investigate: (1) the mechanisms with which robots converge toward the limit cycle while they move in the environment after being placed in a randomly selected position and orientation, (2) the possibility to evolve robots able to return toward the target location by selecting the shortest path, (3) the relation between the topology of the limit cycles, the spatial structure of the environment, and the robots' motor behaviour.

#### ACKNOWLEDGMENT

This work was supported by the ECAgents project, a project funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant 001940. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication.

#### REFERENCES

- [1] Beer, R.D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* 72:173-215.
- [2] Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11, 209-243.
- [3] Floreano D. & Mondada F. (1996) Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics--Part B: Cybernetics*, 26(3):396-407.
- [4] Meyer J-A., Filliat D. (2003). Map-based navigation in mobile robots: II. A Review of map-learning and path planning strategies. *Cognitive Systems Research*, 4: 283-317.
- [5] Mondada R., Franzi E. & Jenne P. (1993). 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.
- [6] Nolfi S. (2002). Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. *Connection Science* (14) 3:231-244.
- [7] Nolfi S. (2005). Categories formation in self-organizing embodied agents. In H. Cohen & C. Lefebvre (Eds), *Handbook of Categorization in Cognitive Science*. Elsevier.
- [8] Nolfi S. (2005). Behaviour as a complex adaptive system: On the role of self-organization in the development of individual and collective behaviour. *ComplexUs*, 3-4: 195-203.
- [9] Nolfi S. & Floreano D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books.
- [10] Nolfi S. & Marocco D. (2001). Evolving robots able to integrate sensory-motor information over time. *Theory in Biosciences*, 120:287-310.
- [11] Slocum, A. C., Downey, D. C., & Beer, R. D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat, & S. Wilson (Eds.), *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press.
- [12] Trullier O., Wiener S.I., Berthoz A., Meyer J-A. (1997). Biologically based artificial navigation systems: Review and prospects. *Progress in Neurobiology* 51:483-544.
- [13] Vickerstaff, R. J., and Di Paolo, E. A., (2005). An evolved agent performing efficient path integration based homing and search. *Proceedings ECAL 2005*, Springer.
- [14] Vickerstaff, R. J., and Di Paolo, E. A., (2005). Evolving neural models of path integration. *Journal of Experimental Biology*, 208: 3349-3366.