Place cognition and active perception: a study with evolved robots

Orazio Miglino [a] , Michela Ponticorvo [a] & Paolo Bartolomeo [b]

[a] Department of Relational Sciences , University of Naples
'Federico II' , Naples, Italy

[b] INSERM UNIT 610, Assistance Publique-Hopitaux de Paris, Hôpital
de la Salpetriere , Paris, France
Published online: 24 Mar 2009.

PLEASE SCROLL DOWN FOR ARTICLE

Place cognition and active perception: a study with evolved robots

Orazio Miglino[a], Michela Ponticorvo[a]* and Paolo Bartolomeo[b]

[a] Department of Relational Sciences, University of Naples 'Federico II', Naples, Italy; [b] INSERM UNIT 610, Assistance Publique-Hopitaux de Paris, Hôpital de la Salpetriere, Paris, France

A study of place cognition and 'place units' in robots produced via artificial evolution is described. Previous studies have investigated the possible role of place cells as building blocks for 'cognitive maps' representing place, distance and direction. Studies also show, however, that when animals are restrained, the spatial selectivity of place cells is partially or completely lost. This suggests that the role of place cells in spatial cognition depends not only on the place cells themselves, but also on representations of the animal's physical interactions with its environment. This hypothesis is tested in a population of evolved robots. The results suggest that successful place cognition requires not only the ability to process spatial information, but also the ability to select the environmental stimuli to which the agent is exposed. If this is so, theories of active perception can make a useful contribution to explaining the role of place cells in spatial cognition.

Keywords: place cognition; evolved robots; evolutionary robotics; embodied cognition; situated cognition; active perception

1. Introduction

Place cognition is fundamental to the adaptive success of higher organisms. It is not surprising, therefore, that researchers have dedicated considerable effort to understanding the underlying neural and functional processes. Among theories of place cognition, we can distinguish between those that posit an allocentric representation of space and those that assume an egocentric representation. From Tolman (1948) to Wang and Spelke (2002), many supporters of the allocentric view have supported the concept of 'cognitive maps'. The idea of a map inside the head has been very successful, perhaps because of the power of the metaphor that helps to simplify reasoning about spatial cognition. Cognitive maps have been ascribed to humans (e.g. Tolman 1948; Herman, Miller and Shiraki 1987; Coucelis, Golledge, Gale and Tobler 1987; Baker 1989; Péruch, Firaudo and Gärling 1989; Gärling 1990; Gallistel 1990), to dogs, rats, chimpanzees (e.g. Tolman 1948; Menzel 1973; O'Keefe and Nadel 1978; Thinus-Blanc 1987; Gallistel 1990), to birds (e.g. Wallraf 1974; Gould 1982; Baker 1984; Wiltschko and Wiltschko 1987; Gallistel 1990) and to insects (Gould 1986; Gallistel 1989, 1990; Poucet 1993).

---

*Corresponding author. Email: michela.ponticorvo@unina.it

It has been proposed that one of the key building blocks for cognitive maps is the place cell. Place cells are neurons in the hippocampus whose rate of firing depends strongly on the rat's location in the environment. When the rat is in the area associated with a specific 'place field', neurons in the field fire faster than other cells. In O'Keefe's words, a place cell is 'a cell which constructs the notion of a place in an environment by connecting together several multi-sensory inputs each of which can be perceived when the animal is in a particular place in the environment' (O'Keefe 1979, p. 425). Complementary information is provided by 'head direction cells' (Taube 1998), which fire only when the animal's head is pointing in a specific direction. Recently discovered 'grid cells' in dorsocaudal medial entorhinal cortex (dMEC) (Fyhn, Molden, Witter, Moser and Moser 2004; Hafting, Fyhn, Molden, Moser and Moser 2005) fire when the animal occupies any one of the vertices of a grid overlaid on the surface of its environment. Together, these findings support the hypothesis that the mammalian brain contains topographic neural maps, representing the spatial environment occupied by the animal.

What, though, is the real nature of this representation? It is true that patterns of firing in 'place cells' and 'grid cells' provide an unambiguous representation of the animal's position in space and would thus be ideal building blocks for topographical maps. All that would be required to construct such maps would be an ensemble of cells, in which each cell is associated with a specific region of space. In such a system the firing of these cells would provide an allocentric representation of the environment. Computational models of place cells and the way they might emerge provide evidence supporting this view (Zipser 1985; Sharp 1991; Shapiro and Hetherington 1993; O'Keefe and Burgess 1996). However, animal experiments have shown that when rats are restrained their place cells completely or partially lose their specificity. In particular, Foster, Castro and McNaughton (1989) have compared location-dependent firing in freely moving and restrained rats, showing that movement and preparation for movement play a strong role in determining spatially selective activity of hippocampal place cells. In a similar vein, Gavrilov, Wiener and Berthoz (1998) have investigated the role of location, movement and directional-selectivity as rats are passively carried around the environment by a mobile robot. In Foster et al.'s study (1989), recordings were made in a small set of locations whose 'place fields' had already been identified in previous work. In Gavrilov et al.'s work (1998), by contrast, the robot (and the rat) follows standard trajectories beginning in one corner of a square arena, and traversing an area much broader than the area associated with the place field being recorded. The results show a graded response to the rat's position. In other words, they do not show either the complete loss of selectivity Foster et al. (1989) observed in restrained rats, or the sharp responses observed in studies with freely moving animals.

These results seem to show that the function of place cells depends on the animal's active behaviour in its environment. We suggest that this effect does not depend directly on the movement of the animal, or on movement preparation, but on its active perceptions as it moves through the world. When it moves, an animal modifies the inputs it receives from the environment and these modified inputs affect the neural responses that result in the next action. By choosing its actions, the animal actively *selects* the sensorial input it receives from the environment. This implies that when it moves freely there are certain combinations of sensorial inputs to which it is never exposed. When the animal is artificially exposed to these inputs, as in the experiments described earlier, it is possible that the place cells will fail to display their usual sensitivity. This is evidence that the animal is using an egocentric rather than an allocentric representation of space and that this representation is part of a process of *active perception*.

Although largely absent in place cell research, the concept of 'active perception' is supported by a large body of theory. For instance, Bajcsy (1988, p. 996) argues that '. . . perception is not passive, but active. Perceptual activity is exploratory, probing, searching; percepts do not simply fall onto sensors as rain falls onto ground. We do not just see, we look. And in the course, our pupils adjust to the level of illumination, our eyes converge or diverge, we move our heads or

change our position to get a better view of something. This adaptiveness is crucial for survival in an uncertain and generally unfriendly world'. Ballard (1991) and O'Regan and Noë (2001) suggest that the external world is scaffolding for perception whose function is to extract regularities in sensory–motor interactions. Some authors (cf. Scheier and Pfeifer 1995; Pfeifer and Scheier 1999; Pfeifer and Bongard 2006) conclude that perception is inseparable from action.

Testing the role of active perception in natural organisms is, however, difficult. In particular, it is virtually impossible to control for all possible sources of information (visual, olfactory, vestibular, proprioceptive, sense of direction, internal compass, etc.). An alternative is to use artificial organisms whose input patterns are under the complete control of the experimenter, and whose internal organisation can be analysed at a level of detail difficult to obtain with more traditional methods (Langton 1995). One way of building such organisms is through evolutionary robotics (Nolfi and Floreano 2000). In embodied and situated agents, such as evolved robots, behaviour depends on the dynamic interaction between the control system, the body and the environment (cf. Beer 1995). Nolfi and Marocco (2002) have shown that the way artificial agents categorise different objects depends on active perception.

Several models have focused on the emergence of spatial cognition in artificial agents. For example, Treves, Miglino and Parisi (1992) have evolved simple agents based on artificial neural networks, which can perform navigation tasks similar to those used in rat studies. Although these agents lack a specific neural architecture for spatial cognition and any form of activity-dependent plasticity, they develop computational units with response characteristics similar to those of place cells in rat hippocampus. This suggests that the emergence of space cells is a 'natural' adaptive solution to a simple spatial task.

It is this work that is built on in the study presented here. The basic hypothesis is that, at a functional level, place cognition does indeed depend on a 'cognitive map' in the brain but that the correct functioning of this map requires a navigational strategy that pre-selects the sensory stimuli to which the brain is exposed. In other words, perception is *active*. To test this hypothesis, we 'evolve' robots whose control system includes an artificial 'place unit' whose level of activation is highly sensitive to the robot's location in the environment. We conduct experiments in which the position and orientation of the robot are determined by the experimenter rather than by the robot itself, as in the work by Foster et al. (1989) and Gavrilov et al. (1998). The behaviour of the robot's 'place unit' provides evidence for the importance of activity and 'active perception' both for place cell function and for spatial cognition in general.

2. Method

2.1. *The robot*

The experiments are based on a software simulation of Khepera, a round, 30-mm tall micro-robot with a diameter of 55 mm. Khepera is produced and sold by the K-team group (Mondada, Franzi and Ienne 1993). Eight infrared sensors, positioned around the robot's circumference, detect obstacles within a range of 3 cm. Activation of the sensor, inversely proportional to distance from the obstacle, is expressed on a scale from zero (absence of obstacles for a distance of at least 3 cm) to unity (obstacle immediately adjacent to the sensor). A round video camera, with a visual field of 270 degrees, is located on the upper part of the robot. To move, Khepera uses two wheels located on opposite sides of its body. Each wheel is controlled by a motor. The robot is equipped with a small on-board computer and a battery, facilitating relatively autonomous movement.

Artificial evolution of physical robots would require large population sizes and a huge number of generations. The usual technique is thus to simulate evolution using software that captures the physical characteristics of the robot and the environment. In the experiment described here,

a modified version of the 'Evorobot' simulator developed by Nolfi (2000) is used – created specifically for experiments with simulated populations of Khepera robots.

2.2. *The neural control system*

The robot control system consists of a two-layer, feed-forward artificial neural network in which every input neuron has a direct connection to all the output neurons (Figure 1). Each of these connections is associated with a weight. At time $t0$, when the evolutionary process begins, each of these 'weights' is assigned a random value between -1 and 1. During evolution, connection weights vary and can theoretically assume any value between $-\infty$ and $+\infty$.

Each input unit produces an output with a value between 0 and 1. Each output unit computes an 'activation level' in the range -1 to 1. This is a sigmoid, logistic function of the sum of the inputs, weighted by the value of the connections between the input and output layer. The function is continuous monotonic and is defined for all values in the range $-\infty$ to $+\infty$.

The sensory layer consists of a total of 18 units: two bias units, eight close-range landmark detectors and eight long-range detectors. The output layer consists of two motor units and one 'place unit'.

Bias units are units that are always on, that is, their level of activation is always 1. Each of the short-range detectors is connected to an infrared sensor and transfers the output from the sensor to the rest of the neural network. Each of the long-range detectors receives a filtered signal from the on-board camera. This signal represents the aggregate output from a group of pixels representing a visual field of approximately 33.75 degrees. If a single pixel in the visual field has a value of '1' (corresponding to a black object), the signal assumes a value of 1. Otherwise it takes on a value of zero (corresponding to a white object). The long-range detectors can pick up signals originating at some distance (up to 1 m) from the robot.

The behaviour of the robot is determined by the activation of the three output units. The speed of each motor is proportional to the activation of one of the output units. When activation is negative, the motor runs anti-clockwise. When it is positive, it runs clockwise. For the robot to move in a straight line, both motors have to run at the same speed in the same direction. The third output unit (the place unit) temporarily stops the robot whenever the unit's level of activation is higher than 0.5. The activation of the place unit can be interpreted as signalling the robot's perception that it has reached a target location.

Whenever the robots perform an action, this activates the input units, the signal from these units is transmitted to the output units, and these send their own signal to the motors. The complete cycle takes roughly 100 ms.
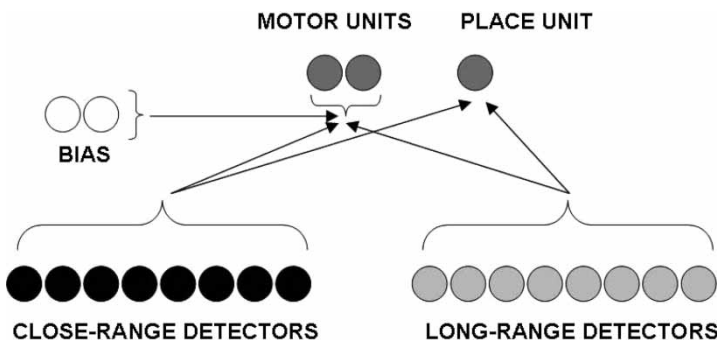

Figure 1. The robot control system: a feed-forward two-layer neural network with 18 input units (two bias units, eight close-range detectors and eight long-range detectors) and three output units (two motor units and a 'place' unit).

2.3. *Experimental task*

In our experiments, robots are trained to perform the 'blue wall' task, an experimental task used in the study of spatial behaviour (Cheng 1986; Vallortigara, Zanforlin and Pasti 1990; Hermer and Spelke 1996; Kelly, Spetch and Heth 1998; Gouteux, Thinus-Blanc and Vauclair 2001; Learmonth, Nadel and Newcombe 2002). At the beginning of the task, the subject is placed in a rectangular room, with three white walls and one blue wall that functions as a 'landmark'. The experimenter places a reward in one corner of the room. The position of the reward is unambiguously identified by the geometry of the room and the presence of the landmark. The subject is allowed to see the reward, which the experimenter subsequently hides. The task for the subject is to find it again.

2.4. *Breeding*

The goal of the breeding process is to create robots whose place units are active when the robot is in the target area and inactive when it is not. Such a unit would be functionally similar to place cells in rat hippocampus or grid cells in dMEC. A genetic algorithm was used (Belew and Mitchell 1996) to implement a process of selective 'breeding'. At the beginning of the process, a population of 100 robots was created whose neural control systems were instantiated with randomly set connection weights. For each robot, 100 trials were conducted on the experimental task. At the beginning of each trial, the robot was positioned at the centre of the room, facing in a random direction. It was then allowed to move freely for 1500 computation cycles. Each time the robot reached and 'identified' the target area (activation of place unit > 0.5) it stopped for five computation cycles (500 ms) and received one 'reward point'. The robot's final score was the sum of reward points received during the trials.

After all the robots had been tested, the 80 robots with the lowest scores were eliminated (truncation selection). Five clones were then produced for each of the remaining 20 robots. During cloning, 35% of connection weights were incremented by random values uniformly distributed in the interval $[-1, +1]$. The new neural control systems were implanted in 100 robot bodies, creating a new generation of robots. This procedure was iterated for 100 generations.

2.5. *Experiments*

Two experiments were performed, based on a version of the blue wall task first used by Sovrano, Bisazza and Vallortigara (2002) in work with *Xenotoca eiseni*, a small freshwater fish. Breeding takes place in a 568×256 cm rectangular arena with three white walls and one long black wall (the landmark). In the first experiment, the final population of robots was tested in this same environment. In the second experiment, breeding proceeded as previously. This time, however, the robots were tested in arenas whose shape and size differed from those of the arena where they were bred. The position of the landmark was also different. In this way, the experiment tests the robot's ability to generalise.

3. Results

3.1. *Experiment 1*

Evolved robots (the best performing robots in the last generation of each trial) exhibited high performance on the experimental task, successfully identifying the target corner in 99.5% of trials. Different robots use different strategies, which produce different trajectories. These can be grouped into two main categories, as illustrated schematically in Figure 2.
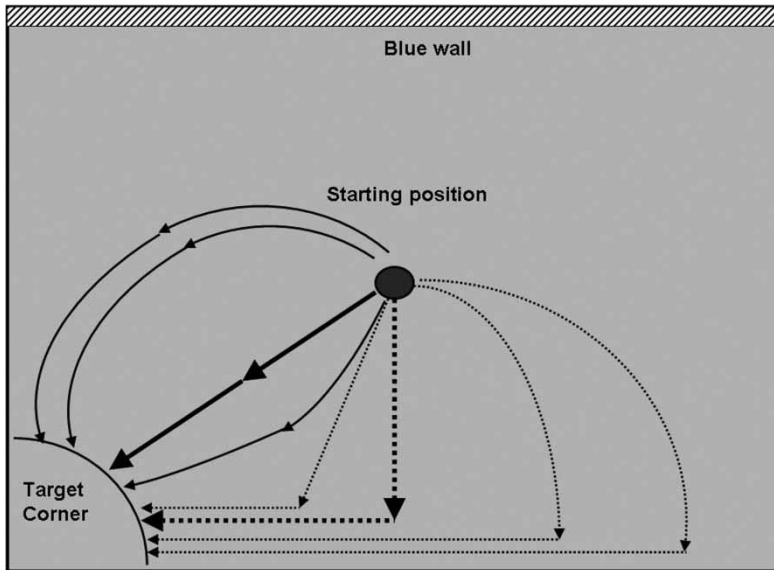
Figure 2.　Prototypical trajectories used to reach the target area; full lines represent the trajectories produced by 'distal' strategies, dotted lines represent 'proximal' (wall-following) strategies (see text for explanation). The bold arrows represent optimal trajectories.

About 40% of the robots exhibit what we term a 'distal' or 'direct' strategy: they rotate to face in the right direction and move directly from their starting point in the centre of the arena to the target area. In some cases, they move in a straight line (the optimal path); in others, they follow a curved trajectory. When they reach the target area, they stop. In the figure, these trajectories are shown as full lines.

About 35% of robots adopt a 'proximal' (wall-following) strategy. First they move towards the long white wall immediately opposite the blue wall. Sometimes they move in a straight line, sometimes in an arc. When they bump into the wall, they turn left and follow the wall until they reach the target corner, where they stop. These trajectories are shown by dotted lines.

In both cases, the robot effectively co-ordinates sensory input and motor behaviour until the task is completed. Interestingly, about 25% of robots display 'mixed' strategies. For example, they use a distal or direct strategy when they begin facing the blue wall and a proximal (wall-following) strategy when they begin facing in other directions.

To understand the processing underlying these results, the arena was divided into a grid of 11×5 cells. Using this grid, the behaviour of the best performing robot with a distal strategy and the best performing robot with a proximal strategy was investigated. Two conditions were compared. In the first ('free movement'), the activity of the place unit as the robot moved freely through the environment was recorded. In the second ('imposed position'), the robot was positioned in the first cell and the activation of the place unit was recorded using a methodology derived from Foster et al. (1989) and Gavrilov et al. (1998). Then the robot was rotated 45 degrees at a time, and the place unit activation after each rotation was recorded until the robot returned to its original orientation. This procedure was repeated for each of the 55 cells in the grid. In this way, we identified which cells a 'restrained' robot recognises as belonging to the target.

The results shown in Figure 3 show that in the 'free movement' condition the place unit fires almost exclusively when the robot is in the reward area. By contrast, in the 'imposed position' condition, the place unit loses its selectivity and is also active in other areas of the arena. In the robot with a proximal (wall-following) strategy, the unit activates along the borders of the arena,
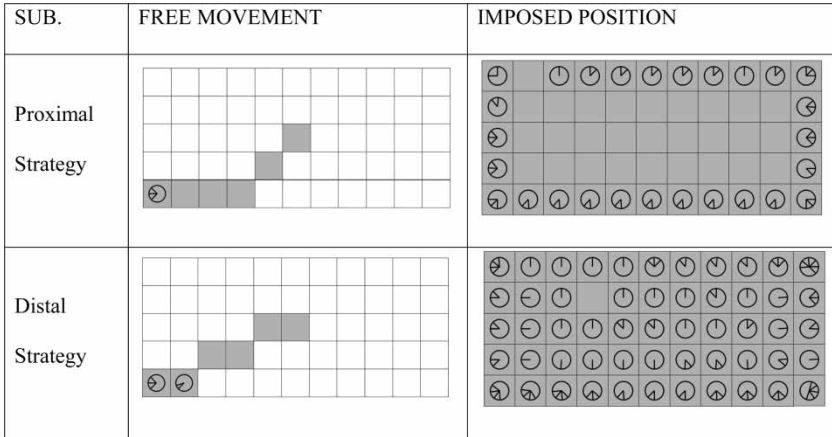
Figure 3. Maps of discrimination behaviour when the robot is allowed to move freely (left column) and when its position is imposed by the experimenter (right column). The behaviour of the robots is shown on different rows. Cells that the robot has visited are coloured in grey. Cells it has not visited are coloured white. When the robot recognises that a cell belongs to the target area, this is shown by a circle inside the cell. The line(s) radiating from the circle show the direction in which the robot is facing (0°, 45°, 90°, 135°, 180°, 225°, 270°, 315°) when the place cell fires.

and is never active when the robot is in the centre of the arena. This result holds regardless of the direction in which the robot is facing. In the robot with a distal strategy, the unit is active in almost all locations, most frequently along the walls. These findings show that, in both cases, the robot's ability to detect the reward area depends on its ability to choose trajectories that do not pass through areas where it would give incorrect results.

3.2. *Experiment 2*

To what extent can the robots generalise their ability to solve the blue wall task? To address this issue, the robots were tested in arenas whose size and shape were different from that of the arena where they were bred and in which the position of the landmark was also different. The characteristics of the test arenas are summarised in Table 1.

Figure 4 shows the percentage of trials in which the two robots successfully reached the target area and the percentage in which they reached the other corners.

The results of the experiment show that the robots perform the task well, even in arenas they have never experienced before. In other words, they acquire a general ability to perform the task in different environments. If the landmark is moved to a new location (e.g. along the short wall), the evolved robot takes account of the new position and moves towards the opposite side. This result is consistent with observations in natural organisms, showing how they can use distal visual information to cope with novel situations (O'Keefe and Conway 1978; Olton, Branch and Best 1978; Miller and Best 1980; Muller, Kubie and Ranck 1987; Hetherington and Shapiro 1997).

Table 1. Characteristics of the test arenas.

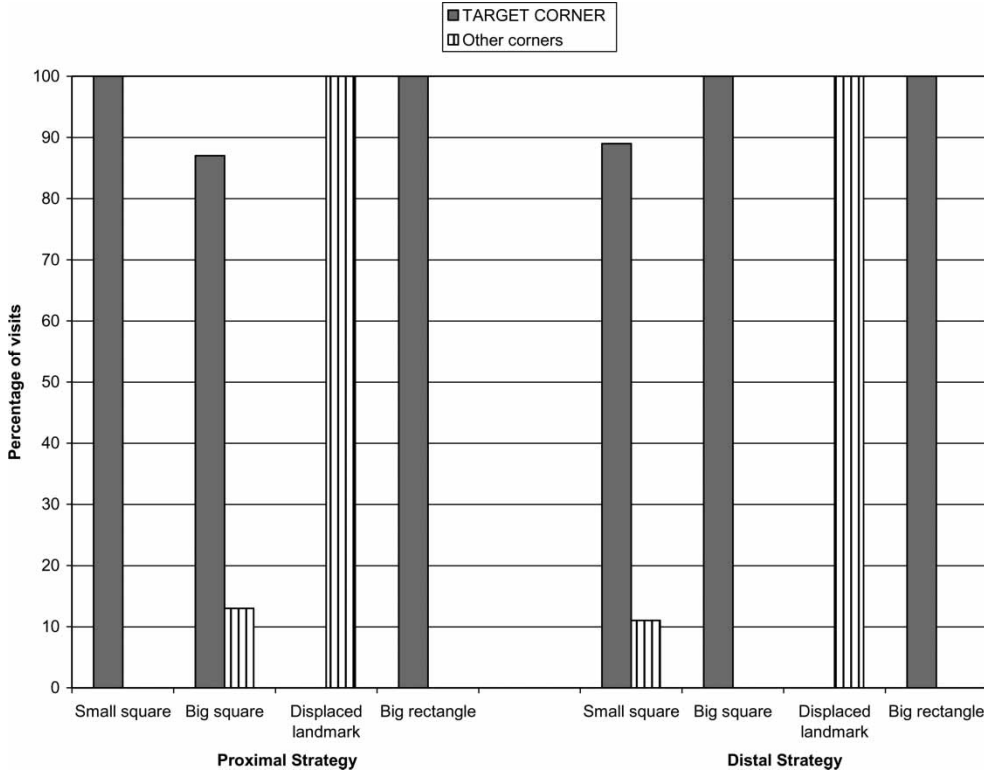| ID	Size (mm)	Shape	Landmark position
Small square environment	300 × 300	Square	Along one of the walls
Big square environment	500 × 500	Square	Along one of the walls
Displaced landmark environment	568 × 256	Rectangle	On a short wall
Big rectangle environment	852 × 384	Rectangle	On a long wall

Figure 4.   Percentage of visits to the target in the different environments used in experiment 2 (best robot adopting distal strategy, best robot adopting proximal strategy).

As in the previous experiment, the robot successfully discriminates the target area only when it is allowed to move freely (Figure 5(a–d)). When the robot is placed directly in a particular position, the place unit reacts to salient features of the arena (the walls). In the robot with a distal (direct) strategy, the place unit displays low-frequency activity over the whole arena. As in experiment 1, it loses selectivity, firing preferentially in corners and close to walls. This result is very similar to observations of place cell activity in rats, as reported by Gavrilov et al. (1998).
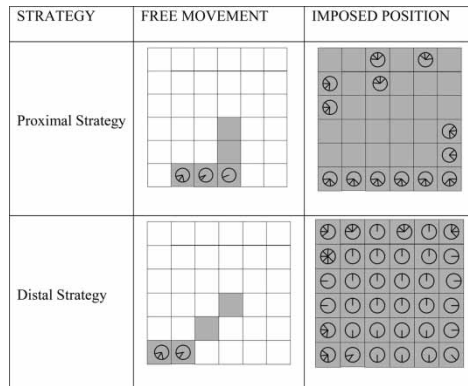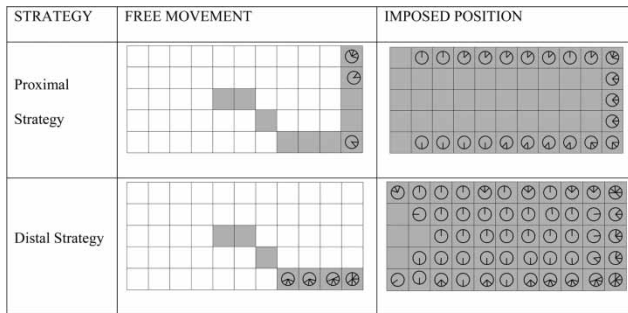

a

Figure 5a.   Maps of discrimination behaviour in the small square arena in the 'free movement' condition (left column) and the 'imposed position' condition (right column). Results for the two robots are shown on different rows. See legend to Figure 3 for explanation.
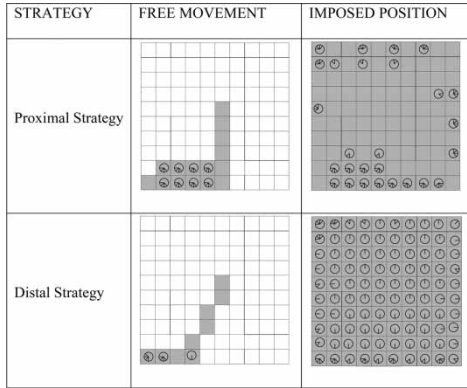
b

Figure 5b. Maps of discrimination behaviour in the big square arena in the 'free movement' condition (left column) and the 'imposed position' condition (right column). Results for the two robots are shown on different rows. See legend to Figure 3 for explanation.

c

Figure 5c. Maps of discrimination behaviour in the rectangular arena with displaced landmark in the 'free movement' condition (left column) and the 'imposed position' condition (right column). Results for the two robots are shown on different rows. See legend to Figure 3 for explanation.

d

Figure 5d. Maps of discrimination behaviour in the big rectangular arena in the 'free movement' condition (left column) and the 'imposed position' condition (right column). Results for the two robots are shown on different rows. See legend to Figure 3 for explanation.
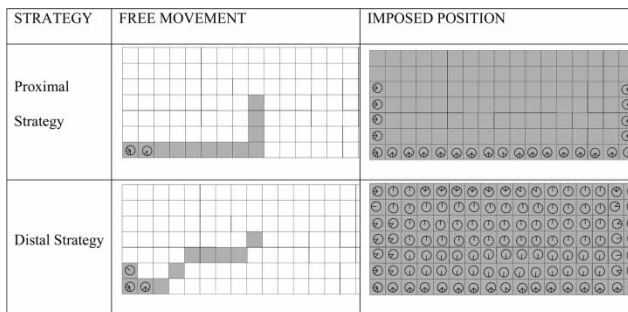
4. Conclusion

In the work described in this paper, simulated evolved robots were used to investigate the emergence of units that fire preferentially when the robot is in a specific location. These units are in some ways analogous to place cells in rat hippocampus or grid cells in dMEC. In our experiments, place units effectively discriminate the target area only when robots are free to move actively in the environment. When they cannot move freely, the place unit fires even when the robot is outside the target zone. Reaching the target seems to depend on two distinct abilities: (1) the use of environmental stimuli to distinguish the target from other regions of the arena; and (2) a navigation strategy that determines or selects the stimuli to which the robot is exposed. It is significant that the pattern of place unit activity depends on robot strategy. Robots that adopt a wall-following strategy conserve the ability to identify unambiguously walls that they need in order to implement their strategy. Robots that adopt the distal strategy lose this ability – which is not important when the robot uses distal cues.

We conclude that at least in our robots, knowledge of space cannot be isolated from action. We are aware that our model is a highly simplified representation of biological reality. For instance, the simulated place units in our robots perform a single function, although we know that biological place cells and grid cells play multiple roles. Furthermore, phenomena such as path integration and place cell remapping suggest that place cognition in animals does not rely on a simple reactive function such as the function implemented by the 'place units' in our robots. Nonetheless, this does not mean that the model is without meaning.

Place and grid cells are usually seen as a neural substrate for map-like representations. The findings reported in this paper do not contradict this view. However, they suggest that the correct functioning of these representations should be interpreted in the light of theories of 'active perception' based on 'situated knowledge' (Clark 1997; Tschacher and Scheier 1999). In sum, if we wish to understand the mounting tide of data from experiments in spatial cognition it is useful to think in terms not only of representation, but also of the coupling between representation and action.

References

Bajcsy, R. (1988), 'Active Perception,' *Proceedings of the IEEE*, 76, 996–1005.
Baker, R.R. (1984), *Bird Navigation: The Solution of a Mystery*, Sevenoaks: Hodder & Stoughton.
———. (1989), *Human Navigation and Magnetoreception*, Manchester: Manchester University Press.
Ballard, D.H. (1991), 'Animate Vision,' *Artificial Intelligence*, 48, 57–86.
Beer, R. (1995), 'A Dynamical Systems Perspective on Agent–Environment Interaction,' *Artificial Intelligence*, 72, 173–215.
Belew, R.K., and Mitchell, M. (eds.) (1996), *Adaptive Individuals in Evolving Populations: Models and Algorithms*, Reading, MA: Addison Wesley.
Cheng, K. (1986), 'A Purely Geometric Module in the Rat's Spatial Representation,' *Cognition*, 23, 149–178.
Clark, A. (1997), *Being There: Putting Brain, Body, and World Together Again*, Cambridge, MA: MIT Press.
Coucelis, H., Golledge, R.G., Gale, N., and Tobler, W. (1987), 'Exploring the Anchor-point Hypothesis of Spatial Cognition,' *Journal of Environmental Psychology*, 7, 99–122.
Foster, T.C., Castro, C.A., and McNaughton, B.L. (1989), 'Spatial Selectivity of Rat Hippocampal Neurons: Dependence on Preparedness for Movement,' *Science*, 244, 1580–1582.
Fyhn, M., Molden, S., Witter, M.P., Moser, E.I., and Moser, M.-B. (2004), 'Spatial Representation in the Entorhinal Cortex,' *Science*, 305, 1258–1264.

Gallistel, C.R. (1989), 'Animal Cognition, the Representation of Space, Time and Number,' *Review of Psychology*, 40, 155–189.

———. (1990), *The Organization of Learning*, Cambridge MA: MIT Press.

Gärling, T. (1990), 'Is Elevation Encoded in Cognitive Maps?,' *Journal of Environmental Psychology*, 10, 341–351.

Gavrilov, V., Wiener, S.I., and Berthoz, A. (1998), 'Discharge Correlates of Hippocampal Complex Spike Neurons in Behaving Rats Passively Displaced on a Mobile Robot,' *Hippocampus*, 8, 475–490.

Gould, J.L. (1982), 'The Map Sense of Pigeons,' *Nature*, 296, 205–211.

———. (1986), 'The Local Map of Honey-bee: Do Insects have Cognitive Maps?,' *Science*, 232, 861–863.

Gouteux, S., Thinus-Blanc, C., and Vauclair, J. (2001), 'Rhesus Monkeys use Geometric and Nongeometric Information during a Reorientation Task,' *Journal of Experimental Psychology: General*, 130, 505–519.

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E.I. (2005), 'Microstructure of a Spatial Map in the Entorhinal Cortex,' *Nature*, 436, 801–806.

Herman, J.F., Miller, B.S., and Shiraki, J.H. (1987), 'The Influence of Affective Associations on the Development of Cognitive Maps of Large Environments,' *Journal of Environmental Psychology*, 7, 89–98.

Hermer, L., and Spelke, E. (1996), 'Modularity and Development: The Case of Spatial Reorientation,' *Cognition*, 61, 195–232.

Hetherington, P.A., and Shapiro, M.L. (1997), 'Hippocampal Place Fields are Altered by the Removal of Single Visual Cues in a Distance Dependent Manner,' *Behavioral Neuroscience*, 111, 20–34.

Kelly, D., Spetch, M.L., and Heth, C.D. (1998), 'Pigeon's Encoding of Geometric and Featural Properties of a Spatial Environment,' *Journal of Comparative Psychology*, 112, 259–269.

Langton, C.G. (1995), *Artificial Life*, Cambridge, MA: MIT Press.

Learmonth, A.E., Nadel, L., and Newcombe, N.S. (2002), 'Children's Use of Landmarks: Implications for Modularity Theory,' *Psychological Science*, 13, 337–341.

Menzel, E.W. (1973), 'Chimpanzee Spatial Memory Organization,' *Science*, 182, 943–945.

Miller, V.M., and Best, P.J. (1980), 'Spatial Correlates of Hippocampal Unit Activity are Altered by Lesions of the Fornix and Entorhinal Cortex,' *Brain Research*, 194, 311–323.

Mondada, F., Franzi, E., and Ienne, P. (1993), 'Mobile Robot Miniaturization: A Tool for Investigation in Control Algorithms,' Paper presented at the *Third International Symposium on Experimental Robotics*.

Muller, R.U., Kubie, J.L., and Ranck, J.B. (1987), 'Spatial Firing Patterns of Hippocampal Complex Spike Cells in a Fixed Environment,' *Journal of Neuroscience*, 7, 1935–1950.

Nolfi, S. (2000), *Evorobot 1.1 User Manual*, Rome: Institute of Psychology, CNR.

Nolfi, S., and Floreano, D. (2000), *Evolutionary Robotics: The Biology, Intelligence, Technology of Self-organizing Machines*, Cambridge, MA: MIT Press.

Nolfi, S., and Marocco, D. (2002), 'Evolving Robots Able to Visually Discriminate Between Objects with Different Size,' *International Journal of Robotics and Automation*, 17, 163–170.

O'Keefe, J. (1979), 'A Review of the Hippocampal Place Cells,' *Progress in Neurobiology*, 13, 419–439.

O'Keefe, J., and Burgess, N. (1996), 'Geometric Determinants of the Place Fields of Hippocampal Neurons,' *Nature*, 381, 425–428.

O'Keefe, J., and Conway, D.H. (1978), 'Hippocampal Place Units in the Freely Moving Rat: Why they Fire where they Fire,' *Experimental Brain Research*, 31, 573–590.

O'Keefe, J., and Nadel, L. (1978), *The Hippocampus as a Cognitive Map*, Oxford: Clarendon.

Olton, D.S., Branch, M., and Best, P.J. (1978), 'Spatial Correlates of Hippocampal Unit Activity,' *Experimental Neurology*, 58, 387–409.

O'Regan, J.K., and Noë A. (2001), 'A Sensorimotor Account of Vision and Visual Consciousness,' *Behavior and Brain Sciences*, 24, 939–1011.

Péruch, P., Firaudo, K.D., and Gärling, T. (1989), 'Distance Cognition by Taxi Drivers and the General Public,' *Journal of Environmental Psychology*, 9, 233–239.

Pfeifer, R., and Bongard, T.C. (2006), *How the Body Shapes the Way We Think. A New View of Intelligence*, Cambridge, MA: MIT Press.

Pfeifer, R., and Scheier, C. (1999), *Understanding Intelligence*, Cambridge, MA: MIT Press.

Poucet, B. (1993), 'Spatial Cognitive Maps in Animals: New Hypotheses on their Structure and Neural Mechanisms,' *Psychological Review*, 100, 163–182.

Scheier, C., and Pfeifer, R. (1995), 'Classification as Sensorimotor Coordination: A Case Study on Autonomous Agents,' in *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, eds. F. Moran, A. Moreno, J.J. Merelo and P. Chacon, Berlin: Springer.

Shapiro, M.L., and Hetherington, P.A. (1993), 'A Simple Network Model Simulates Hippocampal Place Fields. 1. Parametric Analysis And Psychological Predictions,' *Behavioral Neuroscience*, 107, 34–50.

Sharp, P.E. (1991), 'Computer Simulation of Hippocampal Place Cells,' *Psychobiology*, 19, 103–115.

Sovrano, V.A., Bisazza, A., and Vallortigara, G. (2002), 'Modularity and Spatial Reorientation in a Simple Mind: Encoding of Geometric and Non-geometric Properties of Spatial Environment by Fish,' *Cognition*, 85, 51–59.

Taube, J.S. (1998), 'Head Direction Cells and the Neurophysiological Basis for a Sense of Direction,' *Progress in Neurobiology*, 55, 225–256.

Thinus-Blanc, C. (1987), The Cognitive Map Concept and its Consequences,' in *Cognitive Processes in Animal and Man*, eds. P. Ellen and C. Thinus-Blanc, The Hague: Martinus Nijhoff, pp. 1–19.

Tolman, E.C. (1948), 'Cognitive Maps in Rats and Men,' *Psychological Review*, 36, 13–24.

Treves, A., Miglino, O., and Parisi, D. (1992), 'Rats, Nets, Maps, and the Emergence of Place Cells,' *Psychobiology*, 20, 1–8.

Tschacher, W., and Scheier, C. (1999), Situated and Self-organizing Cognition,' in *The Nature of Concepts*, ed. P. Van Loocke, London: Routledge.

Vallortigara, G., Zanforlin, M., and Pasti, G. (1990), 'Geometric Modules in Animals' Spatial Representations: A Test with Chicks (Gallus gallus domesticus),' *Journal of Comparative Psychology*, 104, 248–254.

Wallraf, H.G. (1974), *Das Navigationssystem der Vögel. Ein theoretischer Beitrag zur Analyse ungeklärter Orientierungsleistungen*, München: Schriftenreihe Kybernetik, Oldenbourg.

Wang, R.F., and Spelke, E.S. (2002), 'Human Spatial Representation: Insight from Animals,' *Trends in Cognitive Sciences*, 6, 376–382.

Wiltschko, W., and Wiltschko, R. (1987), 'Cognitive Maps and Navigation in Homing Pigeons,' in *Cognitive Processes and Spatial Orientation in Animal and Man I*, eds. P. Ellen and C. Thinus-Blanc, The Hague: Martinus Nijhoff, pp. 201–216.

Zipser, D. (1985), 'A Computational Model of Hippocampal Place Fields,' *Behavioral Neuroscience*, 99, 1006–1018.