

# Power and the Limits of Reactive Agents

**Stefano Nolfi**

Institute of Psychology, National Research Council  
Viale Marx 15, 00137, Roma, Italy  
Voice:++39-06-86090231, Fax:++39-06-824737  
e-mail:nolfi@ip.rm.cnr.it  
<http://gral.ip.rm.cnr.it/nolfi>

## **Abstract**

*In this paper I will show how reactive agents can solve relatively complex tasks without requiring any internal state and I will demonstrate that this is due to their ability to coordinate perception and action. By acting (i.e. by modifying their position with respect to the external environment and/or the external environment itself), agents partially determine the sensory patterns they receive from the environment. As I will show, agents can take advantage of this ability to: (1) select sensory patterns that are not affected by the aliasing problem and avoiding those that are; (2) select sensory patterns in which groups of patterns requiring different answers do not strongly overlap; (3) exploit the fact that, given a certain behavior, sensory states might indirectly encode information about useful features of the environment; (4) exploit emergent behaviors resulting from a sequence of sensory-motor loops and from the interaction between the robot and the environment. Finally I will discuss the relation between pure reactive agents and pure representational agents and I will argue that a large variety of intermediate cases between these two extremes exists. In particular I will discuss the case of agents which encode in their internal states what they did in the previous portion of their lifetime which, given a certain behavior, might indirectly encode information about the external environment.*

**Keywords:** Sensory-Motor Coordination, Embodied Intelligence, Reactive Systems, Internal Representations, Evolutionary Robotics, Adaptive Behavior, Neural Networks

## **1. Introduction**

Recently, a new research paradigm has challenged the traditional view according to which intelligence is an abstract process that can be studied without taking into consideration the physical aspects of natural systems. In this new paradigm, research tends to stress the importance of ‘situatedness’ (i.e. the importance of studying systems that are situated in an external environment) and embodiment (i.e. the importance of study systems that have bodies, receive input from their sensors and produce motor actions as output) (Brooks, 1991; Hendriks-Jansen, 1996; Clark, 1997; Pfeifer & Scheier, 1999).

Research in this field often involves systems that are reactive or mostly reactive. These are systems in which sensors and motors are directly linked and which always react to the same sensory state with the same motor action. In these systems internal states, if any, play a comparatively limited role in determining the motor output of the system. For this reason, researchers working with systems that are embodied and situated have often been criticized

for rejecting the importance of internal representations<sup>1</sup> of the external world. As claimed by Brooks this criticism is invalid. Researchers in this field in no way deny the importance of internal representations although they tend to have a different view regarding what internal representations stand for. Within the traditional Artificial Intelligence approach, in fact, there is a tendency to view internal representations as explicit and complete representations of the external world, while those who embrace the embodied and situated approach tend to view internal representations as partial models of the world which include only those aspects that are necessary to allow agents to achieve their goals (Brooks 1991; Clark 1997) and/or, as I will claim here, as internal states that encode information about the external world in a very indirect and behavioral-dependent fashion.

The reason why most of the systems developed within the situated and embodied approach are mainly reactive is that embodied and situated systems can solve rather complex tasks without requiring internal states or internal representations. As we will show below, this ability to solve a rather large number of tasks without the need to react differently to the same sensory states in different contexts can be explained by considering that systems that are situated and embodied (agents) can coordinate perception and action.

Agents partially determine the sensory patterns they receive from the environment by executing action that modify their position with respect to the external environment and/or the external environment itself. Agents can take advantage of this ability in different ways. We will refer to the process of exploiting the agent-environment interaction (i.e., the ability to select favorable sensory patterns through motor actions) as *sensory-motor coordination* (for a similar view, see Pfeifer and Scheier 1997).

Agents which are able to rely on internal states or which can modify how they react to sensory states through learning can exploit sensory-motor coordination in additional ways. As shown in (Nolfi & Parisi, 1997), for example, agents that learn during their interaction with the environment can use sensory-motor coordination to select sensory patterns which are useful for learning purposes (i.e. sensory patterns channel learning into appropriate direction or that increase the probability of acquiring a given ability through learning).

In sections 2, 3, and 4 I will restrict my analysis to the case of pure reactive agents which always react in the same way to the same sensory pattern. In section 5 and 6 I will discuss the relation between pure reactive agents and pure representational agents (i.e. agents that have an explicit and complete representation of the external environment). In particular I will argue that a large variety of intermediate cases between these two extremes exists.

## **2. How sensory-motor coordination can cope with perceptual aliasing**

The perceptual aliasing problem is a simple example where sensory-motor coordination becomes very useful. Perceptual aliasing, a term coined by Whitehead and Ballard (1991), refers to the situation wherein two or more objects generate the same sensory pattern, but require different responses. A solution to this situation (i.e., when a agent receives a sensory pattern which requires different motor responses in different circumstances) could be that of performing actions in order to search for other sensory information that is not affected by the aliasing problem (that is, an unambiguous sensory pattern).

---

<sup>1</sup> The term internal representation is often used with different meanings. Here I will use it to distinguish between pure reactive systems (i.e. systems which react to a certain sensory state by always producing the same motor action) and non-reactive systems (i.e. systems which may react to the same sensory state in different manners). In order to react differently to the same sensory state non-reactive systems must have internal states which can vary and which determine how they react differently in different circumstances. I will refer to these states as internal states or internal representation.

Consider for example the case of a robot provided with contact sensors and a camera placed in an environment with two type of objects: one with the top painted in black, which should be avoided, and one with the top painted in white, which should be approached. Every time such a robot approaches an object which does not happen to be in the view angle of its camera, the robot will experience an ambiguous sensory pattern (i.e., a sensory pattern which is affected by the aliasing problem). The same type sensory pattern will be experienced by the robot both in the case of objects to be approached and of objects to be avoided. The obvious solution to this problem is to turn toward the object. By turning until the object is within the view angle of the camera, the robot will eventually receive an unambiguous sensory pattern. In this example the aliasing problem arise from the interaction between the sensory structure of the robot and the environment (an agents with a view angle of  $360^\circ$  will not have to turn in order to discriminate the two types of objects). In other cases, the aliasing problem might arise from the interaction between the position of the agent and the environment. In order to discriminate between  $2 \times 4$ cm and  $2 \times 6$ cm rectangular shapes for example, agents that are facing the 2cm side should move toward one of the two longer side of the object.

The process of selecting sensory patterns which are easy to discriminate through motor actions is usually referred as *active perception* (Bajcsy 1988). Examples of processes falling within this category have been also identified in natural organisms (e.g. to recognise certain visual patterns, the fruit fly *drosophila* moves in order to shift the perceived image to a certain location of the visual field, Dill et al. [1993]).

One thing that should be stressed in these examples is that, when the robot is experiencing a sensory pattern which is affected by the aliasing problem, different types of motor actions might have different effects. Some motor actions might lead the robot to finally experience an unambiguous sensory pattern (e.g. turning toward the object) while other actions may prevent this possibility (e.g. moving away from the object). Some actions might be more effective than others (e.g. in the case of the camera with a limited view angle, turning always in the same direction independently from the relative position of the object is less effective than turning toward the object. The robot will finally face the object with the camera in both cases, but it will be faster in the second case). This implies that, when facing an ambiguous sensory pattern, agents should perform the actions that maximize the chance to select an unambiguous sensory pattern.

Another point that should be made here is that we could describe this process of active vision as sensory-motor coordination. More precisely, we can identify one subpart of the robot behavior (between when the robot starts to perceive an object with the infrared sensors to when the robot perceives the object also with the camera) and describe it as a sensory-motor coordination behavior. On the other hand it is clear that if we look at the way in which the robot reacts to each sensory state there is not a clear distinction between actions that have the function to select useful sensory patterns (i.e., unambiguous sensory pattern in this example) and other actions, such as move away from an object. Each action, in fact, has always two consequences: (1) it partially determines how the agent accomplishes the desired task; (2) it partially determines the sensory patterns that the agent will experience next.

A final remark concerns the limits of this strategy to solve the perceptual aliasing problem. The strategy is effective only when the robot can find at least one sensory state that is not affected by the aliasing problem.

### **3. How sensory-motor coordination can simplify hard problems**

Another problem is the case in which groups of sensory patterns requiring different motor answers strongly overlap even if not completely. In this case, agents are facing hard

perceptual problems that might be difficult or even impossible to solve with a reactive architecture. In this section we will try to show how sensory-motor coordination can turn these hard problems into simpler ones.

The distinction between simple and hard problems has been recently formalized by Clark and Thornton (1997). They introduced the term *type-2* problems to denote hard tasks in which the problem to map input patterns into appropriate output patterns is complicated by the fact that the regularities<sup>2</sup>, which can allow such a mapping, are marginal or hidden in the sensory patterns. On the basis of this consideration they distinguished *type-2* problems (i.e., hard problems) from *type-1* problems (i.e., easy problems) in which many regularities are directly available in the sensory patterns.

*Type-2* problems, which require complex input-output mappings, may be reduced to *type-1* (tractable) problems by recoding sensory states so to enhance useful regularities (Clark and Thornton 1997). This can be obtained in two different ways. One possibility is to internally recode sensory inputs. Elman (1993), for example, showed how complex tasks that cannot be solved by training a feedforward neural network using standard back-propagation can be solved if the network is first trained on a simpler subtask and then on the entire task. This can be explained by considering that the first learning phase affects how the sensory patterns are recoded at the level of the hidden layer (Elman 1993). This recoding, by enhancing the regularities of the sensory patterns, turns the process of learning the entire task to a *type-1* problem (Clark and Thornton 1997).

However, as claimed by Scheier et al. (1998), agents can also transform a *type-2* into *type-1* problems by actively structuring their own input through sensory-motor coordination. As we will see, this strategy cannot be used by systems that are passively exposed to sensory states, that is systems which are not embodied and situated. Scheier et al. (1998) considered the case of a Khepera robot which should approach large, and avoid small, cylindrical objects located in an arena surrounded by walls (Figure 1). The robot receives sensory input from the six frontal proximity sensors.

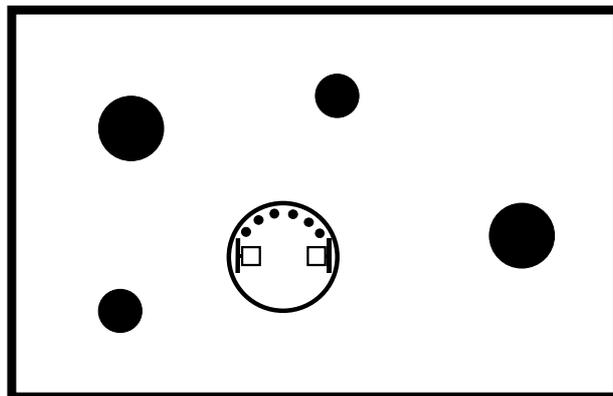


Figure 1. The robot and the environment. Small and large dark circles represent cylinders with a diameter of 1.5 and 4 cm, respectively. The lines represent the walls. The circle represents the Khepera robot.

As I reported in Nolfi (1996, 1997), the task of discriminating between the two types of objects is far from trivial for an agent that is required to passively discriminate between the sensory patterns produced by the two objects. This can be shown by training a neural network

---

<sup>2</sup> The term *regularities* refers to features of the sensory patterns which can be used to discriminate between classes of sensory patterns requiring different answers. The exact meaning of the term will become clear later on.

to discriminate between sensory patterns belonging to cylindrical objects and walls or to cylindrical objects with different diameters.

I tried 3 different architectures: (a) a two-layer feedforward architecture with 6 sensory neurons that encoded the activation of the 6 frontal infrared sensors of the robot and one output neuron which binarily encoded for the object category; (b) an architecture with an additional internal layer with four units; (c) an architecture with an additional internal layer with 8 units. In a first experiments networks were asked to discriminate between cylindrical objects with a diameter of 3 cm and walls while in a second experiment networks were asked to discriminate between cylindrical objects with a diameter of 3 and 5.5 cm respectively. Networks were trained by back-propagation (Rumelhart, Hinton & Williams 1986). During each epoch, networks were exposed to the sensory patterns perceived by the robot at 20 different distances and at 180 different angles with respect to the two objects, for a total of 7200 different patterns<sup>3</sup>.

Figure 2 shows the percentage of positions from which the networks were successfully capable of classifying the two classes of sensory patterns (i.e., to produce an activation value below or above 0.5 depending on the category of the perceived object) for each of the three different architectures. When asked to discriminate between walls and small cylindrical objects, networks without hidden units were capable of correctly classifying only 22% of the cases, whereas networks with 4 or 8 hidden units managed about 35% of the cases. When trained to discriminate between small and large cylindrical objects, the performance was even worse. Independently of their architecture, trained networks were able to correctly discriminate between small and large objects in less than 10% of the cases.

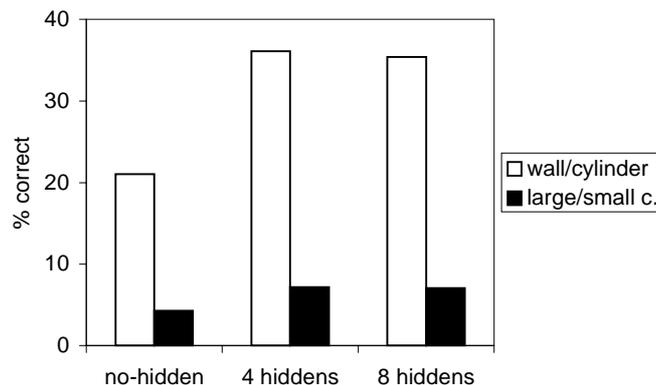


Figure 2. Percentage of positions from which sensory patterns are successfully classified by networks with no hidden units, four hidden units, and eight hidden units. Performance of the best replication out of 10. White histograms represent the performance for the wall/cylinder discrimination task, black histograms represent the performance for the small/large cylinder discrimination task. For 20 different distances and 180 different angles (3600 different positions in all) networks were asked to produce an output above or below the threshold, depending on the type of object (wall or target object). As a consequence, the total number of patterns was 7200. Responses were considered correct if, for a given position, the network was able to correctly classify both sensory patterns corresponding to the two objects.

The fact that these networks can correctly classify the two objects only in the minority of the cases can be explained by considering that variations in distance and angle are larger than variations between the two categories of objects. As a consequence, sensory patterns belonging to different categories largely overlap. “Put differently, the distance in sensor space

<sup>3</sup> The training set was presented 5000 times during learning with learning rate 0.02 and no momentum. For each architecture we ran 10 training trials starting from different randomly-assigned initial weights.

for data originating from one and the same object can be large, while the distance between two objects from different categories can be small” (Scheier et al. 1998, p. 1559).

If we look at Figure 3, that represents the positions (i.e., the combination of angle and distance) from which networks were capable of correctly discriminating (black areas) between sensory patterns belonging to *walls and small cylinders*, we see that the networks produce correct answers when the objects are not more than 120° to the left or the right of the robot face and no more than 32mm away. This result is not surprising if we consider that the robot relies only on 6 frontal infrared sensors, and that when the distance is high the infrared sensors are only slightly activated. However, there are two areas in which objects cannot be correctly discriminated even though they are “within sight”: in the figure these are the white areas enclosed by the black outer stripes. A similar pattern emerge from Figure 4 that represents the positions from which networks are able to correctly discriminate between sensory patterns belonging to *small and large cylinders* although the area from which networks are able to correctly discriminate between the two objects is significantly smaller.

It is also interesting to note that the areas in which stimuli can be correctly disambiguated are not symmetrical. This has to do with the fact that different sensors, even if identical from the electronic point of view, actually respond differently. As a consequence, it is clear that whether stimuli are ambiguous or not is a function of both the structure of the stimuli themselves and of the sensory apparatus of the robot.

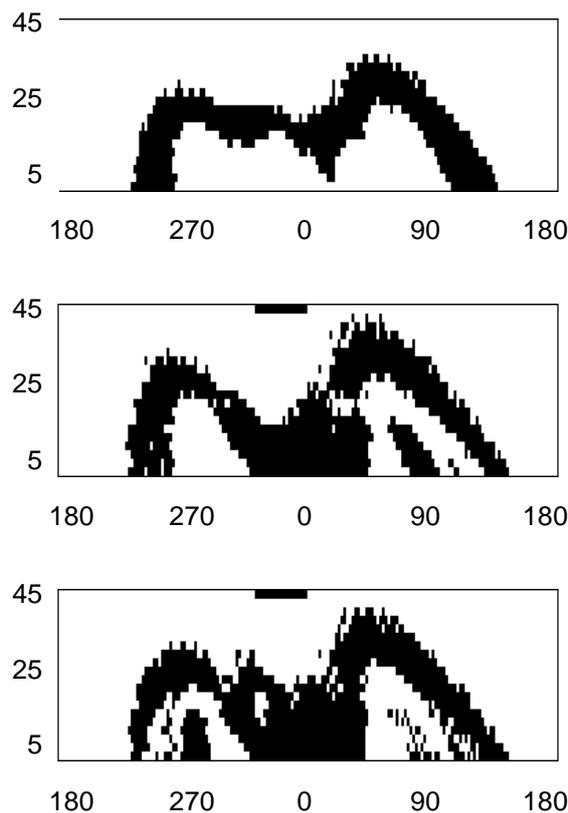


Figure 3. The areas in black represent the relative positions from which networks are able to correctly discriminate the sensory patterns belonging to walls from those belonging to small cylindrical objects. The three pictures (from top to bottom) represent the result for the best simulation with no hidden, 4 hidden, and 8 hidden units, respectively.

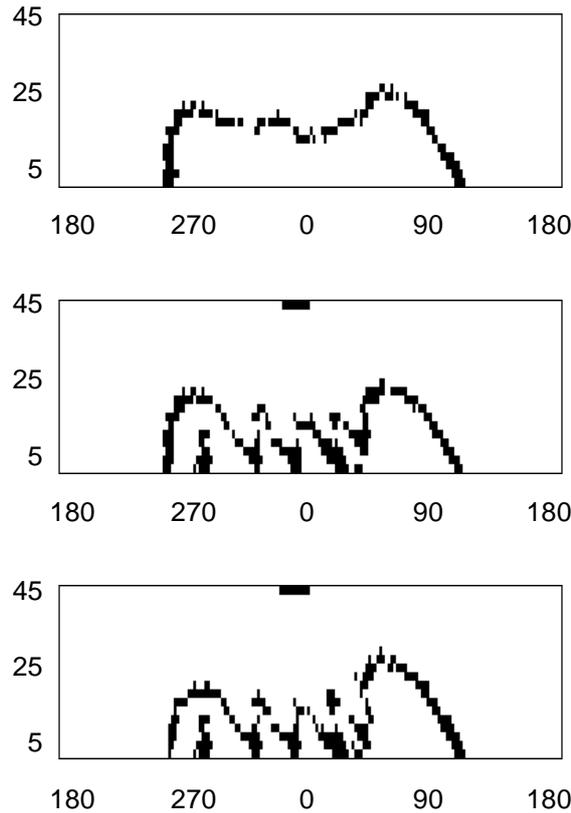


Figure 4. The areas in black represent the relative positions from which networks are able to correctly discriminate the sensory patterns belonging to large from those belonging to small cylindrical objects. The three pictures (from top to bottom) represent the result for the best simulation with no hidden, 4 hidden, and 8 hidden units, respectively.

From these results we can conclude that “passive networks” (i.e., networks which are passively exposed to a set of sensory patterns without being able to interact with the external environment through motor action), are mostly unable to discriminate between different objects. In other words, this is a type-2 problem (on this point see also below). As we will show in this and in the next section however, this problem can easily be solved by agents that are left free to exploit sensory-motor coordination.

To develop robots able to approach large and avoid small cylindrical objects (see Figure 1) Scheier, Pfeifer & Kuniyoshi (1998) used artificial evolution to select the weights of the robot's neural controllers. As the authors show, performances increase during the first generations and stabilize around near optimal performance after about 40 generations. In other words, agents which are allowed to exploit sensory-motor coordination easily solve the task. The fact that the coordination between the sensory and the motor process is crucial in solving this task can be clearly demonstrated by observing the behavior of evolved individuals and by observing how the distribution of the sensory-patterns change through out generations<sup>4</sup>.

<sup>4</sup> Evolving individuals were allowed to “live” for 5 epochs with each epoch consisting of 5000 actions. Individuals' fitness was increased at each time they were close to a large object and decreased when they were close to a small object or a wall. Connections were represented in the genotype by a 6-bit string where 1 bit determined whether the connection was to be used or not and 5 bits coded for the strength of the corresponding weight. Population size was 50. The ten best individuals of each generation were allowed to reproduce by generating 5 copies of their genotype which were mutated by replacing 5% of randomly selected bits with a new randomly chosen value. The experiment was replicated 30 times using 4 different network

As reported in Scheier, Pfeifer & Kuniyoshi (1998), the fittest individuals in 86% of the runs move in the environment until they start to perceive an object (large or small) and then start to turn around the object by circling around it (the other 14% stop in front of the objects, however, these individuals display significant poorer performances). At this point robots continue to circle around large objects while avoiding and abandoning small objects. This circling behavior is crucial to accomplish the discrimination between the two type of objects given that the sensory patterns that robots experience while circling the small objects are significantly different from those that robots experiences while circling the large objects. In other words, the sequence of motor actions that lead to the circling behavior allows the robot to select sensory patterns that can be easily discriminated.

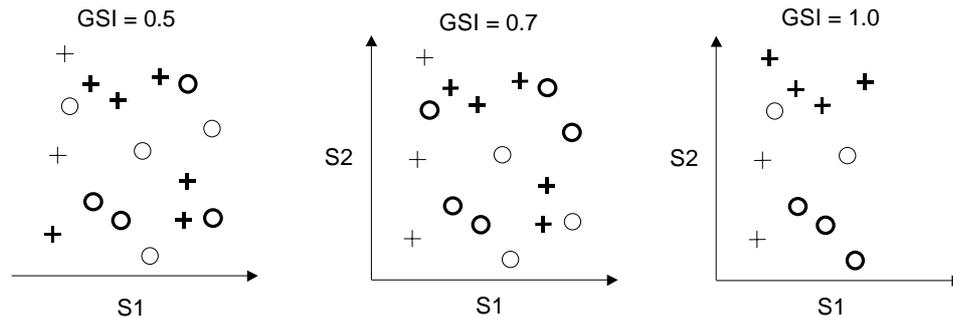


Figure 5. A schematic representation of the distribution of sensory patterns. For sake of simplicity the sensory space has only two dimensions (S1 and S2). Crosses and circles represent the sensory patterns belonging to two classes (e.g., small and large cylinders). Dark crosses and circles represent the sensory patterns experienced by an individual (notice that the sensory patterns experienced by an individual depend on how it react to previous experienced sensory states). The three figures indicate the sensory patterns experienced by three different individuals. As shown in the figure GSI vary from 0.5 to 1.0 depending on how much the two groups of sensory patterns overlap.

The role of sensory-motor coordination has been further demonstrated by analyzing the sensory-motor patterns experienced during the lifetime of individuals of successive generations. Indeed, given that the type of sensory patterns that an individual receives from the environment partially depend on how the individual reacts to each sensory state, individuals who behave differently may face harder or simpler discrimination tasks. To quantify the complexity of the discrimination task faced by different individuals, the authors measured how much the two classes of sensory patterns corresponding to the two objects (small and large cylinders) were separated in the input space. To accomplish this measure the authors used the Geometric Separability Index (GSI) proposed by Thornton (1997) to quantify the distinction between type-1 and type-2 problems introduced above. In the case of this experiment, the GSI gives a measure of the separation between the two classes and of the compactness of each class. GSI is computed by storing all sensory patterns experienced by an individual during N lifecycles by checking for every sensory pattern whether the nearest pattern (euclidean distance) belong to the same class. The total number is then normalized by N. If the nearest pattern in sensory space always belongs to the same class of the currently perceived object the GSI value is 1: this means the patterns of the two classes are well separated. Values close to 1 thus indicate that the sensory patterns belonging to the two categories are quite separated in the input space and easy to discriminate while value close to 0.5 indicate that the sensory patterns corresponding to the two categories completely overlap

---

architectures (with and without recurrent connections and with and without hidden units). Similar results were obtained for all types of architecture.

(see Figure 5). In the case of patterns corresponding to only two categories (small and large objects) this can be calculated in the following way:

$$GSI(f) = \frac{\sum_{i=1}^N (f(x_i) + f(x_{i'}) + 1) \bmod 2}{N}$$

where  $f$  is the category of the object (0 = small 1 = large cylinder),  $x_i$  are the sensory patterns consisting of  $N$  vectors, and  $x_{i'}$  is the nearest neighbor of  $x_i$  (i.e., the pattern within the  $N$  set which has the minimum euclidean distance with respect to  $x_i$ ).

As reported in Scheier, Pfeifer & Kuniyoshi (1998) the GSI value starts from about 0.5 and monotonically increases during the first 40 generations until it reaches a stable state around 0.9 (notice that also performance increases during the first 40 generations). This means that individuals of successive generations increase their ability to coordinate the sensory and motor processes so that experienced sensory patterns corresponding to one object become more similar between them and more different from sensory patterns corresponding to the other object.

A further indication that sensory-motor coordination is able to solve hard problems is that the GSI for “passive networks” (i.e., the GSI computed by taking into account all possible sensory states) is 0.564722 for the wall/small-cylinder discrimination and 0.500556 for the small/large cylinder discrimination. This implies that sensory patterns belonging to different objects overlap extensively in the first case and completely in the second case. This is a further indication that we are dealing with type-2 problems in which regularities are almost or entirely absent. Evolved agents use sensory-motor coordination to select a subset of the data containing enough regularities to discriminate sensory patterns belonging to different objects. In other words, evolved individuals are capable of transforming type-2 problems into type-1 problems through sensory-motor coordination.

#### 4. Exploiting emergent solutions

In this section we will describe another case of hard perceptual discrimination, similar to that described above, where sensory-motor coordination is used in a different way. In this case evolved individuals do not try to select sensory patterns that are not affected by the aliasing problem as in the case of section 2 (indeed, as we will see, evolution find a strategy that is effective also when all sensory patterns are affected by the aliasing problem). Similarly, evolved individuals do not try to select groups of sensory patterns that do not strongly overlap and therefore are easy to discriminate. The key mechanism is to produce discrimination as an emergent result of a sequence of sensory-motor loops.

Consider the case of a Khepera robot which is asked to discriminate between walls and cylindrical objects by finding and remaining close to the latter (Nolfi 1996, 1997). The environment is an arena of 60×35 cm surrounded by walls containing a cylindrical object with a diameter of about 2.3 cm located at a random position. To discriminate between walls and cylinders is difficult given that, as shown in the previous section, the sensory patterns corresponding to the two objects largely overlap.

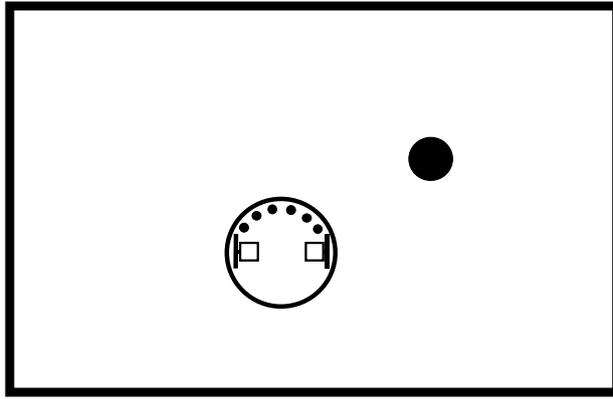


Figure 6. The environment and the robot. The full circle represents a cylinder with a diameter of 2.3 cm and a height of about 3 cm. Lines represent the walls. The large circle represents the Khepera robot.

Individual robots are controlled by neural networks with 6 sensory neurons encoding the states of the 6 frontal infrared sensors and 2 output neurons encoding the speed of the two wheels. If we use artificial evolution to select the weights of the controllers, after few generations we obtain individuals which are capable of spending most of their lifetime close to the cylindrical object (i.e., individuals which are able to discriminate between walls and cylindrical objects avoiding the former and remaining close to the latter). In particular, as shown in Figure 7 (thick line), after a few generations the best individuals are able to end up close to the cylindrical object within 500 cycles most of the times<sup>5</sup>.

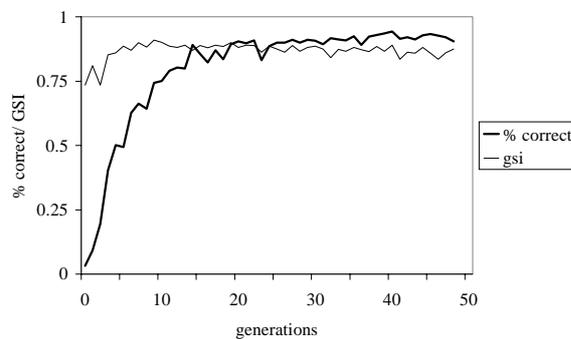


Figure 7. **Thick line:** Percentage of epochs in which the best individuals of each generation are able to end up close to the cylindrical object after 500 cycles. **Thin line:** GSI of the sensory patterns experienced by individuals of successive generations. Average results of 10 replications.

Evolved individuals do not circle around objects (as in the case of the experiment described in the previous section). On the other hand, when approaching the cylinder, evolved individuals move back and forth and/or left and right. Figure 8 displays the behavior of a typical evolved individual. The robot is placed in front of a wall on the bottom side of the environment. During the first 100 cycles it encounters and avoids the south, west, and north

<sup>5</sup> Evolving individuals were allowed to “live” for 5 epochs, each epoch consisting of 500 actions (the speed of the two wheels was updated each 100ms). Individuals’ fitness was increased at each time step when they were close to the cylindrical object. Connection weights were represented in the genotype by 8-bits strings and normalized between -10 and +10. Population size was 100. The best 20 individuals of each generation were allowed to reproduce by generating 5 copies of their genotype with 3% of their bits replaced with a new randomly selected value. The experiment was replicated 10 times.

walls. Later on, it encounters the cylindrical object and remains close to it. All individuals, like the one shown in the figure, never stop in front of the target, but begin moving back and forth, and left and right by remaining at a given angle and distance with respect to the cylindrical object.

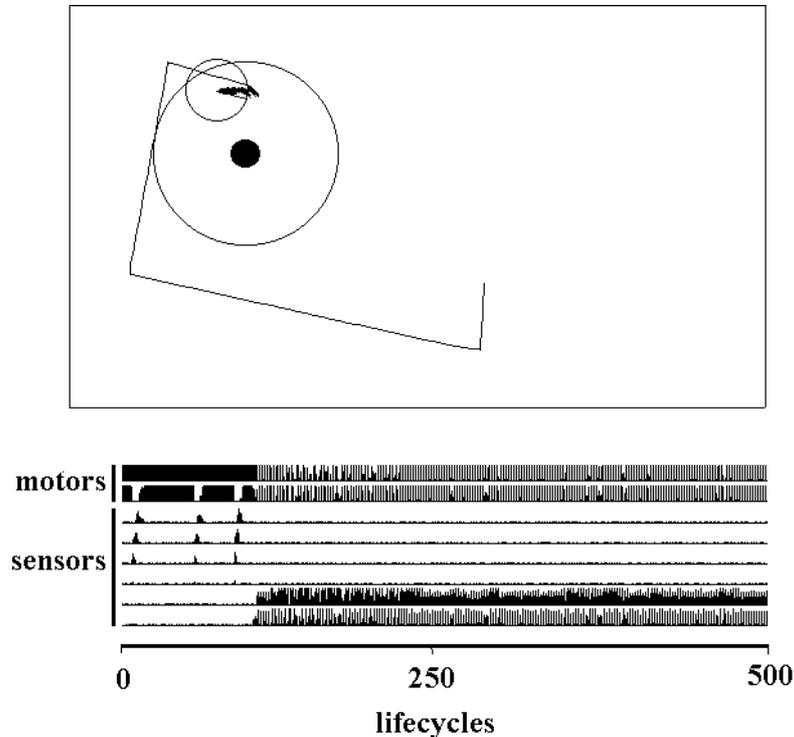


Figure 8. **Top:** Behavior of a typical evolved individual. The lines represent the walls, the full circle represents a cylindrical object, the large empty circle around it represents the area in which the robot is given fitness values, the small empty circle represents the position of the robot after 500 cycles, and the trace on the terrain represents the trajectory of the robot. **Bottom:** State of the two motor neurons and of the six infrared sensors during 500 lifecycles recorded during the trajectory displayed above.

This emergent behavior could be described as a dynamical system. The positions with respect to an object at which individuals start to oscillate while remaining in about the same position is an attractor; the robot's trajectory will always converge to these positions regardless of the direction of approach to the target. Figure 9 shows the movement displayed by an evolved individual close to walls or cylinders. When the individual reaches a distance of about 20mm from an object, it avoids walls but approaches cylinders until it reaches the attractor area located at a distance of about 15 mm and an angle of about 60 degrees. The resulting trajectories converge to the center of the area allowing the individual to keep more or less the same position relative to the cylinder.

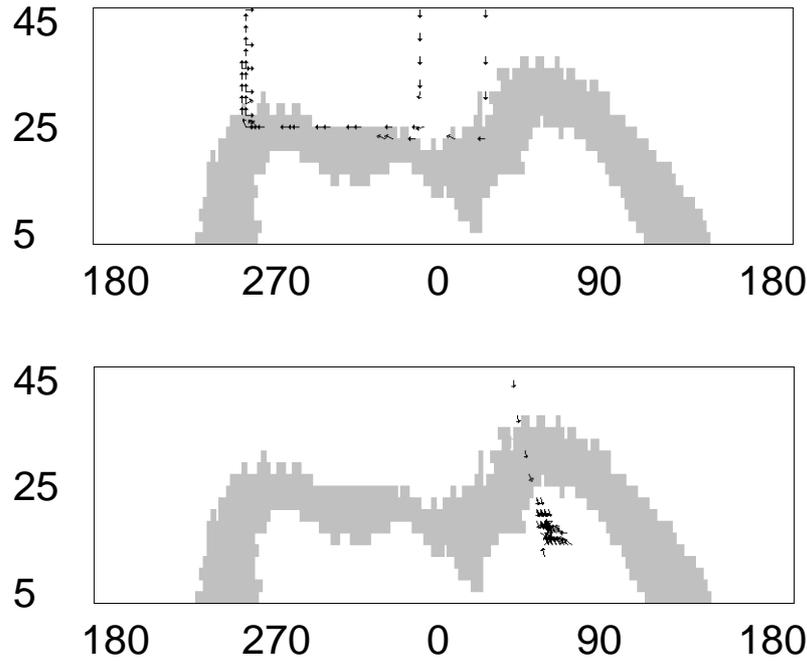


Figure 9. Angular trajectories of an evolved individual close to a wall (top graph) and to a cylinder (bottom graph). The picture was obtained by placing the individual at a random position in the environment, leaving it free to move for 500 cycles, and recording displacements at positions relative to the two types of objects for distances smaller than 45 mm. For sake of clarity, arrows are used to indicate the relative direction but not the amplitude of movements. The gray area represents the positions from which the two objects can be discriminated by a passive agent (same data of Figure 3, top).

To ascertain if these individuals, as in the case of the individuals of the experiment described in the previous section, solve the task by self-selecting sensory patterns which are easy to discriminate we calculated the GSI index for the best individuals of successive generations. As shown in Figure 7 (thin line), GSI index stabilizes at around 0.8 after the very first generations while performance keeps growing throughout 50 generations. These contrasting trends, and the fact that GSI stabilizes on a lower value with respect to the experiments described in the previous section (about 0.8 instead of 0.9), suggest that here the ability to select sensory patterns which are easy to discriminate plays a less important role. This hypothesis is also supported by the fact that the attractor area is located outside the gray area representing the relative positions from which is possible to correctly discriminate between the two type of objects without relying on sensory-motor coordination (same data shown on the top of Figure 3)<sup>6</sup>. These individuals can solve the problem of remaining close to the cylinder while avoiding walls by reacting to sensory states in a way that produce an attractor area close to the cylinder but not close to a wall.

To better clarify the role of sensory-motor coordination in this strategy it is useful to analyse a simpler case designed so to prevent the possibility to increase the separability between groups of sensory patterns belonging to different objects through sensory-motor coordination (see Figure 10). As we will see, in fact, evolved individuals solve the problem by relying on a behavioral attractor also in this case.

<sup>6</sup> The position of the attractor area varies in different replications of the experiment. Often it overlaps the border between areas in which the sensory patterns can and cannot be discriminated by a passive agent. Only in a minority of the cases the attractor area is located within the discriminative area.

Consider the case of a simulated agent which lives in a circular stripe divided into 40 cells (20 cells on the left and 20 on the right side). For each side of the environment, cells are labelled by numbers in the range 0-19 in *different random* orders (see Figure 10, left). At each time step the agent occupies one single cell and perceives the number corresponding to that cell. The agent can react to the current sensory state in two different ways: move one cell clockwise or one cell counter-clockwise. Starting at any position on the circular stripe, the goal of the agent is to reach and/or remain in the left part of the environment.

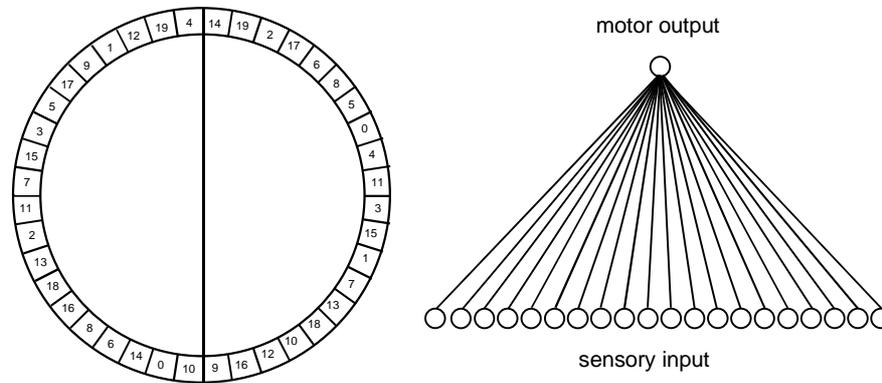


Figure 10. **Left:** The environment. **Right:** The agents' neural controller.

Agents have a neural network with 20 input units locally encoding the number of the current cell, and 1 output unit binarily encoding one of the two possible actions (see Figure 10, right). Only one sensory unit is activated each time step. Weights can assume only two values (0 or 1). Therefore, the weight from the input unit activated by the current sensory state determines the action of the agent. If the weight is 0 the output will be 0 and the agent will move clockwise; conversely, if the weight is 1 the output will be 1 and the agent will move counter-clockwise. Individuals cannot keep trace of previously experienced sensory states (i.e., they react to a certain sensory state always in the same way).

What is interesting about this experimental situation is that all sensory states are affected by the aliasing problem. Whatever the sensory state is, in fact, the agent has 50% of probability to be in the left or in the right part of the environment. In other words, it is impossible to select sensory states which are not affected by the aliasing problem. Moreover, being all sensory states equidistant in the input space, it is impossible to self-select the sensory states so that the two groups belonging to the two objects do not largely overlap.

However, if we evolve a population of such agents by selecting those who spend a larger number of cycles in the left part of the environment, after few generations we obtain individuals which are able to move away from the right part of the environment and to remain in the left part<sup>7</sup>. The way in which evolved individuals solve this problem can be seen by observing the arrows in Figure 11 which display two typical evolved strategies. In the right part of the environment individuals consistently move clockwise or counter-clockwise until they abandon the right side. On the contrary, in some areas of the left side of the

<sup>7</sup> Evolving individuals were allowed to “live” for 100 epochs with each epoch consisting of 200 actions. Each epoch individual were initially placed in a randomly selected cell of the environment. Individuals' fitness was increased each time individuals end up in the left part of the environment after 200 cycles. Connection weights were binarily represented in the genotype which was 20 bits long. Population size was 100. The best 20 individuals of each generation were allowed to reproduce by generating 5 copies of their genotype with 2% of their bits replaced with a new randomly selected value. The experiment was replicated 10 times.

environment, individuals start to move back and forth while remaining in the left side for the remaining part of their lifetime<sup>8</sup>.

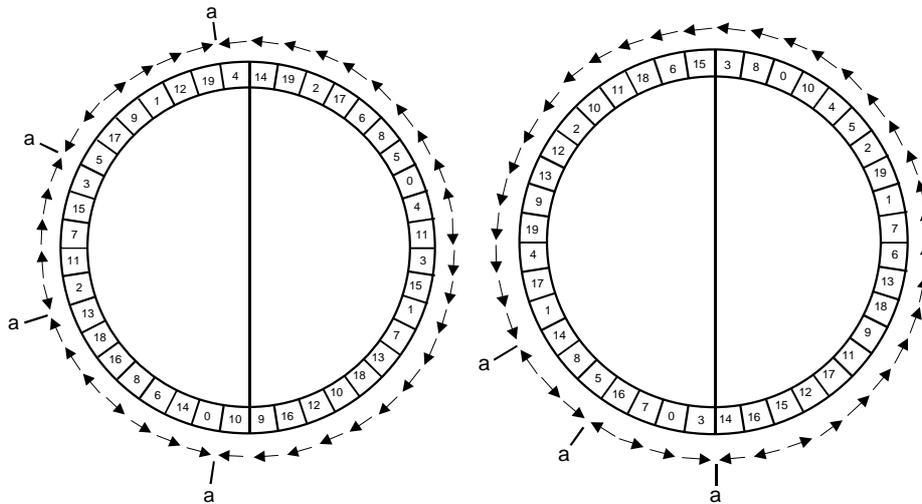


Figure 11. Two typical strategies obtained with artificial evolution. The numbers represents the sensory state experienced by the agent in each cell. Each different sensory state is present one time in the left and one time in the right part of the environment, but at different random locations. Arrows indicate the motor reaction of a typical evolved agent. Points indicated with an “a” are attractors. The two pictures represents the strategies obtained in two different replications of the experiment. In each replication the sequence of 20 cells in the two side of the environment are randomly distributed.

Agents solve their task by reacting to sensory states so to produce attractors in the left part and not in the right part of the environment. Attractors consists of two adjacent cells to which the agent react clockwise and counterclockwise (following the clockwise direction, the robot should respond clockwise to the first cell and counterclockwise to the second cell, see points indicated with an “a” in Figure 11). When the agent encounter an attractor point, it start to move back and forth while remaining there. This strategy ensures that agents will leave the right part, but will remain in one of the attractors located in the left part.

This strategy evolved by these simulated agents closely resembles the strategy evolved by the robots described above in this section. The environment has only one dimension in this experiments while it has two dimensions in the previous experiments. This explains while agents move clockwise or counterclockwise in these experiments, while robots move back and forth and/or left and right in the previous experiments. However, the type of strategy is the same: react to sensory states in order to produce attractors (i.e., set of motor actions that, taken together, allow the individual to remain in the same position) in the left part but not in the right part of the environment in this experiments; react to sensory states to produce an attractor close to cylindrical objects, but not close to walls, in the previous experiment.

It should be noted that how an individual reacts to a particular sensory state does not have any function on its own. For example the fact that the individual shown on the left side of Figure 11) reacts clockwise to the sensory state “3” does not have a function by itself. The way in which an evolved individual reacts to a certain sensory state makes sense only if we

<sup>8</sup> In the case of the individual shown in the right part of Figure 11, one of the three attractors is located at border between the two side of the environment. When the agent falls in this attractor it starts to oscillate between the left and the right side of the environment. However this does not happen often because the basin of attraction of this attractor is quite small. In other words, there are few starting positions from which the agent arrives here.

consider how it reacts to all other sensory states. Another interesting aspect to notice is that in these experiments we cannot distinguish between a phase in which the agent discriminates between the left and right part of the environment and a phase in which it moves to the desired part of the environment. As in the experiment with robot described above, the ability to discriminate and to remain in the target area are the emergent result of the same sequence of sensory-motor loops.

When we claim that a sensory pattern is affected by the aliasing problem or that two groups of sensory patterns which require different answers strongly overlap we implicitly assume a certain behavioral solution to a given task. In other words, we are assuming that the agent should react differently to patterns belonging to different objects or to different sides of the environment. However, tasks can usually be solved in a varieties of different ways and only some of these ways might present problems such as aliasing or lack of regularities between groups of sensory patterns that require different answers. Let us consider the last experiment of this section. When we first think of this task we assume that the only way to solve it is to react differently to sensory states present in the left and in the right part of the environment (at least to some of them). When we then realize that all sensory patterns are both present in the two sides and equally distant in the input space we feel that there is no way to solve the problem for a reactive agent (i.e., for an agent which do not have internal states which can keep a trace of previously experienced sensory states). However, when we observe the behavior of the evolved agent, we see that there is a completely different way to solve the problem that does not require to react differently to sensory states which lie in the two sides.

When we left individuals free to find their own way to solve a task by interacting with the external environment two different processes take place. On one hand, evolution, by selecting the most successful individuals, tend to select the individuals whose strategies are less affected by the aliasing problem or by the lack of regularities within groups of sensory patterns that require different answers. This is the case, for example, of the individuals who start to move back and forth in the left part of the circular stripe which do not need to react differently to the sensory patterns encountered in the two sides to leave the right side and to remain in the left part of the environment. On the other hand, given a certain selected strategy, evolution tend to select the individuals that exploit sensory-motor coordination to avoid sensory pattern affected by the aliasing problems and to select the sensory patterns so to increase regularities within groups of sensory patterns which require different answers. This is the case, for example, of the experiment described in the previous section in which evolutionary improvements are highly correlated with an increasing ability to select sensory patterns easy to discriminate.

## **5. Exploiting constraints**

From an engineering point of view the characteristics of the environment and the body structure of the robot are referred to as constraints and are generally viewed as problems to be solved in the design of control software. Typically, an engineer comes up with a solution to the problem based on a model of the environment and of the robot body, and then checks whether that solution works for the specific environment and robot (e.g. sensors with a certain level of precision, obstacles of a given shape arranged in certain configurations, etc.). Conversely, in artificial evolution constraints may be viewed as opportunities (Ashby 1956; Scheier et al. 1998). To better illustrate this point we will consider an experiment carried out by Miglino (Miglino 1996; see also Lund and Miglino 1998) where the author used artificial

evolution to train a Khepera robot to perform a navigation task that had been previously studied with rats.

Gallistel, Cheng and other researchers (Gallistel 1990; Cheng 1986; Margules and Gallistel 1988) used the experimental settings displayed in Figure 12 to show that rats build a map<sup>9</sup> of the environment and use it to navigate. In one of their experiments, they let rats locate and eat a visible food patch in an open field box. They then tested the rats for their ability to locate food patches in boxes identical to the one experienced previously with the food placed in the same position, but this time buried.

When the box was covered and its orientation changed randomly to prevent the use of external cues, Margules and Gallistel (1988) observed that the rats were able to correctly accomplish the task in 35% of the trials (cf. Table 1). Moreover, they observed that some of the errors made by the rats were systematic. In 31% of cases, rats looked for the food in an area that had the same geometric relationships as the correct location with respect to the rectangular shape of the box. For example, if the food was located in a corner, in about one third of the trials rats looked for the food in the correct location while in another third of the trials, they looked in the opposite corner with respect to the correct location of the food (Figure 12). The authors called these *rotational errors* because rats confused the correct location of the food with a location that would be correct if the box were rotated by 180 degrees. The remaining one third of the cases, rats looked elsewhere. According to the authors, since rotational errors can be explained only in Euclidean terms, these results would show that rats rely on an internal representation of their environment (a map) based on Euclidean geometry.

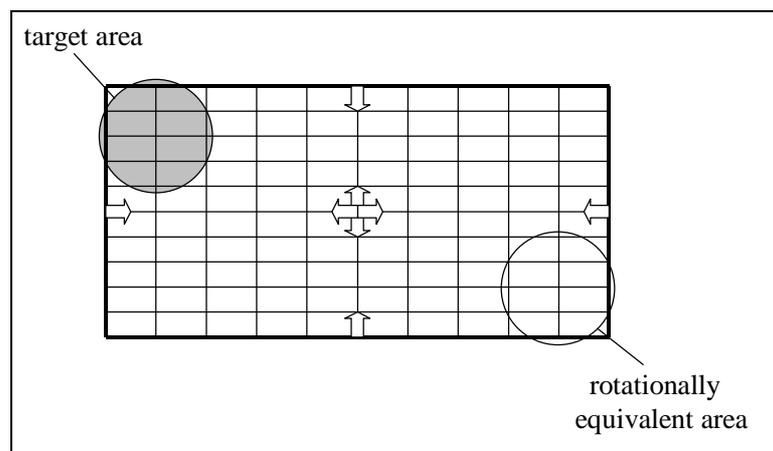


Figure 12. Top view of the open field box experiment. The thick lines represent the box (120×60 cm). The thin lines show how the box can be ideally divided into 100 12×6 cm cells. The arrows indicate the eight different starting positions and orientations used during the test phase. The gray circle represents the area in which the food is located (i.e., the target area) while the empty circle represents the rotationally equivalent area. In the experiments with robots the target area has a diameter of 15 cm and can be positioned at any intersection except for the central one. A new evolutionary run is performed for each target location. Redrawn from Miglino (1996).

Miglino set up a similar experiment in which a Khepera robot was evolved to navigate toward a target area using artificial evolution (Miglino 1996; see also Lund and Miglino 1998). The robot was placed in a rectangular 60×30 cm environment ideally divided into 100 6×3 cm cells. Each robot was allowed to move in the environment for 8 epochs, with each

<sup>9</sup> The term map or cognitive map is often used with different meanings. In this particular case it is used to indicate a two-dimensional map of the environment in which information is represented in Euclidean terms. See also below.

epoch consisting of 80 actions. For each epoch, the robot was initially placed in one of the 8 starting positions and orientations displayed in Figure 12. Robots were provided with a neural network with 8 sensory neurons encoding the state of the infrared sensors of the robot and 2 motor neurons encoding the speeds of the two wheels (which were updated every 100 ms). A new evolutionary run was repeated for every new position of the target on the 80 intersections of the lines dividing the floor into 100 cells, with the exception of the central intersection. Therefore the experiment was replicated 79 times and in each replication the position of the food was different<sup>10</sup>.

Similarly to rats, evolved robots were able to locate the food area or the rotational equivalent area in 82% of the cases and missed one of the two areas in only 18% of the cases (cf. Table 1). Note that, as in the case of rats, pure reactive robots navigated to the rotational equivalent area as many times as they did to the correct target area. It is important to clarify that these results do not imply that rats do not use internal representations to solve this type of task. In fact the same task could be solved in a variety of ways including those that rely on internal representations.

	<b>Hits</b>	<b>Rotational errors</b>	<b>Misses</b>
<b>Rats</b>	35	31	33
<b>Robots</b>	41	41	18

Table 1. Percent of navigation to the target area (hits), to the rotationally equivalent area (rotational errors), and to other areas (misses) for robots and for rats. Rat data from Margules and Gallistel (1988). Robot data from Miglino (1996).

How can purely reactive robot solve this task? One solution could be that of discriminating between long and short walls and navigate accordingly. For example, if the food is located in the upper left corner, as in Figure 12, an agent could follow a long wall on its own right side until a corner (or other similar strategies like following a short wall on its own left side until a corner). In the case of this food location, this strategy will produce 50% of success and 50% of rotational errors, depending on which of the two walls the agent happens to follow. However, a Khepera robot with a purely reactive control system and sensors that can detect obstacles up to a distance of about 5 cm, cannot discriminate between long and short walls. Once again evolution exploit sensory-motor coordination to find the solution to a hard perceptual problem.

---

<sup>10</sup> The architecture of the network was fixed and consisted of a fully connected perceptron with 8 sensory and 2 motor neurons. Individuals' fitness was increased each time individuals ended up in the target area (i.e., each time their distance from the intersection corresponding to the position of the food was less than 7.5 cm after 80 cycles). Each connection weight was represented in the genotype by a 8-bit string and normalized between -10 and +10. Population size was 100. The best 20 individuals of each generation were allowed to reproduce by generating 5 copies of their genotype with 4% of their bits replaced with a new randomly selected value.

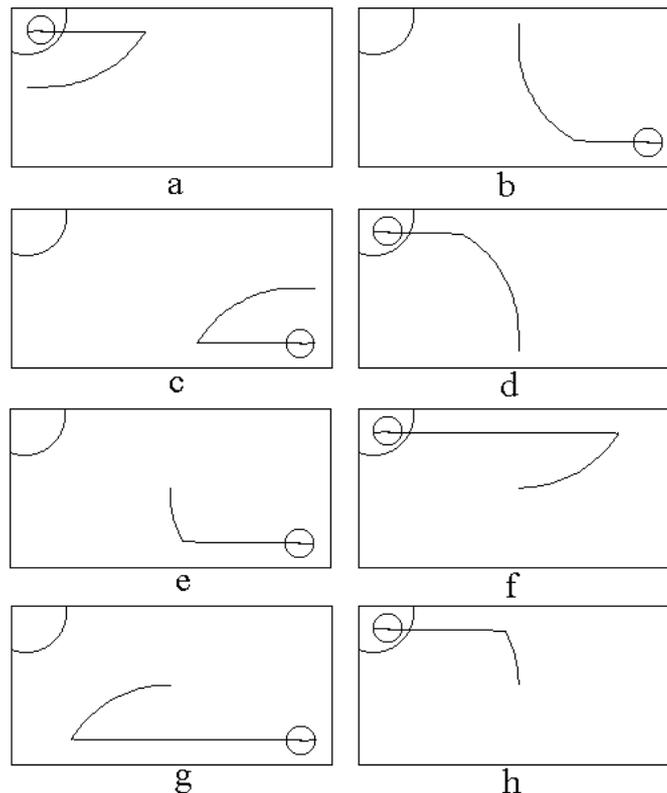


Figure 13. Behavior of a typical evolved robot. The arcs indicate the target area which is located in the upper left corner. In conditions (a), (b), (c), and (d) the robot is initially placed along one of the four walls delimiting the arena and is oriented toward the center of the arena itself. In conditions (e), (f), (g), and (h) the robot is initially placed in the center of the environment with four different orientations. The robot ends up four times in the target area and four times in the rotationally equivalent area. Data obtained from a replication of the experiment described in Miglino (1986).

Figure 13 shows the behavior of an evolved individual in an experiment in which the target area is located in the upper left corner. This individual follows one of the strategies described above (follow a long wall on the right until a corner), but it does not need to discriminate between short and long walls because it never encounters short walls. This favorable condition is achieved by producing a curvilinear trajectory when far from walls (i.e., when the sensors are all off). The result is that the robot always encounter long walls first.

This evolved strategy clearly exploits environmental constraints (i.e., the fact that the environment has certain dimensions, that individuals start from only 8 different positions, and that long walls are twice the length of short walls). If we test the same individual in a 1200×600 cm environment, in conditions (e) and (h) the robot actually produces two misses by performing a circular trajectory and never encountering any wall. Similarly, if we test the same individual in a 300×150 cm environment, in conditions (f) and (g) the robot produces two misses by encountering and following the short walls. In other words, the angular trajectory produced by the evolved robot far from walls exploits the fact that the environment has certain dimensions. This result clarifies the sense in which behavior is the result of the interaction between the robot and the environment. The ability of this robot to get to the target area, or to the rotationally equivalent area, is due both to how the robot reacts to sensory patterns and to the characteristics of the environment (its dimensions, in this case).

These agents can solve relatively complex tasks by developing strategies that exploit environmental constraints. But what happens if some of these features are allowed to vary? In other words, what happens in less constrained environments?

To investigate this issue I ran a new sets of experiments in which the environmental constraints present in the previous experiment were relaxed. Evolved robots were initially placed in a randomly selected location of the environment with a randomly selected orientation (instead of being sistematically placed in one of the eight starting positions shown in Figure 12). Moreover, the length of short walls and the proportion between long and short walls was randomly varied in each epoch (the length of short walls was randomly selected between 15 and 45 cm and the length of long walls was obtained by multiplying the length of short wall by a number randomly selected between 1.1 and 2.0). The aim of these experiments was not to compare the behavior of robots and rats. Therefore, the analysis was restricted to the case in which the food was located in the upper left corner of the environment. In this situation individuals evolved in the previous experiment reported 4 hits and 4 rotational errors out of 10 replications (see Figure 13).

If we keep all other parameters the same, obtained performance are significantly worse (see Nolfi and Floreano, in press). However, if we increase the lifetime of individuals up to 500 cycles for each epoch, evolving individuals are able to navigate to the target or rotational equivalent area in most of the cases also in environment varying in shape and size (see Figure 15).

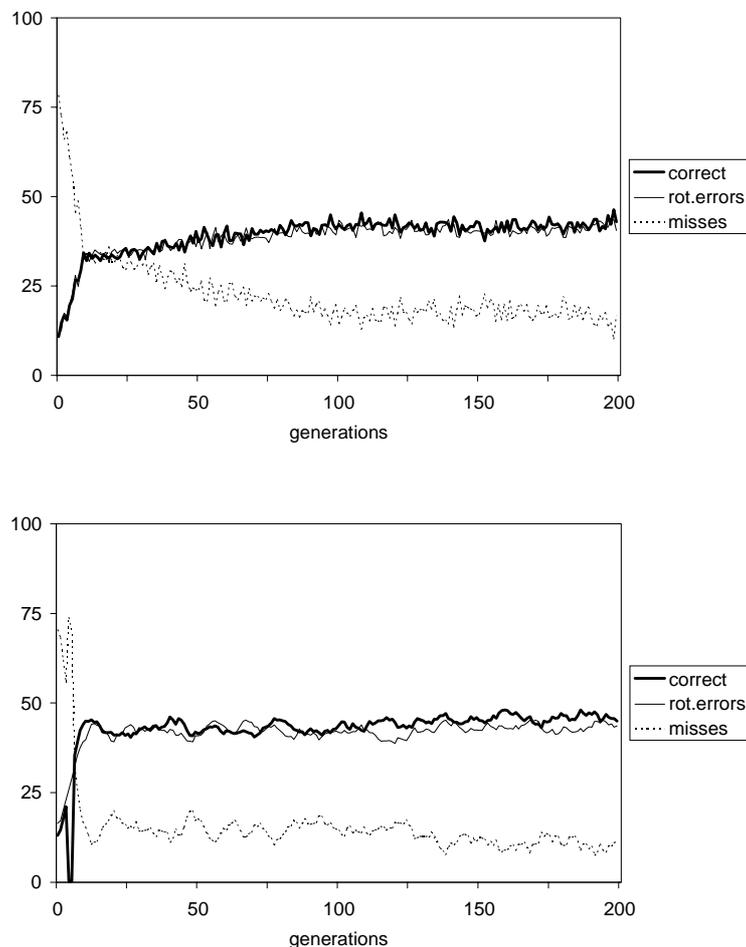


Figure 15. Performance (i.e., percentage of hits, rotational errors, and misses) throughout generations obtained by testing the best individual of each generation for 200 different starting positions. Results of an experiment with 500 cycles of lifetime for each epoch in which the starting positions and orientations, the size of the environment, and the proportion between long and short walls varied randomly for each epoch. **Top:** Average results of 10 replications. **Bottom:** Result of the best replication (rolling average over 7 data points).

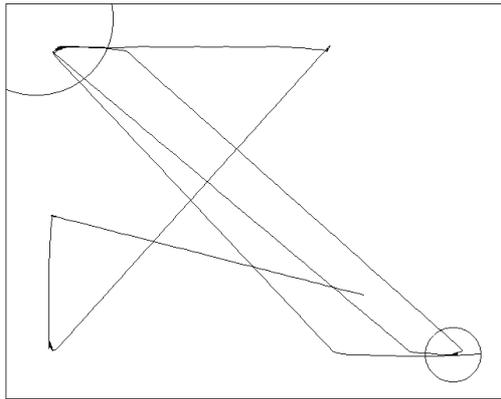


Figure 16. A typical strategy displayed by evolved individuals.

Figure 16 shows a typical strategy displayed by evolved individuals. The robot explores the environment by avoiding walls and moving forward in straight lines until it finds one of the corners (in the case of the Figure it finds the bottom-left corner first). When it is facing a corner it move left and right and back and fourth for a while and then leaves the corner with an angle of about  $45^\circ$  from the two walls that forms the corner. Finally, when it encounter a wall with an angle of about  $45^\circ$  on its left or right side it turns anti-clockwise or clockwise respectively and then follow the wall until the corner.

This strategy ensures that after the first corner (that might be any corner given that the initial position and orientation of the robot that is randomly chosen) the robot will always move toward the target or the rotational equivalent corner. Notice that this strategy works for any rectangular environment independently of the relative dimensions of the long versus short walls and of the absolute dimension of the environment.

Indeed leaving corners with an angle of  $45^\circ$  is smart way of measuring the relative length of the walls. In fact, if the robot find a wall at about  $45^\circ$  on its left side and it previously left a corner, it means that the actual wall is one of the two longer walls. Conversely, if it encounters a wall at  $45^\circ$  on its right side, the actual wall is necessarily one of the two shorter walls. What it is interesting is that the robot "measures" the relative length of the walls through action (i.e. by exploiting sensory-motor coordination) and to do so it does not need any internal states. Also notice that this strategy allows the robot to "measure" the *relative length* of the walls and therefore avoid the need to measure two adjacent walls and to compare their length in order to ascertain which is the longer or shorter wall.

Also notice how, *given the selected behavior, sensory states might indirectly encode information about features of the environment that do not directly affect the sensory state themselves*. In this case, a sensory states experienced by the robot while it encounter a wall at about  $45^\circ$  on the left or on the right side indirectly indicate the relative length of current wall. This seems to imply jet another possible way to exploit sensory-motor coordination. Individuals might select a behavior (e.g. leaving corner at a certain angle) that assures that the presence of a certain sensory state later on (e.g. for example the sensory state corresponding to a wall at about  $45^\circ$  on the left side) might be used to infer information about the environmental structure (e.g. the fact that the wall is one of the two longer walls).

Another thing to notice is that, as in the case of the experiment described in the previous section, these robots cannot stop. In fact, given that the there are not sensory states that unambiguously identify the target area and given that agents do not rely on internal states, they can only produce the desired behavior (move toward the target area) as an emergent result of a set of sensory-motor actions. Interestingly, in this case, emergent behaviors arise at

two different time scales. At a short time scale, individuals oscillate in front of a corner in order to remain there for a while and then leave (this is accomplished by creating a sort of uncompleted behavioral attractor). This oscillatory behavior is necessary given that robot cannot know if the actual corner is the target or its rotational equivalent corner or one of the two wrong corners. As we said above, in fact, pure sensory-motor robots that do not have any internal states cannot "remember" if they already left a corner and therefore cannot be sure that they are actually facing the target or the rotational equivalent corners. For the same reason, at a longer time scale, the robot has to oscillate between the target and the rotational equivalent corner as shown in Figure 16. In fact, by quickly moving toward the target or the rotational equivalent corner and by spending some time in each corner individual maximise their chance to end their lifetime in one of the two corners.

These consideration takes us to a limit of purely reactive agents. Often pure sensory-motor strategies are effective in most of the cases but are not optimal. This strategy, for example, fails when the robot encounter a short wall at about  $45^\circ$  on the left side or a long wall at about  $45^\circ$  on the right side before leaving a corner (evolved individuals act in a way to minimise the risk to experience this unfavourable condition but they cannot solve the problem when they are initially placed at an angle of about  $45^\circ$  with respect to the nearest wall). Apparently this problem might only be solved by a robot provided with at least an internal state. This hypothetical robot should turn its internal state on the first time it leaves a corner and should use the internal state to not leave corners anymore.

Interestingly, at least in this case, the addition of internal states does not imply a profound re-organization of the strategy itself. The new strategy will still be mostly reactive (i.e. the internal state will play a comparatively limited role). This is interestingly from two perspectives. From an evolutionary perspective it might help us to understand how evolution may synthesise individuals that rely on internal states starting from pure sensory-motor individuals (i.e. it might help us to understand how complex forms of life emerged). From a theoretical point of view it might help to clarify in which sense agents might have incomplete and implicit representations of the environment. Representations might be incomplete not only because do not include all the characteristics of the external environment (e.g. the shape, the dimension, the actual position of the robot in the environment etc.) but also because they might encode information about the external environment in very indirect way that are strongly dependent on the robot behavior. As in the case of this experiment, it might be convenient to encode *what the robot did instead of how the environment looks like*.

## 6. Conclusions

We have presented a set of experiments showing that rather complex tasks can be solved by exploiting sensory-motor coordination. In particular, we have shown that agents can use sensory-motor coordination to: (a) select sensory patterns which are not affected by the aliasing problem and avoid those that are (section 2); (b) select sensory patterns for which groups of patterns requiring different answers do not overlap too much (section 3); (c) exploit the fact that, by selecting a given behavior, the sensory states that are experienced later on might indirectly encode information about features of the environment (section 5); (d) exploit emergent behaviors resulting from a sequence of sensory-motor loops and/or from the interaction between the robot and the environment (section 4 and 5). Emergent forms of behavior might even display a hierarchical organization in which at a lower level emergent behavior arises from the interaction of a sequence of sensory-motor loops (such as the behavior that allow the robot to remain close to the corner for a while that arises from a sequence of left and right movement in front of a corner) and at a higher level arises from the

interaction between emergent behavior themselves (such as the oscillatory behavior between two corners arising from the interaction between the semi-attractors behavior displayed in front of two opposite corners).

We have also seen how sensory-motor coordination allow agents to exploit environmental constraints (section 5). If many constraints are available, individuals develop strategies that rely on the constraints that are more easily available (such as the dimension and the exact shape of the environment). If such constraints are removed, individuals are forced to discover and to rely on more hidden constraints (such as the relative dimension between long and short walls).

Sensory-motor coordination allows pure reactive agents to solve problems that, from an external perspective, apparently require more complex solutions that rely on internal states or internal representations of the external environment. Despite of that, pure sensory-motor strategies might display sub-optimal performance in certain cases. Interestingly however, as we claimed in section 5, individuals relying on pure sensory-motor strategies can move to more complex strategies that also rely on internal states without necessarily undergoing to a profound reorganization of their control system and of their actual strategy. At least in the case of the experiment presented in section 5, a nice smooth path exist in which an internal state can be easily integrated with the pre-existing pure sensory-motor strategy. The key aspect of this smooth transition seems to be centered around the idea that internal states might be used to encode what the agent did in the past instead of how the environment looks like.

Between pure reactive agents and pure representational agents (i.e. agents that explicitly represent in their internal states how the external environment looks like and which is their actual relation with respect to the environment) a large variety of intermediate cases exists. As has been claimed by many authors, agents might explicitly represent in their internal states only the characteristics of the external environment that are relevant for achieving their goal with a given strategy (Brooks, 1991; Clark, 1987). In addition, as we showed above, agents might encode in their internal states what they did in the previous portion of their lifetime which, given a certain behavior, indirectly encode relative information about the environment. Finally, agents might rely on both internal states (which might partially and indirectly encode information about the external environment as explained above), and on their actual sensory pattern that, given a certain behavior, might also indirectly encode information about features of the environment that do not affect the current sensory pattern itself.

## Acknowledgments

The author thanks Dario Floreano and Andy Clark for useful comments on a first draft of this paper.

## References

- Ashby, W.R. (1956) *An introduction to cybernetics*. London, Chapman and Hall.
- Bajcsy, R. (1988) Active Perception. *Proceedings of the IEEE* (76) 8:996-1005.
- Brooks, R.A. (1991) Intelligence without reason. In: *Proceedings of 12th International Joint Conference on Artificial Intelligence*, eds. J. Mylopoulos & R. Reiter. San Mateo, CA: Morgan Kaufmann.
- Clark A. (1997) *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Clark, A. & Thornton, C. (1997) Trading spaces: Computation, representation, and the limits of uniformed learning. *Behavioral and Brain Sciences* 20:57-90.

- Cheng, K. (1986) A purely geometric module in the rat's spatial representation. *Cognition* 23:149-178.
- Dill, M., Wolf, R. & Heisenberg, M. (1993) Visual pattern recognition in drosophila involves retinotopic matching. *Nature* 355:751-753.
- Elman, J.L. (1993) Learning and development in neural networks: The importance of starting small. *Cognition* 48:71-99.
- Gallistel, C. R. (1990) *The Organization of Learning*. Cambridge, MA: MIT Press.
- Hendriks-Jansen, H. 1996. *Catching Ourselves in the Act*. Cambridge, MA:MIT Press.
- Lund, H.H. & Miglino, O. (1997) Evolving and breeding robots. In: *Proceedings of the First European Workshop on Evolutionary Robotics*, eds. P. Husbands & J-A Meyer. Berlin: Springer Verlag.
- Margules, J. & Gallistel, C. R. (1988) Heading in the rat: determination by environmental shape. *Animal Learning and Behavior*, 16:404-410.
- Miglino, O. (1996) Orientamento Spaziale in a open field box. Organismi biologici ed artificiali a confronto. *Ph.D. Thesis*, University of Palermo, Italy.
- Mondada, R., Franzi, E. & Ienne, P. (1993) Mobile robot miniaturization: A tool for investigation in control algorithms. In: *Proceedings of the Third International Symposium on Experimental Robots*, eds. T. Yoshikawa & F. Miyazaki. Kyoto, Japan.
- Nolfi, S. & Parisi, D. (1997) Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior* 1:99-105.
- Nolfi, S. (1996) Adaptation as a more powerful tool than decomposition and integration. In: *Proceedings of the workshop on Evolutionary computing and Machine Learning -13th International Conference on Machine Learning*, eds. T. Yoshikawa & F. Miyazaki: Bari, Italy.
- Nolfi, S. (1997) Evolving non-trivial behavior on autonomous robots: Adaptation is more powerful than decomposition and integration. In: *Evolutionary Robotics*, ed. T. Gomi. Kanata, Canada: AAI Books.
- Nolfi S. & Floreano D. (in press). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books
- Pfeifer, R. & Scheier, C. (1997) Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems* 20:157-178.
- Pfeifer R. & Scheier C. (1999) *Understanding Intelligence*. Cambridge, MA:MIT Press
- Rumelhart, D.E., Hinton G.E. & Williams, R.J. (1986) Learning internal representations by error propagation. In: *Parallel Distributed Processing*, eds. D.E. Rumelhart & J.L. McClelland. Cambridge, MA, MIT Press.
- Scheier, C., Pfeifer R. & Kuniyoshi Y. (1998) Embedded neural networks: exploiting constraints. *Neural Networks* 11:1551-1596.
- Thornton, C. (1997) Separability is a learner's best friend. In: *Proceedings of the Neural Computation and Psychology Workshop: Connectionist Representations*, eds. J.A. Bullinaria, D.W. Glasspool, & G. Houghton. London, Springer.
- Whitehead, S.D. & Ballard, D. H. (1991) Learning to perceive and act by trial and error. *Machine Learning* 7:45-83.