Categories Formation in Self-Organizing Embodied Agents

Stefano Nolfi

Institute of Cognitive Sciences and Technologies
National Research Council (CNR)
Viale Marx, 15, 00137, Rome, Italy
s.nolfi@istc.cnr.it

Keywords: Self-organization, Embodiment, sensory-motor coordination.

Abstract

In this paper we describe the implications from a general theory of category formation of a set of experiments in which embodied artificial agents are evolved for the ability to accomplish simple tasks. In particular we will focus on how categories might emerge from the dynamical interaction between an agent and its environment and on the relation between categories and behavior. Finally, we will introduce and discuss the notion of action-mediated categories, that is the notion of internal states that provide indirect and implicit information about the external environment and/or the agent/environment relation by exploiting the effects resulting from a stereotypic way of interacting with the environment.

1. Introduction

A new research paradigm, that has been called *Embodied Cognitive Science* (Pfeifer and Scheier, 1999), has recently 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, researchers tend to stress *situatedness*, i.e., the importance of studying systems that are situated in an environment (Brooks, 1991, Clark, 1997), *embodiment*, i.e., the importance of study systems that have bodies, receive input from their sensors and produce motor actions as output (Brooks, 1991; Clark, 1997), and *emergence*, i.e. the importance of viewing behavior and intelligence as the emergent result of fine-grained interactions between the control system of an agent including its constituents parts, the body structure, and the external environment. An importance consequence of this view is that the agent and the environment constitutes a single system, i.e. the two aspects are so intimately connected that a description of each of them in isolation does not make much sense (Maturana and Varela, 1980, 1988; Beer, 1995).

An attractive way of studying embodied and situated agents consists in developing these systems through a self-organizing process, such as artificial evolution, that allows them to develop their skills autonomously in close interaction with the environment and without human intervention (Nolfi and Floreano, 2000). Recent experimental research in this area showed how self-organizing artificial agents might develop simple cognitive abilities such us the ability to integrate sensory-motor information over time and form internal categories (Nolfi and Tani, 1999, Slocum, Downey and Beer, 2000; Nolfi and Marocco, 2001; Beer 2003).

In this chapter we will review these recent findings and their implications from the point of view of category formation. Rather then focusing on how a shared language can selforganize in a population of interacting embodied agents (on this issue see the chapter of Cangelosi in this book) we will focus on how categories might emerge from the dynamical interaction between an agent and its environment and on the relation between categories and behavior. In particular, we will introduce and discuss the notion of *action-mediated* states, that is the notion of sensory or internal states that provide indirect and implicit information about the external environment and/or the agent/environment relation by exploiting the effects resulting from stereotypic ways of interacting with the environment.

2. The Method

One effective way to build artificial agents able to develop their skill autonomously in close interaction with the environment is to rely on evolutionary computation and more specifically on evolutionary robotics techniques (Nolfi and Floreano, 2000).

The basic idea behind this approach is the following: an initial population of different artificial genotype, each encoding the control system (and sometimes the morphology) of a robot, are created randomly. Each robot (physical or simulated) is placed in the environment and it is left free to act (move, look around, manipulate) while its performance on various tasks is automatically evaluated. The fittest robots are allowed to reproduce by generating copies of their genotypes with the addition of changes introduced by some genetic operators (e.g., mutations, crossover, duplication). This process is repeated for a number of generations until an individual is born which satisfies the performance criterion (*fitness function*) set by the experimenter.

The experimenter must design the fitness function, which is a criterion that is used to measure how much an individual robot is able to accomplish the desired task. Moreover, the experimenter must specify how genetic information (which is usually encoded as a sequence of binary values) is translated into the corresponding phenotipical robot. However, the mapping between the genotype and phenotype is usually task independent and evolving individuals are selected only on the basis of the overall efficacy of their behaviour. This allows to minimize our a priori commitments on how a given problem should be solved thus exploring the space of possibilities in a relatively unbiased way.

3. Categories emerging from the interaction between the agent and the environment

In this section we will show how problems that apparently require agents able to discriminate different categories (i.e. different classes of environmental situations) can be solved by relying on simple control strategies that does not require to internally partition environmental situations into distinct classes eliciting different motor responses.

3.1 Finding and remaining into favorable environmental areas.

Consider the case of a simulated agent which lives in a circular strip divided into 40 cells (20 cells on the left and 20 on the right). At each time step the agent occupies one single cell and perceives a sensory state corresponding to the cell types. There are 20 different cell types and 20 different sensory states that the agent can perceive, numbered from 0 to 19. Cell types are distributed in a randomly generated fashion, but each cell type is present once in both the left and in the right part of the environment (see Figure 1, left). The agent can react to the current sensory state in two different ways (move one cell clockwise or counterclockwise). The goal of the agent is reaching and remaining in the left part of the environment (Nolfi, 2002).

Agents have a neural network with 20 sensory neurons which locally encodes the corresponding perceived sensory state and 1 output unit which binarily encodes one of the two possible actions (see Figure 1, right). As a consequence only one sensory neuron is activated each time step. Weights can assume only two values (-1 or 1). As a consequence, the weight of the connection between the current activated sensory neuron and the motor neuron locally encode how the agent reacts to the current sensory state (i.e. the agent moves clockwise and counterclockwise when the connection weight is -1 or 1, respectively). Agents do not have any memory of the previously experienced sensory states (i.e. they always react in the same way to a given sensory state).

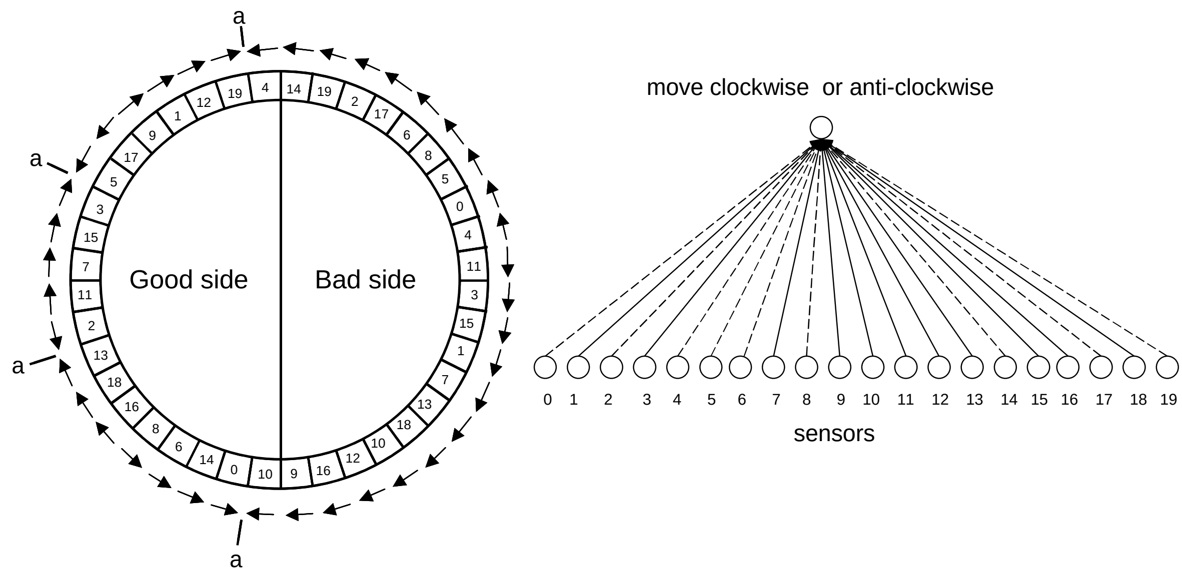

Figure 1. **Left:** The circular environment. Arrows indicate the movement produced by a typical evolved agent in the different cells of the environment. Points indicated with an “a” are attractors in the behaviour of the agent (i.e., points in which the agent starts to oscillate by moving clockwise and counterclockwise). **Right**: the neural controller of evolved agent. Dashed and full lines correspond to negative and positive connection weights that elicit counterclockwise and clockwise movement, respectively, when the corresponding sensory neuron is activated. The connection weights shown in the right figure correspond to the genotype of the same evolved individual shown in the left figure.

What is interesting about this experimental situation is that all possible sensory states do not provide by themselves any indication on the current location of the agent. For each possible sensory state, in fact, agents have a 50% probability of being in the left or in the right part of the environment. Apparently, therefore, agents that only decide how to act on the basis of the current sensory state cannot solve this problem.

However, by evolving agents for the ability to move toward the left part of the environment[1], we observed that, after few generations, evolving individuals are able to move away from the right part and to remain in the left part of the environment. The way in which evolved individuals solve this problem can be seen by observing the arrows in the right part of Figure 1. In the right part of the environment individuals consistently move clockwise or counterclockwise until they abandon the right side. Conversely, in some areas of the left side of the environment, individuals start to move back and forth by remaining there for the rest of the epoch.

---

1 Evolving individuals were allowed to "live" for 100 epochs, each epoch consisting of 200 actions. At the beginning of each epoch agents are placed in a randomly selected location of the environment. Fitness is computed by counting the number or epochs in which individuals were located in the left part of the environment after 200 cycles. Connection weights were binarily encoded 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.

Individuals react to the same sensory state always in the same manner. Despite of that, the way in which they react to the different sensory states allow them to produce behavioral attractors in the left but 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 1, left). When the agent encounters an attractor point, it remains there by moving back and forth. For an example in which the same type of strategy emerges by evolving a Khepera robot for the ability to discriminate between object with different shapes see Nolfi (2002).

3.2 Discriminating objects with different shapes on the basis of tactile information

As second example consider the case of a robot with an artificial finger that has to discriminate objects with different shapes on the basis of rather rough tactile information (Nolfi and Marocco, 2002).

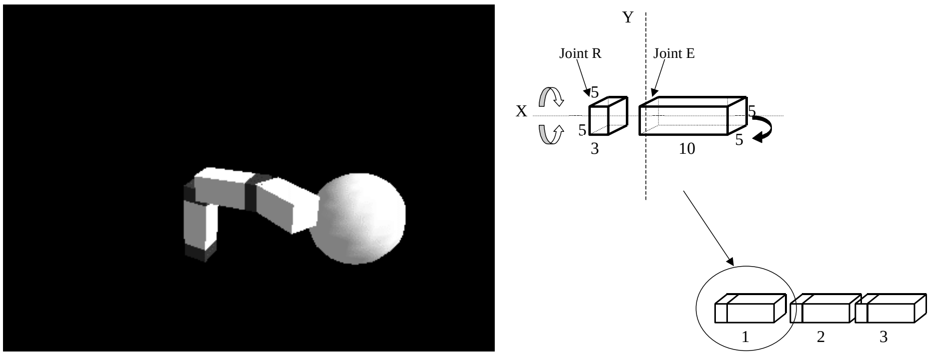

Figure 2. **Left:** The artificial finger and a spherical object. **Right:** A schematic representation of the finger.

The finger consists of 3-segments with 6 degrees of freedom (DOF) and extremely coarse touch sensors (see Figure 2, left). More precisely, the artificial finger consists in a basic structure of two bodies and two joints replicated for three times (see Figure 2, right). These two bodies are connected by means of a joint (i.e. the *Joint E* in Figure 2, right) that allows only one DOF on axis *Y* , while the shorter body is connected at the floor, or at the longer body, by means of a joint (i.e. the *Joint R*) that provides one DOF on axis *X*. In practice, the *Joint E* allows to elevate and to lower the connected segments and the *Joint R* allows to rotate them in both direction. *Joint E and Joint R* are free to moves only in a range between [0 and $\pi/2$] and [$-\pi/2$, $+\pi/2$], respectively. Each actuator is provided with a corresponding motor that can apply a varying force. Therefore, to reach every position in the environment the control system has to appropriately control several joints and to deal with the constraints due to gravity.

The sensory system consists of three simple contact sensors placed on each longer body that detect when these bodies collides with obstacles or other bodies and six proprioceptive sensors that provide the current position of each joint. The controller of each individual consists of a neural network with 10 sensory neurons directly connected to 7 motor neurons and 2 internal neurons receiving connections from the sensory neurons and from themselves and projecting connections to the motor neurons. The first 9 sensory neurons encode the angular position (normalized between 0.0 and 1.0) of the 6 DOF of the joints and the state of the three contact sensors located in the three corresponding segments of the finger. The last sensory neuron is a copy of the last motor neuron that encodes the current classification produced by the individual (see below). The first 6 motor neurons control the actuators of the

6 corresponding joints. The output of the neurons is normalized between [0, +π/2] and [-π/2, +π/2] in the case of elevation and rotational joints respectively and is used to encode the desired position of the corresponding joint. The motor is activated so to apply a force proportional to the difference between the current and the desired position of the joint. The seventh motor neuron encodes the classification of the object produced by the individual (value below or above 0.5 are interpreted as classifications corresponding to a cubic or spherical object respectively).[2]

By running 10 replications of the experiment and by evolving individuals for 50 generations we observed that in many of the replications evolved individuals display a good ability to discriminate the two objects and, in some cases, produce close to optimal performance.

By analyzing the obtained behaviors one can clearly see that in all experiments evolved individuals select a well defined behavior that assures that perceived sensory states corresponding to different objects can be easily discriminated and allows robust and effective categorizations. Figure 3 shows how a typical evolved individual behave with a spherical and a cubic object (left and right sides of the Figure, respectively). As can be seen, first the finger bends on the left side and move to the right so to start to feel the object with the touch sensor of the third segment. Then the finger moves so to follow the curvilinear surface of the sphere or so to keep touching one of the angles of the cubic object.

The fact that such behavior significantly simplifies the discrimination task can be explained by considering that the finger ends in very different conditions in the case of a sphere or of a cubic object. In particular, after a certain amount of time in which the finger is negotiating with the object, it ends almost fully extended in the case of spherical objects and almost fully bended in the case of cubic objects. This implies that, given such a behavior, the state of the proprioceptive sensors after a certain amount of time can be used as a direct and straightforward indication of the category of the object. The fact that such behavior allows evolved individuals to effectively discriminate the two objects can be explained by considering that the discrimination process is not the result of a single decision but is the end result of an interaction between the agent and the object that last several cycles. Indeed, evolved individuals that display shorter negotiation periods with spherical objects also produce worse performance (result not shown). A similar temporally-extended decision process has been observed by Beer (2003) in evolved agents asked to catch diamonds-shaped objects and avoid circular objects.

[2] Evolving individuals were allowed to "live" for 36 epochs, each epoch consisting of 150 actions. Each individual of the population was tested for 36 epochs, each epoch consisting of 150 lifecycles. At the beginning of each epoch the finger is fully extended and a spherical or a cubic object is placed in a random selected position in front of the finger (the position of the object is randomly selected between the following intervals: 20.0 >= X <= 30.0; 7.5 >= Y <= 17.5; -10.0 >= Z <= 10.0). The object is a sphere (15 units in diameter) during even epochs and a cube (15 units in side) during odd epochs so that each individual has to discriminate the same number of spherical and cubic objects during its "lifetime" Fitness is computed by counting the number or epochs in which individuals correctly classify the object (i.e. the number of times in which at the end of the epoch the activation of the last motor units is below 0.5 and the object is a cube or is above 0.5 and the object is a sphere). Population size was 100. The best 20 individuals of each generation were allowed to reproduce by generating 5 copies of their genotype with 1% of their bits replaced with a new randomly selected value.

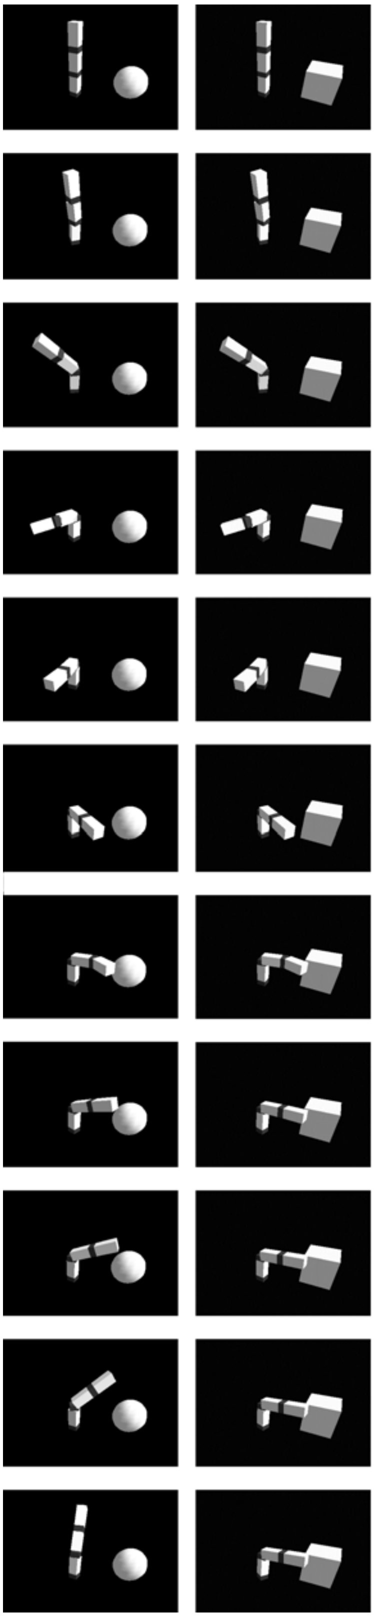

Figure 3. Behaviour of a typical evolved individual during an epoch (150 cycles) in which the object consists of a sphere (left pictures) and of a cube (right pictures). For reason of space, the pictures show the position of the finger each 15 cycles.

3.3 Behavior emerge from the dynamical interaction between the agent and the environment

The two examples reported in the previous sections show how the ability to categorize objects or environmental situations does not necessarily require agents able to partition sensory states or sequence of sensory states into different internal categories.

To understand this apparent paradox we should distinguish two ways of describing behavior. A *distal description of behavior* is a description from the observer's point of view in which high level terms such as "approach" or "discriminate" are used to describe the result of a sequence of sensory-motor loops. A *proximal description of behavior* is a description from the point of view of the agent's sensory-motor system that describes how the agent reacts to different sensory and internal states (see Figure 4).

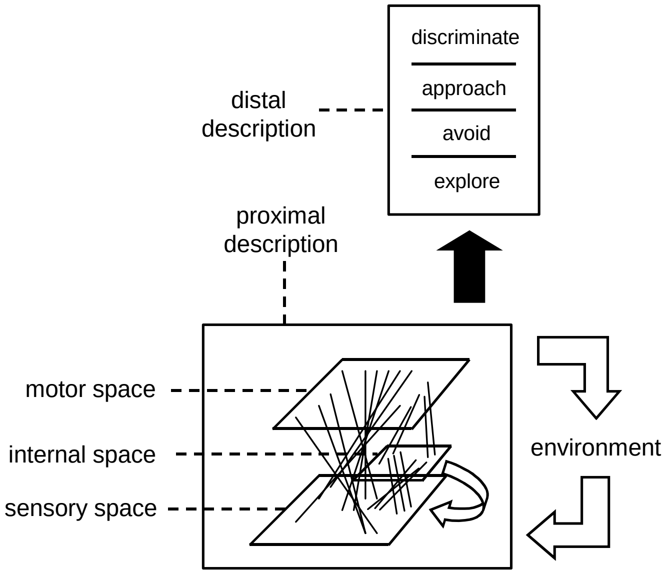

Figure 4. Proximal and distal description of behavior. **Bottom**: lines between planes indicate the mapping between sensory, internal and motor space. The curvilinear arrows indicate that internal states might influence each other and might influence sensory states thus producing an internal dynamics. The two arrows on the left indicate that the motor actions influence the environment or the relation between the agent and then environment that, in turn, influence the sensory states of the agent. **Top:** the full arrow indicates that behavior, from a distal description, emerges from the dynamical interaction between the agent control system (i.e. the behavior from a proximal description point of view) and the environment.

It should be noted that behavior from the point of view of a distal description is the result not only of behavior from a proximal description point of view but also of the environment. More precisely it is the result of the dynamical interaction between the agent and the environment. The sensory patterns that the environment provides to the agent partially determine the agent's motor reactions. These motor reactions in turn, by modifying the environment or the relative position of the agent in the environment, partially determine the type of sensory patterns that the agent will receive from the environment.

The fact that behavior, from a distal perspective, emerge from the dynamical interaction between the agent and the environment implies that there is not necessarily a one-to-one correspondence between the distal and the proximal description of behaviour and there is no reason to expect that what makes sense at the distant level of behaviour will also makes sense at the proximal level.

The fact that an agent discriminates different types of objects or discriminate different environmental situations (by producing different labels for them or by reacting differently in different environmental situations) from the point of view of a distal description of behaviour, therefore, does not necessarily imply that a discrimination process is occurring in the agent control system (i.e. at the level of the proximal description of behaviour).

Indeed, the evolved agent described in section 3.1 never "knows" whether it is located in the good or bad side of its environment and reacts in the same way in the two sides of the environment. It simply acts in a way that, by interacting with that environment, assures that it always leave the bad side and remains in the good side. Similarly the evolved agent described in section 3.2 does not "knows" whether the object it is currently touching is a cube or a sphere. It simply acts in a way that, by interacting with the object, produces two qualitatively different behaviours in the two cases (i.e. to keep touching the cube and leaving the sphere).

These two examples are a straightforward demonstration of the importance of embodiment, situatedness and emergence. Regarding embodiment we should consider that, for example, the behaviour of the artificial finger and its ability to discriminate different shapes strongly depend on the physical shape of the finger itself and on the results of the physical collisions between the finger and the objects. Regarding situatedness we should consider that the dynamical interaction with the environment and the structure of the environment play a crucial role in the way in which the problem is solved. In the case of the agent living in a circular environment, for instance, the ability to produce dynamical movement (e.g. moving back and fourth in the attractors area of the environment) clearly results from the dynamical interaction between the agent control system and the environment. Finally, regarding emergence, we should consider that evolved solutions are typically qualitatively different from the solutions that we, as external designer, tend to develop. Solving the problem of the circular environment for a human observer with a simple solution qualitatively similar to those discovered by artificial evolution, in fact, is extremely hard if not impossible.

4. Action-mediated sensory states

The case of the artificial finger described in section 3.2 is interesting also for another aspect that we will discuss in more detail in this section. The well defined way with which evolved individuals interact with their environment does not only allow them to display two different behaviors in the case of the two categories (i.e. cubes and spheres). It also assures that, after a period of time, the activation state of the sensors is well differentiated for different type of objects. After a short interaction with the objects, in fact, the finger is bended in the case of cubes and extended in the case of spheres. After this short negotiation phase, therefore, the activation state of the proprioceptive sensors that encode the positions of the 6 actuators become well differentiated, in the two cases. This indeed explains how the artificial finger is able to appropriately label the two objects.

More generally we can say that well defined ways to interact with the environment might allow sensory states to indirectly convey complex information about the external environment that would not become available without such interaction process. These states are action-mediated given that they do not convey such information by themselves, they only acquire their meaning after an appropriate interaction with the environment took place. For example, the state of the proprioceptive sensors corresponding to a bended finger do not provide any information about the shape of the object placed close to the finger, they only provide this information if the finger previously interacted with that object on the basis of the simple behavior described above.

In this section we will describe other two experiments in which evolved agents are able to solve their adaptive task by selecting simple ways of interacting with the environment that, in turn, assures that they will later experience useful action-mediated sensory states (i.e. sensory states providing ready to use information for discriminating different type of objects or different environmental locations).

4.1 Discriminating larger and smaller cylindrical objects

Consider the case of a mobile robot placed on an environment surrounded by walls that should be able to find and stay close to large cylindrical objects by avoiding small cylindrical objects. The robot is Khepera (Mondada, Franzi, Ienne, 1993) a miniature mobile robot with a diameter of 55 mm and a weight of 70 g. It is supported by two lateral wheels that can rotate in both directions and two rigid pivots in the front and in the back. The sensory system employs eight infrared sensors that are able to detect obstacles up to about four cm.

As demonstrated in (Scheier, Pfeifer, and Kunyioshi 1998; Nolfi, 2002) this problem is far from trivial given that the two categories corresponding to large and small objects largely overlap in sensory space. Indeed, distance in sensor space for sensory patterns originating from one object can be large, while the distance in sensory space for sensory patterns originating from two different objects can be small. Despite of that, evolved robots are able to solve the problem on the basis of a simple control strategy (Scheier, Pfeifer, and Kunyioshi 1998).

As in the case of the experiments reported in section 3, the authors evolved the connection weights of the robots neural controllers[3]. By analyzing the performance of robots through out generations, Scheier, Pfeifer & Kunyioshi, observed that they increase rather quickly during the first generations and stabilize around near optimal performance after about 40 generations. 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 individuals of the other 14% runs 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. In fact, while the sensory patterns corresponding to small and large cylinders largely overlap overall, the subset of sensory patterns experienced while the robot is circling small and large cylinders are nicely separated in sensory space. In other words, the circling behavior allows the robot to select sensory patterns that can be easily discriminated.

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

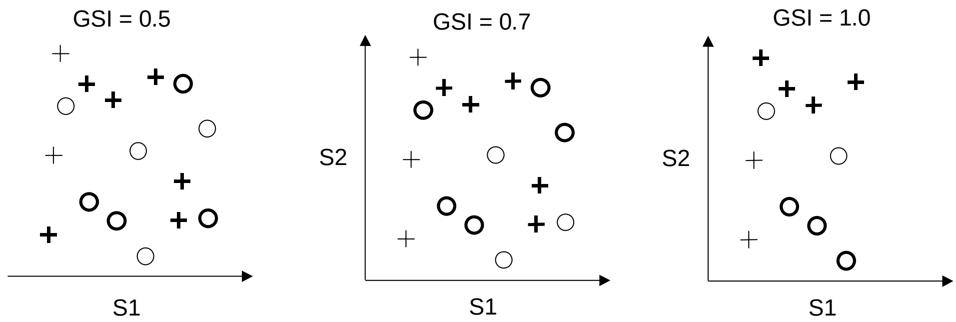

Figure 5. A schematic representation of the distribution of sensory patterns. For sake of simplicity the sensory space has only two dimensions (S1 and S2). Crosses and circles represent all possible sensory patterns originating from small and large cylinders, respectively. Dark crosses and circles represent the sensory patterns experienced by an individual. Notice that individuals tend to experience only a subset of all possible sensory patterns and that the patterns forming the subset depends from the type of behaviour displayed by the individual. The three figures indicate the sensory patterns experienced by three different individuals. As shown in the figure GSI might vary from 0.5 to 1.0 depending on how much the two groups of patterns overlap in sensory space.

The importance of the circling behaviour can be further demonstrated by analysing the complexity of the discrimination task for individual of successive generations. To understand this point we should consider that, 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 that differ in their way of interacting with the environment might face simpler or harder discrimination task.

To quantify the complexity of the discrimination task, the authors measured how much the two classes of sensory patterns corresponding to the two objects (small and large cylinders) were separated in the sensory space. This measure can be obtained by using the Geometric Separability Index (GSI) proposed by Thornton (1997) that provides a quantitative measure of the separation in space of two or more classes of sensory patterns.

In the case of these experiments, GSI can be computed by storing all sensory patterns experienced by an individual during N lifecycles and 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 categories 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 largely overlap (see Figure 5).

As reported in Scheier, Pfeifer & Kunyioshi (1998) the GSI value of the best individuals of successive generations 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 the ability of evolving individuals to find and stay close to large cylinders while avoiding small cylinders is mainly due to an ability to act so to experience sensory patterns that can easily be discriminated.

Sensory patterns experienced by evolved robots showing the circling behavior, therefore, are another example of action mediated sensory states. The sensory states experienced while the robot is circling around one large cylindrical object, for example, cannot be separated by all the other possible sensory states. However, they can be easily separated by the other sensory states that the robot experience in other situations if the robot interact with the environment according to a well specified way.

By evolving agents to visually discriminate between circular and diamond-shaped object by catching the former and avoiding the latter, Beer (2003) observed that evolved agents foveated and actively scanned any object before eventually catching or avoiding it. According to the author, the scanning behavior might have the same functional role of the circling behavior described above. Indeed, at page 214, Beer (2003) claims "…it is likely that this scanning accentuates the small differences between a circle and a diamond".

4.2 Navigating toward a target area of the environment

As a second example, consider the case of a Khepera robot placed in a randomly selected location of a rectangular environment that should navigate toward the north-west or the south-east corners of the environment (Figure 6). The size of the environment and the proportion between long and short walls randomly vary in each trial within a given range.

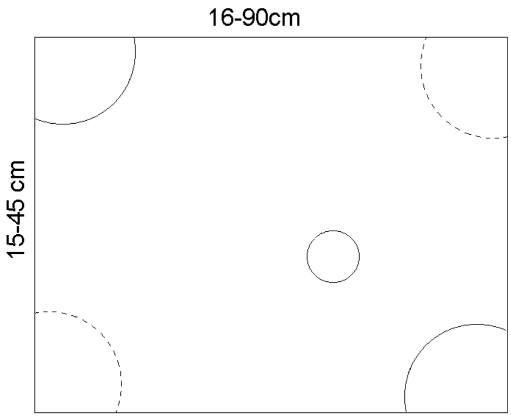

Figure 6. The environment and the robot. The lines represent the walls. Full and dashed arcs represent the right and wrong areas respectively. The small circle represents the Khepera robot.

Apparently, the only way to solve this problem is to discriminate between long and short walls and to navigate accordingly. For example the robot 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). Given that sensors can only provide information about the local portion of the environment surrounding the robot (i.e. they are activated by obstacles up to four cm) and given that the size of the environment might vary, the ability to detect long or short walls seems to require an ability to: (1) "measure" the length of two adjacent walls by moving along them and by identifying the beginning and the end of each wall, (2) "memorize" the measured length into internal states of the robot's controller, and (3) "compare" the two measured length stored into internal states.

By selecting robots for the ability to reach the two target areas[4], however, Nolfi and Marocco (2002) observed that evolving robots provided with simple reactive neural

[4] The architecture of the neural controller was fixed and consisted of a fully connected perceptron with 8 sensory and 2 motor neurons encoding the state of the 8 infrared sensors of the robot and the speed of the 2 motors controlling the two wheels. Individuals' fitness was increased or decreased of 1 point each time individuals ended their lifetime in one of the two right or wrong corners respectively. The genotype encoded the connection weights and biases of the neural controller. Each 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. Individuals were tested for 10 epochs. At the beginning of each epoch the length short and long walls was randomly selected within [15,45] and [16, 90] cm, respectively and the robot was placed in a randomly selected location within the environment with a randomly selected orientation.

controllers can solve this problem to a good extent (up to 85% of correct navigations in the case of the best replications of the experiment). These simple neural controllers do not have any internal states and therefore cannot accomplish the complex measuring behavior described above.

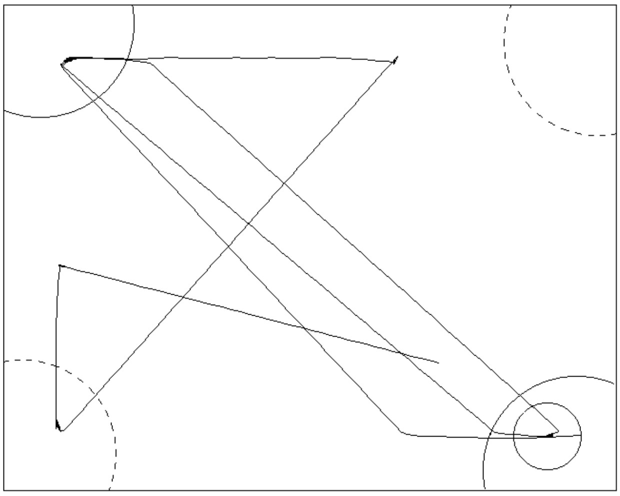

Figure 7. A typical strategy displayed by evolved individuals. The line within the rectangle represents the trajectory of a typical evolved individual during a trial.

Figure 7 shows a typical strategy displayed by evolved individuals. The robot explores the environment by avoiding walls and by moving forward in straight lines until it finds one of the corners (in the case shown in the figure the robot starts from the right side of the environment and encounters the bottom-left corner first). When it is facing a corner, the robot moves left and right and back and fourth for a while and then leaves the corner with an angle of about 45° with respect to the two walls forming the corner. Finally, when it encounters a wall it turns counterclockwise until the wall is located on its right side, and then follows 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 is randomly chosen), the robot will always reach one of the two long walls. At this point it can easily find the target area by turning until the wall is on its right side, and then following the wall. Notice that this strategy works for any rectangular environment independently of the relative length of long versus short walls and of the size of the environment. Indeed leaving corners with an angle of 45° is smart way of measuring the relative length of the walls. Once again, action mediation (i.e. leaving corners with an angle of 45° in the case of this experiment) allows sensory states experienced later on to assume useful meanings (i.e. sensory state corresponding to walls uniquely identify long walls and sensory states corresponding to corners uniquely identify the two target corners).

This simple strategy however does not allow evolving individuals to remain in the target corners. They spend sometime there, moving back and fourth, but they later abandon the current corner by quickly moving toward the other right corner. Indeed this is the reason why they do not reach optimal performance (fitness is computed by looking at how many trials end with the robot in one of the two target corners). This inability to remain in target corners can be explained by considering that evolved robots "know" how to move to reach the two target corners but do not "know" whether the corner in which they are currently located is a correct or not.

Further experiments conducted by evolving robots provided with a neural network with internal neurons and recurrent connections shown how, in this case, evolving individuals are also able to stop on one of the two target corners after abandoning one of two other corners (Nolfi and Marocco, 2002). Interestingly, the analysis of how evolved individuals solve the problem of finding and remaining in one of the two correct corners indicates that the same strategy described above is used to reach the correct target corners. Internal neurons only keep track of how much time is passed since the robot started to interact with the environment. If enough time has passed and the robot is on a corner, the robot stops there. This simple behaviour exploits the fact that leaving corners with an angle of about 45° and then following walls by keeping them on the right side guarantees that the robot will only encounter target corners after a while.

This example also shows how sub-optimal simple strategies based on action-mediated sensory states might be complemented with simple additional internal mechanisms. This possibility is important from an evolutionary or developmental perspective. In fact, it implies that simple strategies based on action-mediated sensory states might later be enhanced in an incremental fashion without necessarily undergoing profound re-organizations (Nolfi and Marocco, 2002).

5. Integrating sensory-motor information over time and the emergence of complex internal categories

The examples described in section 3 and 4 shows how non trivial problems can be solved without relying on internal categories but rather by exploiting action mediated sensory states that provide the information necessary to behave correctly at the right time and in a ready to use fashion. At this point we might be interested in trying to understand in what conditions embodied agents might be unable to solve their adaptive problems by only relying on simple reactive or quasi-reactive solutions.

From this perspective, interesting candidate situations are those in which agents cannot freely select their way to interact with the environment. Limitations on the interaction, in fact, by reducing the chances that useful properties emerging from the interactions can be exploited, might prevent the possibility to exploit action-mediated sensory states. In these cases, more complex strategies based on internal states, might be the only option available.

Several causes might prevent agents to freely determine their way to interact with the environment. One case is the case in which the environmental structure strongly limits the degrees of freedom of the agent behavior. A second class of cases might be constituted by situations in which the agent behavior should satisfy several constraints at the same time. Finally, another class of cases might be constituted by agents that should be able to communicate their sensory-motor experience to other agents when requested. One case in which the interaction between the agent and the environment is limited by both the environmental structure and the need to communicate will be reviewed in the next section.

5.1 The self-localization problem

Consider the case of a Khepera robot that should be able to travel along a loopy corridor (see Figure 8, left) and to self-localize by identifying its current location in the environment (Nolfi, 2002).

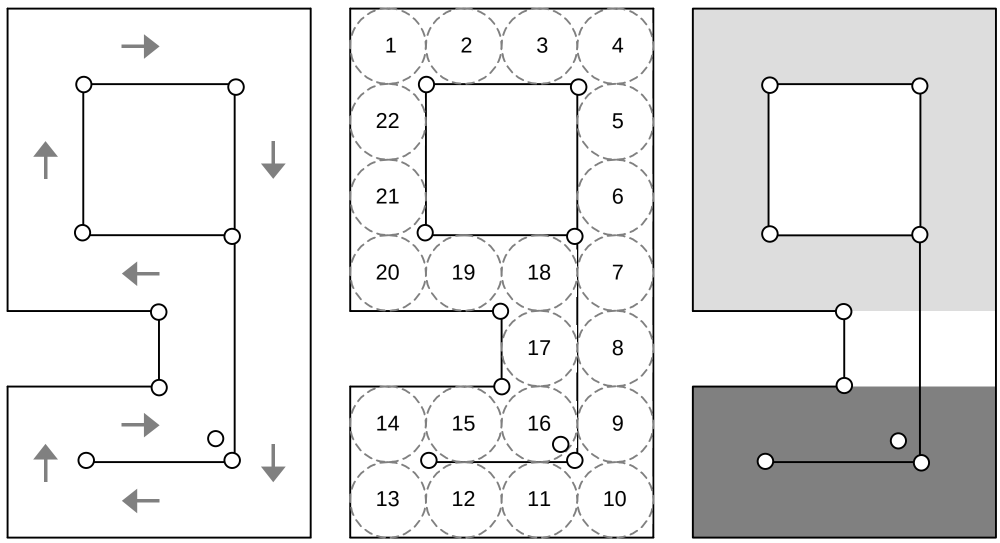

Figure 7. Left: The environment consists of a loopy corridor of 40x70 cm. Lines represent walls. Circles represent cylindrical objects. Arrows represent the starting positions and orientations in which the robot is placed at the beginning of each trial. **Center:** The environment is divided into 22 idealized regions placed along the corridor clockwise. **Right:** The environment is ideally divided into two rooms indicated with light and dark grey colours.

The controller of each individual consists of a neural network with nine sensory neurons directly connected to three motor neurons and five internal neurons receiving connections from the sensory neurons and sending connections to the motor neurons and to themselves (see Figure 8). The first three sensory neurons encode the state of the three corresponding motor neurons at the previous time step, the other six sensory neurons encode the six frontal infrared sensors (normalized between [0.0, 1.0]). The first two motor neurons encode the desired speed of the two corresponding wheels and the last motor neuron encodes the robot's self-localization output (see below). Internal neurons were updated according to an activation function with a genetically encoded time constant parameter (that allows neurons to change their activation state at different time rates) and a thresholded activation function (see Nolfi [2002b] for details).

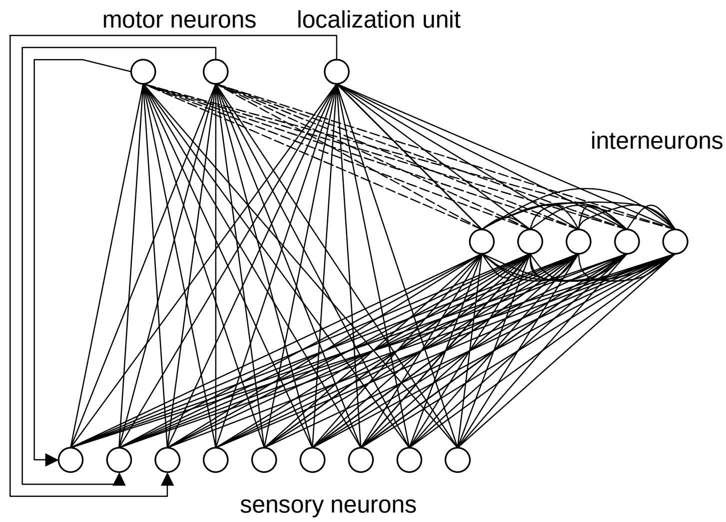

Figure 8. The architecture of the neural controller.

The fitness function has two components that reward, respectively, the ability to travel clockwise along the corridor and the ability to indicate the current position of the robot in the environment (see Nolfi, 2002 for details).

Evolved individuals show close to optimal performance on both the navigation and localization task. Figure 9 displays the behavior and the neural activity of one evolved individual. As shown in the Figure, the internal neuron i1 is turned off when the robot negotiates right handed corners (see the locations indicated with the letter **A** on the left side of the Figure) and increases its activation while the robot travels along straight corridors. Thanks to a recurrent positive connection, however, the neuron is turned off on right-handed corners only if its activation level is below a given threshold (on left-handed corners, instead, this neuron is turned off independently from its activation state, see the point indicated with **C**). The final result is that this neuron is always below a given threshold in the light gray room due to the reset of its activity occurring in C and in A corners and is always over that threshold in the dark gray room. Notice that internal neurons i1 is used to capture sensory-motor regularities that extend over rather long time scales (ranging from few to several seconds). Indeed, in order to display self-localization this robot is able to detect regularities such us corners (that extend over a period of few hundreds of milliseconds) and regularities such as corridors of different length (that extend from few to several seconds).

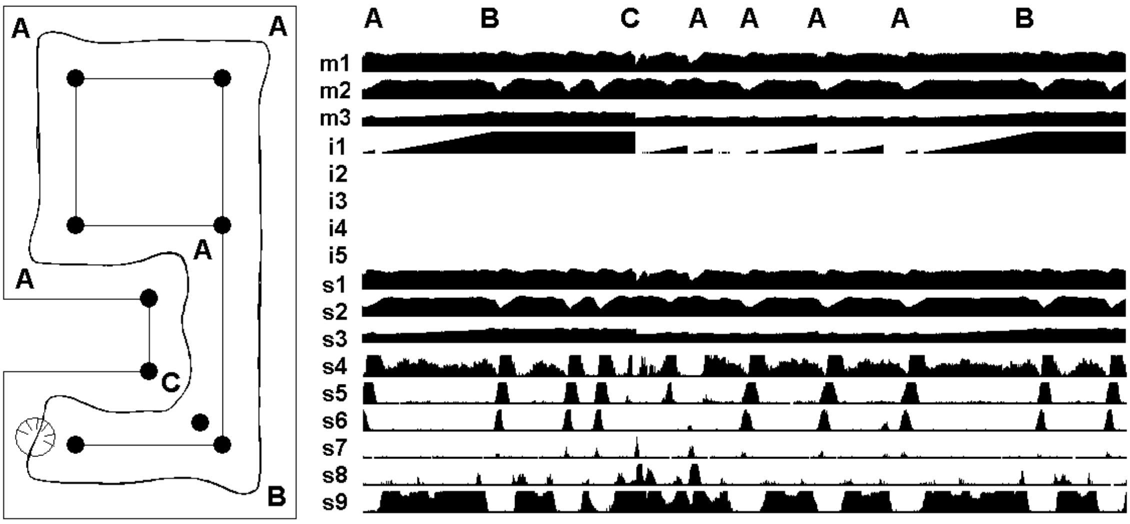

Figure 9. The architecture of the neural controller. Neural outputs of a successful evolved robot travelling in the environment. **Left**: The environment and the robot trajectory during a few laps of the corridor. A and C indicate the approximate position of the robot when the output of the first internal unit is reset to 0. B indicates the position of the robot when the first internal unit reaches its maximum activation level. **Right**: The output value of motor (m1-m3), internal (i1-i5), and sensory (s1-s9) neurons while the robot is travelling along the corridor (the output value is represented by the height with respect to the baseline).

Notice how, contrary to the experiments described in section 3 and 4, in the experiment described in this section evolved robots develop internal categories, that is internal states that integrate sensory-motor information through time and co-varies with the current position of the robot in the environment. In the case of the evolved individual described in Figure 9, for example, the internal neuron *i1* encodes the distance travelled by the robot since the last left-handed corner or the last right-handed corner followed by short corridors. In other evolved individuals internal neurons simply encode the distance travelled from the last left-handed corner (see Nolfi, 2002), or the frequency with which the robot encountered left and right-handed corners during its previous movements weighted by the type of corner (see Croon, Nolfi and Postma, in press).

The possibility to rely on simpler strategies exploiting action-mediated sensory states is prevented by the need to move fast and by the structure of the environment that, by being constituted by tight corridors, leave very little degrees of freedom in the way in which the robot can move in the environment. Indeed, in experiments in which agents were unable to extract relatively “complex” internal categories like those described above we observed that evolving individuals: (1) displayed very poor performance on the self-localization task when asked to travel fast (i.e. when asked to visit at least 1000 succeeding areas of the environment), and (2) displayed good but sub-optimal performance when allowed to travel at a lower speed (Croon, Nolfi and Postma, in press). In the latter case, evolved individuals exploited action-mediated sensory states that allowed them to partially solve the self-localization process without relying on internal categories. For example, some evolved individuals travelled along corridors by slightly moving from the left to the right side of the corridor. This allowed them to experience a unique sensory state toward the end of the long corridor that, in turn, allowed them to detect the beginning the dark grey room without encoding internally the length of corridors.

6. Conclusions

In this paper we described the results of a set of experiments in which embodied artificial agents autonomously develop their abilities, in interaction with the environment, thanks to a self-organizing process based on artificial evolution.

By analyzing the evolved individuals we observed that, by exploiting properties emerging from well specified ways to interact with the environment, they can solve non trivial problems without the need to develop internal categories and more generally without the need to internally process sensory-motor information. By selecting well defined ways to interact with the environment, in fact, evolved individuals are able to experience *action-mediated* sensory states that provide ready to use information when needed (i.e. information that can be transformed directly into the appropriate motor actions without significant further elaboration).

We also showed how simple reactive strategies based on the exploitation of action mediated sensory states might be complemented with an ability to integrate sensory-motor information over time into internal states that can later be used to appropriately modulate the agent's behavior.

Finally we showed how the need to rely on internal categorization and more generally on internal elaboration of sensory-motor information tends to be particularly compelling in the case of agents that, due to environmental and adaptive constraints, cannot freely chose between different ways of interacting with the environment.

The results and the analysis reported in this paper demonstrate how the evolutionary method is a powerful tool for understanding adaptive behavior in embodied and situated agents. It provides a way to understand how behavior emerges from the interaction between the control system, the body, and the environment.

Acknowledgements

This research has been partially supported by the ECAGENTS project founded by the Future and Emerging Technologies program (IST-FET) of the European Community under EU R&D contract 001940. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that may be made of data appearing in this publication.

References

Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence, 72*, 173–215.

Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior, 11*, 209–243.

Brooks, R. A. (1991). Intelligence without reason. In J. Mylopoulos & R. Reiter (Eds.), *Proceedings of 12th international joint conference on artificial intelligence* (pp. 569–595). San Mateo, CA: Morgan Kaufmann.

Cangelosi (this volume). Grounding symbols in perceptual and sensorimotor categories: Connectionist and embodied approaches. In H. Cohen & C. Lefebvre (Eds), *Handbook of Categorization in Cognitive Science*. Elsevier.

Clark, A. (1997). *Being there: Putting brain, body and world together again.* Cambridge, MA: MIT Press.

De Croon, G., Nolfi, S., & Postma, E. O. (in press). Toward pro-active embodied agents: On the importance of neural mechanisms suitable to process information in time. In D. Braha, A. Minai, & Y. Bar-Yam (Eds.), *Complex engineering systems*. City: Perseus Books Groups Press.

Maturana, H. R., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. Dordrecht: Reidel.

Maturana, H. R., & Varela, F. J. (1988). *The tree of knowledge: The biological roots of human understanding*. Boston: New Science Library.

Mondada, R., Franzi, E., & Ienne, P. (1993). Mobile robot miniaturization: A tool for investigation in control algorithms. In T. Y. Yoshikawa & F. Miyazaki (Eds.), *Proceedings of the third international symposium on experimental robots* (pp. 501–513). Berlin: Springer-Verlag.

Nolfi S. (2002a). Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time scales. *Connection Science* (14) 3:231-244.

Nolfi, S. (2002b). Power and limits of reactive agents. *Neurocomputing*, *49*, 119–145.

Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.

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

Nolfi, S., & Marocco, D. (2002). Active perception: A sensorimotor account of object categorization. In B. Hallam, D. Floreano, J. Hallam, G. Hayes, & J.-A. Meyer (Eds.), *From animals to animats 7: Proceedings of the VII international conference on simulation of adaptive behavior* (pp. 266–271). Cambridge, MA: MIT Press.

Nolfi, S., & Tani, J. (1999). Extracting regularities in space and time through a cascade of prediction networks: The case of a mobile robot navigating in a structured environment. *Connection Science*, *11*, 129–152.

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

Scheier, C., Pfeifer, R., & Kuniyoshi, Y. (1998). Embedded neural networks: Exploiting constraints. *Neural Networks*, *11*, 1551–1596.

Slocum, A. C., Downey, D. C., & Beer, R. D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat, & S. Wilson (Eds.), *From animals to animats 6: Proceedings of the sixth international conference on simulation of adaptive behavior* (pp. 430–439). Cambridge, MA: MIT Press.

Thornton, C. (1997). Separability is a learner's best friend. In J.A. Bullinaria, D.W. Glasspool, & G. Houghton (Eds.) Proceedings of the Neural Computation and Psychology Workshop: Connectionist Representations. London: Springer Verlag.