Behavior and Cognition as a Complex Adaptive System: Insights from Robotic Experiments

Stefano Nolfi

Institute of Cognitive Sciences and Technologies
National Research Council (ISTC-CNR)
Via S. Martino della Battaglia, 44, Rome, Italy
stefano.nolfi@istc.cnr.it

1. Introduction

Recent advances in different disciplines, ranging from cognitive sciences and robotics, biology and neurosciences, to social sciences and philosophy are clarifying that intelligence resides in the circular relationship between the brain of an individual organism, its body, and the environment (including the social environment).

In this chapter we will focus our attention on the evidence collected in robotic research with particular reference to results obtained in experiments in which the robots develop their skill autonomously while they interact with the external environment through an adaptive process. In particular we will demonstrate how the behavioural and cognitive skills developed by the robots can be properly characterized as complex adaptive systems which: (a) arise from the interactions between the brain of the robots, their body, and the environment and eventually between the dynamical process occurring within the robot and within the environment, (b) display a multi-level and a multi-scale organization in which the interaction between behavioural and cognitive properties at a certain level of organization lead to higher-level properties and in which higher-level properties later affect the lower-level properties.

The complex system nature of behaviour and cognition has important consequences both from an engineering and a modelling point of view. From the point of view of developing effective robotic artefacts it implies the need to rely on “design for emergence” techniques, i.e. techniques allowing the development of robots which are able to exploit useful emergent properties.

From the point of view of modelling biological systems, it implies the need to conceptualize behaviour and cognition as dynamical processes which unfold in time while the organism interacts with the environment.

Viewing behaviour and cognition as a complex adaptive system represents a new conceptual framework that can have profound consequences for cognitive science. In particular, as we will discuss in the concluding section, it might allow us to clarify better the notion and the role of embodiment and situatedness.

In section 2, we will demonstrate how behavioural and cognitive skills developed by robots which adapt to their task-environment through an adaptive process can be properly characterised as a complex system. In section 3, we will discuss the relation between adaptation and the complex system nature of behaviour. Finally in section 4, we will discuss the implications of the complex adaptive systems of behavioural and cognitive skills.

2. Behavior and cognition as complex systems

In agents which are embodied (i.e. have a physical body) and are situated (i.e. are located in a physical environment with which they interact) behavioural and cognitive skills are dynamical properties which unfold in time and which arise from the interaction between agents' nervous system, body, and the environment (Beer, 1995; Chiel and Beer, 1997; Keijzer, 2001; Nolfi and Floreano, 2000; Nolfi 2005) and from the interaction between dynamical processes occurring within the agents' control system, the agents' body, and within the environment (Beer, 2003; Tani & Fukumura 1997; Gigliotta and Nolfi, 2008). Moreover, behavioural and cognitive skills typically display a multi-level and multi-scale organization involving bottom-up and top-down influences between entities at different levels of organization. These properties imply that behavioural and cognitive skills in embodied and situated agents can be properly characterized as complex systems (Nolfi, 2005).

These aspects and the complex system nature of behaviour and cognition will be illustrated in more detail in the next section with the help of examples involving robotic experiments.

2.1 Behavior and cognition as emergent dynamical properties

Behaviour and cognition are dynamical properties which unfold in time and which emerge from high-frequency non-linear interactions between the agent, its body, and the external environment (Chiel and Beer, 1997).

At any time step, the environmental and the agent/environmental relation co-determine the body and the motor reaction of the agent which, in turn, co-determines how the environment and/or the agent/environmental relation vary. Sequences of these interactions, occurring at a fast time rate, lead to a dynamical process – behaviour – which extends over significantly larger time spans than the interactions (Figure 1).

Since interactions between the agent's control system, the agent's body, and the external environment are nonlinear (i.e. small variations in sensory states might lead to significantly different motor actions) and dynamical (i.e. small variations in the action performed at time t might significantly impact later interactions at time t_{+x}) the relation between the rules that govern the interactions and the behavioural and cognitive skills originating from the interactions tend to be very indirect. Behavioural and cognitive skills thus emerge from the interactions between the three foundational elements and cannot be traced back to any of the three elements taken in isolation. Indeed, the behaviour displayed by an embodied and situated agent can hardly be predicted or inferred by an external observer even on the basis of a complete knowledge of the interacting elements and of the rules governing the interactions.

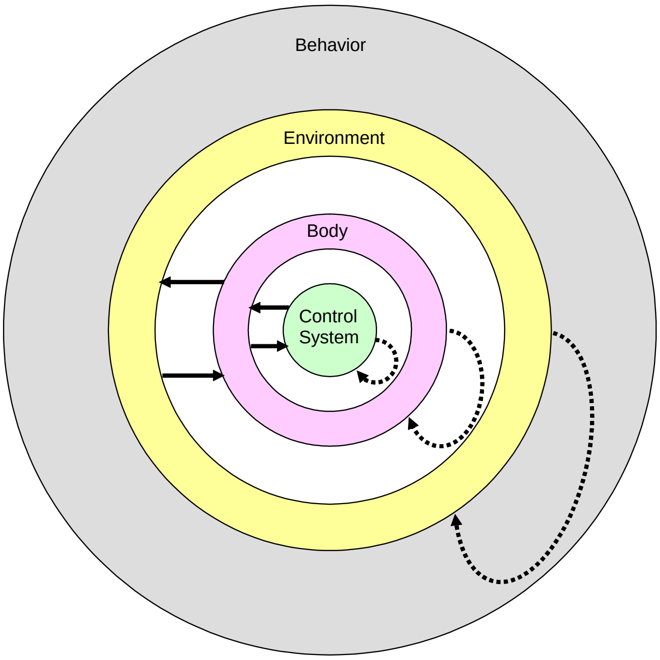

Figure 1. A schematic representation of the relation between agent's control system, agent's body, and the environment. The behavioural and cognitive skills displayed by the agent are the emergent result of the bi-directional interactions (represented with full arrows) between the three constituting elements – agents' control system, agent's body, and environment. The dotted arrows indicate that the three constituting elements might be dynamical systems on their own. In this case, agents' behavioural and cognitive skills result from the dynamics originating from the agent/body/environmental interactions but also from the combination and the interaction between dynamical processes occurring within the agents' body, within the agents' control system, and within the environment (see section 3)

A clear example of how behavioural skill might emerge from the interaction between the agents' body and the environment is constituted by the passive walking machines developed in simulation by McGeer (1990) --- two-dimensional bipedal machines able to walk down a four-degree slope with no motors and no control system (Figure 2). The walking behaviour arises from the fact that the physical forces resulting from gravity and from the collision between the machine and the slope produce a movement of the robot and the fact that the robot's movements produce a variation of the agent-environmental relation which in turn produces a modification of the physical forces to which the machine will be subjected in the next time step. The sequence of bi-directional effects between the robot's body and the environment can lead to a stable dynamical process --- the walking behaviour.

The type of behaviour which arises from the robot/environmental interaction depends on the characteristics of the environment, the physics laws which regulate the interaction between the body and the environment, and the characteristics of the body. The first two factors can be considered as fixed but the third factor, the body structure, can be adapted to achieve a given function. Indeed, in the case of this biped robot, the author carefully selected the leg length, the leg mass, and the foot size to obtain the desired walking behaviour. In more general terms, this example shows how the role of regulating the interaction between the robot and the environment in the appropriate way can be played not only by the control system but also by the body itself, providing that the characteristics of the body have been shaped to favour the exhibition of the desired behaviour. This property, i.e. the ability of the body to control its interaction with the environment, has been named 'morphological computation' (Pfeifer, Ida & Gomez, 2006). For related work which demonstrates how effective walking machines can be obtained by integrating passive walking techniques with simple control mechanisms, see Bongard and Paul (2001), Endo (2002), Vaughan et al. (2004). For related work which shows the role of elastic material and elastic actuators for morphological computing see Schmitz et al. (2007), Massera et al. (2007).

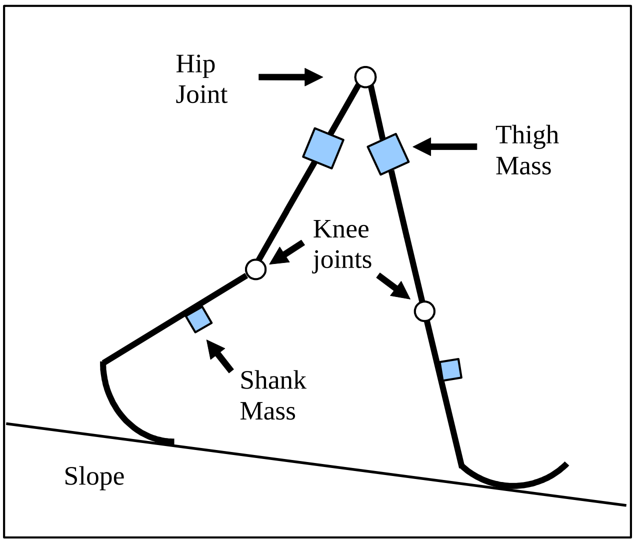

Figure 2. A schematization of the passive walking machine developed by McGear (1990). The machine includes two passive knee joints and a passive hip joint.

To illustrate how behavioural and cognitive skills might emerge from agent's body, agent's control system, and environmental interactions we describe a simple experiment in which a small wheeled robot situated in an arena surrounded by walls has been evolved to find and to remain close to a cylindrical object (Nolfi, 2002). The Khepera robot (Mondada, Franzi & Ienne, 1993) is provided with eight infrared sensors and two motors controlling the two corresponding wheels (Figure 3).

From the point of view of an external observer, solving this problem requires robots able to: (a) explore the environment until an obstacle is detected, (b) discriminate whether the obstacle detected is a wall or a cylindrical object, and (c) approach or avoid objects depending on the object type. Some of these behaviours (e.g. the wall-avoidance behaviour) can be obtained through simple control mechanisms but others require non trivial control mechanisms. Indeed, a detailed analysis of the sensory patterns experienced by the robot indicated that the task of discriminating the two objects is far from trivial since the two classes of sensory patterns experienced by robots close to a wall and close to cylindrical objects largely overlap.

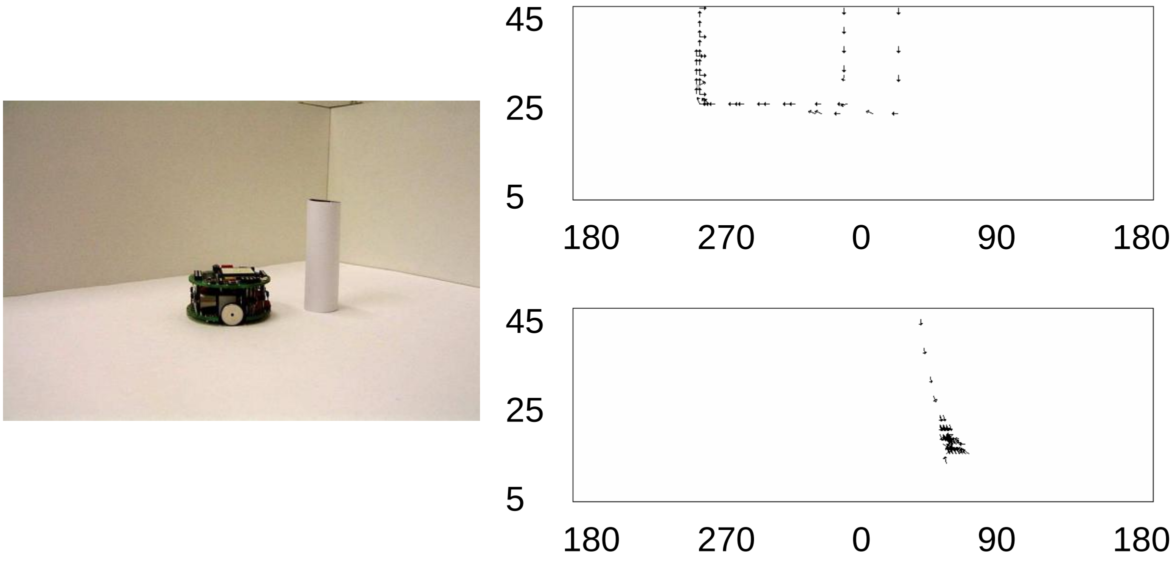

Figure 3: **Left**: The agent situated in the environment. The agent is a Khepera robot (Mondada, Franzi & Ienne, 1993). The environment consists of an arena of 60 × 35 cm containing a cylindrical object placed in a randomly selected location. **Right:** Angular trajectories of an evolved robot close to a wall (top graph) and to a cylinder (bottom graph). The picture was obtained by placing the robot at a random position in the environment, leaving it free to move for 500 time steps each lasting 100ms, and recording its relative movements with respect to the two types of objects for distances smaller than 45 mm. The x-axis and the y-axis indicate the relative angle (in degrees) and distance (in mm) between the robot and the corresponding object. For sake of clarity, arrows are used to indicate the relative direction, but not the amplitude of movements.

The attempt to solve this problem through an evolutionary adaptive method (see section 4) in which the free parameters (i.e. the parameters which regulate the fine-grained interaction between the robot and the environment) are varied randomly and in which variations are retained or discarded on the basis on an evaluation of the overall ability of the robot (i.e. on the basis of the time spent by the robot close to the cylindrical object) demonstrated how adaptive robots can find solutions which are robust and parsimonious in terms of control mechanisms (Nolfi, 2002). Indeed, in all replications of these experiments, evolved robots solve the problem by moving forward, by avoiding walls, and by oscillating back and fourth and left and right close to cylindrical objects (Figure 3, right). All these behaviours result from sequences of interactions between the robot and the environment mediated by four types of simple control rules which consist in: turning left when the right infrared sensors are activated, turning right when the left infrared sensors are activated, moving back when the frontal infrared sensors are activated, and moving forward when the frontal infrared sensors are not activated.

To understand how these simple control rules can produce the required behaviours and the required arbitration between behaviours we should consider that the same motor responses produce different effects on different agent/environmental situations. For example, the execution of a left-turning action close to a cylindrical object and the subsequent modification of the robot/object relative position produce a new sensory state that triggers a right-turning action. Then, the execution of the latter action and the subsequent modification of the robot/object relative position produce a new sensory state that triggers a left-turning action. The combination and the alternation of these left and right-turning actions over time produce an attractor in the agent/environmental dynamics (Figure 3, right, bottom graph) that allows the robot to remain close to the cylindrical object. On the other hand the execution of a left-turning behaviour close to a wall object and the subsequent modification of the robot/wall position produce a new sensory state that triggers the reiteration of the same motor action. The execution of a sequence of left-turning actions then leads to the avoidance of the object and to a modification of the robot/environmental relation that finally leads to a sensory state that triggers a move-forward behaviour (Figure 4, right, top graph).

Before concluding the description of this experiment, it is important to notice that, although the rough classification of the robot motor responses into four different types of actions is useful to describe the strategy with which these robots solve the problem qualitatively, the quantitative aspects which characterize the robot motor reactions (e.g. how sharply a robot turns given a certain pattern of activation of the infrared sensors) are crucial for determining whether the robot will be able to solve the problem or not. Indeed, small differences in the robot's motor response tend to accumulate in time and might prevent the robot from producing successful behaviour (e.g. might prevent the robot producing a behavioural attractor close to cylindrical objects).

This experiment clearly exemplifies some important aspects which characterize all adaptive behavioural system, i.e. systems which are embodied and situated and which have been designed or adapted so to exploit the properties that emerge from the interaction between their control system, their body, and the external environment. In particular, it demonstrates how required behavioural and cognitive skills (i.e. object categorization skills) might emerge from the fine-grained interaction between the robot's control system, body, and the external environment without the need for dedicated control mechanisms. Moreover, it demonstrates how the relation between the control rules which mediate the interaction between the robot body and the environment and the behavioural skills exhibited by the agents are rather indirect. This means, for example, that an external human observer can hardly predict the behaviours that will be produced by the robot, before observing the robot interacting with the environment, even on the basis of a complete description of the characteristics of the body, of the control rules, and of the environment.

2.2 Behaviour and cognition as phenomena originating from the interaction between coupled dynamical processes

Up to this point we restricted our analysis to the dynamics originating from the agent's control system, agents' body, and environmental interactions. However, the body of an agent, its control system, and the environment might have their own dynamics (dotted arrows in Figure 1). For the sake of clarity, we will refer to the dynamical processes occurring within the agent control system, within the agent body, or within the environment as *internal dynamics* and to the dynamics originating from the agent/body/environmental interaction as *external dynamics*. In cases in which agents' body, agents' control system, or the environment have their own dynamics, behaviour should be characterized as a property emerging from the combination of several coupled dynamical processes.

The existence of several concurrent dynamical processes represents an important opportunity for the possibility to exploit emergent features. Indeed, behavioural and cognitive skills might emerge not only from the external dynamics, as we showed in the previous section, but also from the internal dynamical processes or from the interaction between different dynamical processes.

As an example which illustrates how complex cognitive skills can emerge from the interaction between a simple agent/body/environmental dynamic and a simple agent's internal dynamic consider the case of a wheeled robot placed in a maze environment (Figure 4) which has been trained to: (a) produce a wall-following behaviour which allows the robot to periodically visit and re-visit all environmental areas, (b) identify a target object constituted by a black disk placed in a randomly selected position in the environment for a limited time duration, and (c) recognize the location in which the target object was previously found every time the robot re-visits the corresponding location (Gigliotta & Nolfi, 2008).

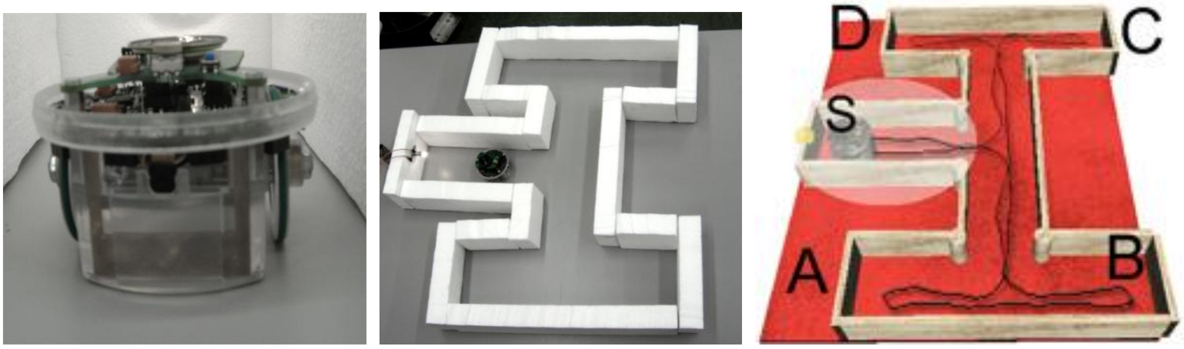

Figure 4. **Left**: The e-puck robot developed at EPFL, Switzerland (http://www.e-puck.org/). **Centre**: The environment which has a size of 52cm by 60cm. The light produced by the light bulb located on the left side of the central corridor cannot be perceived from the other two corridors. **Right**: The motor trajectory produced by the robot during a complete lap of the environment.

The robot has infrared sensors (which provide information about nearby obstacles), light sensors (which provide information about the light gradient generated by the light bulb placed in the central corridor), ground sensors (which detect the colour of the ground), two motors (which control the desired speed of the two corresponding wheels), and one additional output unit which should be turned on when the robot re-visit the environmental area in which the black disk was previously found. The robot's controller consists of a three-layer neural network which includes a layer of sensory neurons (which encode the state of the corresponding sensors), a layer of motor neurons which encode the state of the actuators, and a layer of internal neurons which consist of leaky integrators operating at tuneable time scale (Beer, 1995; Gigliotta & Nolfi, 2008). The free parameters of the robot's neural controllers (i.e. the connection weights, and the time constant of the internal neurons which regulate the time rate at which these neurons change their state over time) were adapted through an evolutionary technique (Nolfi & Floreano, 2000).

By analysing the evolved robot the authors observed how it is able to generate a spatial representation of the environment and of its location in the environment while it is situated in the environment itself. Indeed, while the robot travel by performing different laps of the environment (see Figure 4, right), the states of the two internal neurons converge on a periodic limit cycle dynamic in which different states correspond to different locations of the robot in the environment (Figure 5).

As we mentioned above, the ability to generate this form of representation, that allows the robot to solve its adaptive problem, originates from the coupling between a simple robot's internal dynamics and a simple robot/body/environmental dynamics. The former dynamics is characterized by the fact that the state of the two internal neurons tends to move slowly toward different fixed point attractors, in the robot's internal dynamics, which correspond to different types of sensory states exemplified in Figure 5. The latter dynamics originate from the fact that different types of sensory states last for different time durations and alternate with a given order while the robot moves in the environment. The interaction between these two dynamical processes leads to a transient dynamics of an agent's internal state that moves slowly toward the current fixed point attractor without never fully reaching it (thus preserving information about previously experienced sensory states, the time duration of these states, and the order with which they have been experienced). The coupling between the two dynamical processes originates from the fact that the free parameters which regulate the agent/environmental dynamics (e.g. the trajectory and the speed with which the robot moves in the environment) and the agent internal dynamics (e.g. the direction and the speed with which the internal neurons change their state) have been co-adapted and co-shaped during the adaptive process.

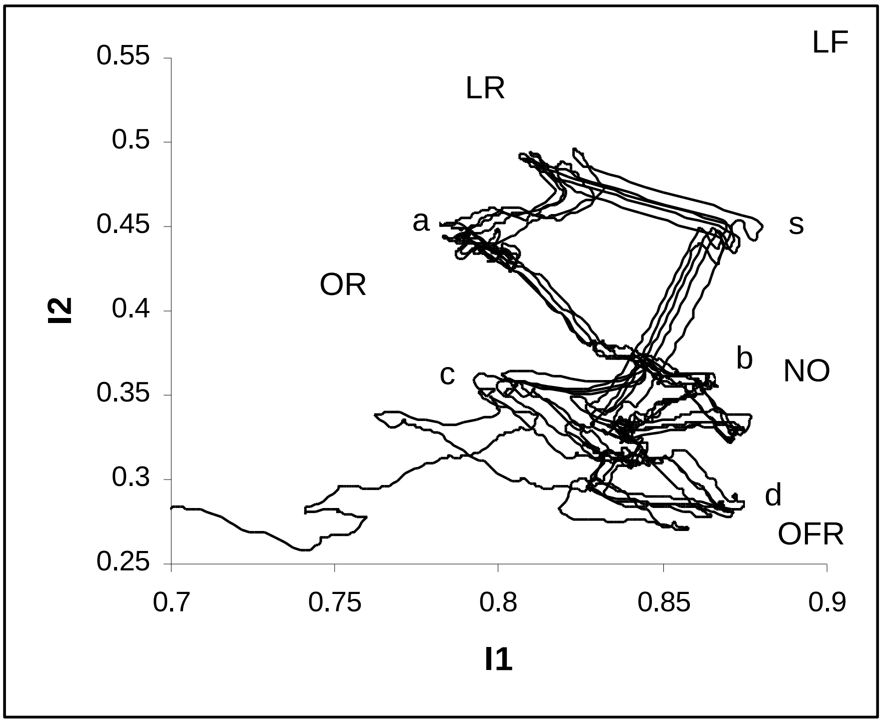

Figure 5. The state of the two internal neurons (i1 and i2) of the robot recorded for 330s while the robot performs about 5 laps of the environment. The s, a, b, c, and d labels indicate the internal states corresponding to five different positions of the robot in the environment shown in Figure 4. The other labels indicate the position of the fixed point attractors in the robot's internal dynamics corresponding to five types of sensory states experienced by the robot when it detects: a light in its frontal side (LF), a light on its rear side (LR), an obstacle on its right and frontal side (OFR), an obstacle on its right side (OR), no obstacles and no lights (NO).

For related works which show how navigation and localization skills might emerge from the coupling between agent's internal and external dynamics, see (Tani and Fukumura, 1997). For other works addressing other behavioural/cognitive capabilities see Beer (2003) for what concerns categorization, Goldenberg at al. (2004) and Slocum et al. (2000) for what concerns selective attention, Sugita and Tani (2005) for what concern language and compositionality.

2.3 Behaviour and cognition as phenomena with a multi-level and multi-scale organization

Another fundamental feature that characterizes behaviour is the fact that it is a multi-layer system with different levels of organizations extending at different time scales (Keijzer, 2001; Nolfi, 2005). More precisely, as exemplified in Figure 6, the behaviour of an agent or of a group of agents involve both lower- and higher- level behaviours that extend for shorter or longer time spans, respectively. Lower-level behaviours arise from few agent/environmental interactions and short term internal dynamical processes. Higher-level behaviours, instead, arise from the combination and interaction of lower-level behaviours and/or from long term internal dynamical processes.

The multi-level and multi-scale organization of agents' behaviour plays important roles: it is one of the factors which allow agents to produce functionally useful behaviour without necessarily developing dedicated control mechanisms (Brooks, 1991; Nolfi, 2005), it might favour the development of new behavioural and/or cognitive skills thanks to the recruitment of pre-existing capabilities (Marocco and Nolfi, 2007), it allow agents to generalize their skills in new task/environmental conditions (Nolfi, 2005).

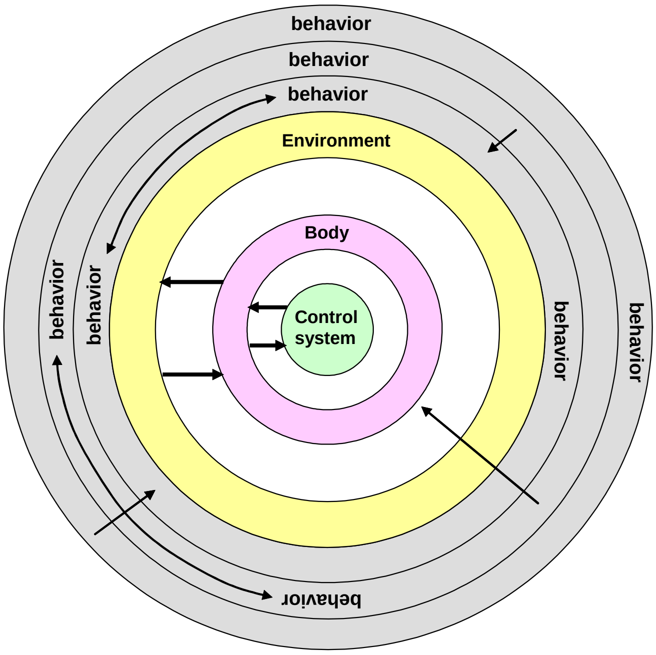

Figure 6. A schematic representation of multi-level and multi-scale organization of behaviour. The behaviours represented in the inner circles represent elementary behaviours which arise from fine-grained interactions between the control system, the body, and the environment, and which extend over limited time spans. The behaviours represented in the external circles represent higher-level behaviours which arise from the combination and interaction between lower-level behaviours and which extend over longer time spans. The arrows which go from higher-level behaviour toward lower levels indicate the fact that the behaviours currently exhibited by the agents later affect the lower-level behaviours and/or the fine-grained interaction between the constituting elements (agent's control system, agent's body, and the environment)

An exemplification of how the multi-level and multi-scale organization of behaviour allow agents to generalize their skill in new environmental conditions is represented by the experiments carried our by Baldassarre et al. (2006) in which the authors evolved the control system of a group of robots assembled into a linear structure (Figure 7) for the ability to move in a coordinated manner and for the ability to display a coordinated light approaching behaviour.

Each robot (Mondada et al, 2004) consists of a mobile base (chassis) and a main body (turret) that can rotate with respect to the chassis around a vertical axis. The chassis has two drive mechanisms that control the two corresponding tracks and teethed wheels. The turret has one gripper, which allows robots to assemble together and to grasp objects, and a motor controlling the rotation of the turret with respect to the chassis. Robots are provided with a traction sensor, placed at the turret-chassis junction, that detects the intensity and the direction of the force that the turret exerts on the chassis (along the plane orthogonal to the vertical axis) and light sensors. Given that the orientations of individual robots might vary and given that the target light might be out of sight, robots need to coordinate to choose a common direction of movement and to change their direction as soon as one or few robots start to detect a light gradient.

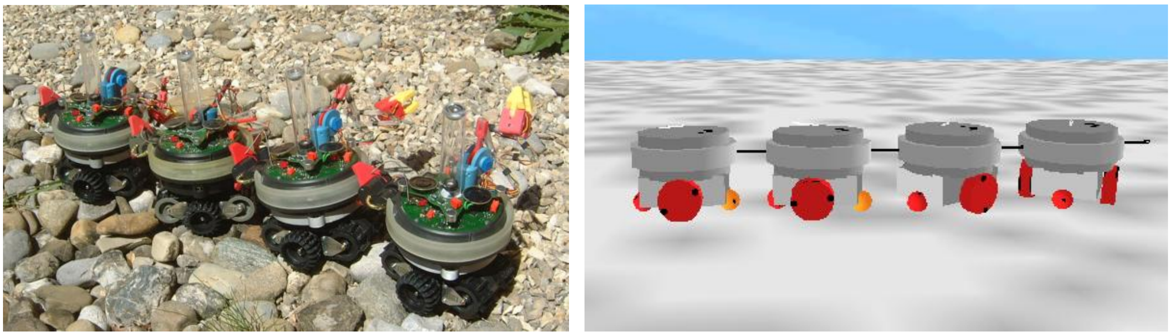

Figure 7. **Left:** Four robots assembled into a linear structure. **Right**: A simulation of the robots shown in the left part of the figure.

Evolved individuals show the ability to negotiate a common direction of movement and by approaching light targets as soon as a light gradient is detected. By testing evolved robots in different conditions the authors observed that they are able to generalize their skills in new conditions and also to spontaneously produce new behaviours which have not been rewarded during the evolutionary process. More precisely, groups of assembled robots display a capacity to generalize their skills with respect to the number of robots which are assembled together and to the shape formed by the assembled robots. Moreover, when the evolved controllers are embodied in eight robots assembled so as to form a circular structure and situated in the maze environment shown in Figure 8, the robots display an ability to collectively avoid obstacles, to rearrange their shape so to pass through narrow passages, and to explore the environment. The ability to display all these behavioural skills allow the robots to reach the light target even in large maze environments, i.e. even in environmental conditions which are rather different from the conditions that they experienced during the training process (Figure 8).

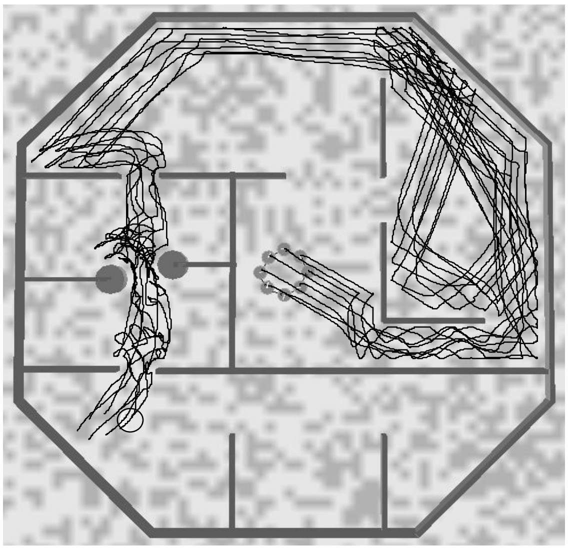

Figure 8. The behaviour produced by eight robots assembled into a circular structure in a maze environment including walls and cylindrical objects (represented with grey lines and circles). The robots start in the central portion of the maze and reach the light target located in the bottom-left side of the environment (represented with an empty circle) by exhibiting a combination of coordinated-movement behaviours, collective obstacle-avoidance, and collective light-approaching behaviours. The irregular lines, that represent the trajectories of the individual

robots, show how the shape of the assembled robots changes during motion by adapting to the local structure of the environment.

By analysing the behaviour displayed by the evolved robots tested in the maze environment, a complex multi-level organization can be observed. The simpler behaviours that can be identified consist of low level individual behaviours which extend over short time spans:

1) A *move-forward behaviour* which consists of the individuals' ability to move forward when the robot is coordinated with the rest of the team, is oriented toward the direction of the light gradient (if any), and does not collide with obstacles. This behaviour results from the combination of: (a) a control rule which produces a move forward action when the perceived traction has a low intensity and when difference between the intensity of the light perceived on the left and the right side of the robot is low, and (b) the sensory effects of the execution of the move forward action selected mediated by the external environment (which does not produce a variation of the state of the sensors until the conditions that should be satisfied to produce this behaviour hold).
2) A *conformistic behaviour* which consists of the individuals' ability to conform its orientation with that of the rest of the team when the two orientations differ significantly. This behaviour results from the combination of: (a) a control rule that makes the robot turns toward the direction of the traction when its intensity is significant, and (b) the sensory effects produced by the execution of this action mediated by the external environment, that lead to a progressive reduction of the intensity of the traction until the orientation of the robot conforms with the orientation of the rest of the group.
3) A *phototaxis behaviour* which consists of the individuals' ability to orient toward the direction of the light target. This behaviour results from the combination of: (a) a control rule that makes the robot turns toward the direction in which the intensity of the light gradient is higher, and (b) the sensory effects produced by the execution of this action mediated by the external environment, that lead to a progressive reduction of the difference in the light intensity detected on the two sides of the robot until the orientation of the robot conforms with the direction of the light gradient.
4) An *obstacle-avoidance behaviour* which consists of the individuals' ability to change direction of motion when the execution of a motor action produced a collision with an obstacle. This behaviour results from the combination of: (a) the same control rule that led to behaviour #2 and that makes the robot turn toward the direction of the perceived traction (which in this case is caused by the collision with the obstacle, while in the case of behaviour #2 it is caused by the forces exerted by the other assembled robots), and (b) the sensory effects produced by the execution of the turning action mediated by the external environment, that make the robot turn until collisions no longer prevent the execution of a moving forward behaviour.

The combination and the interaction between these four behaviours produce the following higher-level collective behaviours that extend over a longer time span:

5) A *coordinated-motion behaviour* that consists in the ability of the robots to negotiate a common direction of movement and to keep moving along such direction by compensating further misalignments originating during motion. This behaviour emerges from the combination and the interaction of the conformistic behaviour (which plays the main role when robots are misaligned) and the move-forward behaviour (which plays the main role when robots are aligned).

6) A *coordinated-light-approaching behaviour* which consists in the ability of the robots to co-ordinately move toward a light target. This behaviour emerges from the combination of the conformistic, the move-forward, and the phototaxis behaviours (which is triggered when the robots detect a light gradient). The relative importance of the three control rules which lead to the three corresponding behaviours depends both on the strength of the corresponding triggering condition (i.e. the extent of lack of traction forces, the intensity of traction forces, and the intensity of the light gradient, respectively) and on a priority relation among behaviours (i.e. the fact that the conformistic behaviour tends to play a stronger role than the phototaxis behaviour).
7) A *coordinated-obstacle-avoidance behaviour* which consists in the ability of the robots to co-ordinately turn to avoid nearby obstacles. This behaviour arises as the result of the combination of the obstacle avoidance, the conformistic and the move-forward behaviours.

The combination and the interaction between these behaviours leads in turn to the following higher-level collective behaviours that extend over still longer time spans:

8) A *collective-exploration-behaviour* that consists in the ability of the robots to visit different areas of the environment when the light target cannot be detected. This behaviour emerges from the combination of the coordinated-motion behaviour and the coordinated-obstacle-avoidance behaviour that ensures that the assembled robots can move in the environment without getting stuck and without entering into limit cycle trajectories.
9) A *shape-re-arrangement behaviour* which consists in the ability of the assembled robots to dynamically adapt their shape to the current structure of the environment so to pass through narrow passages especially when the passages to be negotiated are in the direction of the light gradient. This behaviour emerges from the combination and the interaction between coordinated motion and coordinated-light-approaching behaviours mediated by the effects produced by relative differences in motion between robots resulting from the execution of different motor actions and/or from differences in the collisions. The fact that the shape of the assembled robots adapt to the current environmental structure so as to facilitate the overcoming of narrow passages can be explained by considering that collisions produce a modification of the shape that affects the relative positions of the colliding robots, in particular with respect to the axis of the narrow passage.

The combination and the interaction of all these behaviours leads to a still higher-level behaviour:

10) A *collective-navigation-behaviour* which consists in the ability of the assembled robots to navigate toward the light target by producing coordinated movements, exploring the environment, passing through narrow passages, and producing a coordinated-light-approaching behaviour (Figure 8).

This analysis illustrates two important mechanisms that explain the remarkable generalization abilities of these robots. The first mechanism consists in the fact that the control rules that regulate the interaction between the agents and the environment so as to produce certain behavioural skills in certain environmental conditions, will produce different but related behavioural skills in other environmental conditions. In particular, the control rules that generate the behaviours #5 and #6, for which evolving robots have been evolved in an environment without obstacles, also produce behaviour #7 in an environment with obstacles. The second mechanism consists in the fact that the development of certain behaviours at a given level of organization that extend for a given time span will, through their combination and interaction, automatically lead to the exhibition of related higher-level behaviours extending over longer time spans (even if these higher-level behaviours have not been rewarded during the adaptation process). In particular, the combination and the interaction of behaviours #5, #6 and #7 (that have been rewarded during the evolutionary process or that arise from the same control rules that lead to the generation of rewarded behaviours) automatically lead to the production of behaviours #8, #9, and #10 (that have not been rewarded). Obviously, there is no warranty that the new behaviours obtained as a result of these generalization processes will play useful functions. However, the fact that these behaviours are related to the other functional behavioural skills implies that the probabilities that these new behaviours will play useful functions are significant.

In principle, these generalization mechanisms can also be exploited by agents during their adaptive process to generate behavioural skills which play new functionalities and which emerge from the combination and the interaction between pre-existing behavioural skills playing different functions.

2.4 On the top-down effect from higher to lower levels of organization

In the previous sections we have discussed how the interactions between the agents' body, the agents' control system, and the environment lead to behavioural and cognitive skills and how such skills have a multi-level and multi-scale organization in which the interaction between lower-level skills leads to the emergence of higher-level skills. However, higher-level skills also affect lower-level skills up to fine-grained interactions between the constituting elements (agents' body, agents' control system, and environment). More precisely, the behaviours that originate from the interaction between the agent and the environment and from the interaction between lower-level behaviours, later affect the lower-level behaviours and the interaction from which they originate. These bottom-up and top-down influences between different levels of organization can lead to circular causality (Kelso, 1995) where high-level processes act as independent entities that constraint the lower-level processes from which they originate.

One of the most important effects of this top-down influence consists in the fact that the behaviour exhibited by an agent constrains the type of sensory patterns that the agent will experience later on (i.e. constrains the fine-grained agent/environmental interactions that determine the behaviour that will be later exhibited by the agent). Since the complexity of the problem faced by an agent depends on the sensory information experienced by the agent itself, this top-down influence can be exploited in order to turn hard problems into simple ones.

One neat demonstration of this type of phenomena is given by the experiments conducted by Marocco and Nolfi (2002) in which a simulated finger robot with six degree of freedom provided with sensors of its joint positions and with rough touch sensors is asked to discriminate between cubic and spherical objects varying in size. The problem is not trivial since, in general terms, the sensory patterns experienced by the robot do not provide clear regularities for discriminating between the two types of objects. However, the type of sensory states which are experienced by the agent also depend on the behaviour previously exhibited by the agent itself --- agents exhibiting different behaviour might face simpler or harder problems. By evolving the robots in simulation for the ability to solve this problem and by analyzing the complexity of the problem faced by robots of successive generations, the authors observed that the evolved robots manage to solve their adaptive problem on the basis of simple control rules which allow the robot to approach the object and to move following the surface of the object from left to right, independently of the object shape. The exhibition of this behaviour in interaction with objects characterized by a smooth or irregular surface (in the case of spherical or cubic objects, respectively) ensures that the same control rules lead to two types of behaviours depending on the type of the object. These behaviours consist in following the surface of the object and then moving away from the object in the case of spherical objects, and in following the surface of the object by getting stuck in a corner in the case of cubic objects. The exhibition of these two behaviours allows the agent to experience rather different proprioceptors states as a consequence of having interacted with spherical or cubic objects that nicely encode the regularities that are necessary to differentiate the two types of object.

For other examples which shows how adaptive agents can exploit the fact that behavioural and cognitive processes which arise from the interaction between lower-level behaviours or between the constituting elements later affect these lower-level processes, see Scheier et al. (1998), Nolfi (2002), Beer, (2003).

3. Behaviour and cognition as adaptive systems

The complex system nature of behaviour and cognition can also help us to understand why embodied and situated agents are difficult to handcraft by a human designer. The fact that the relation between the characteristics of the robot and of the environment and the behavioural and cognitive skills that emerge from robot/environmental interactions is so complex and indirect, in fact, implies that a human observer cannot predict the behaviour that will be exhibited by the robot (without having had the chance to observe the behaviour exhibited by the robot in interaction with the environment) even on the basis of a detailed description of the characteristics of the robot and of the environment and of the rules that regulate the robot/environmental interaction. This in turn implies that the task faced by the robot designer which consists in inferring the fine-grained characteristics that the robot and the robot's control system should have in order to produce a given desired behaviour is extremely hard if not completely hopeless (Funes et al., 2003; Nolfi, 2005).

On the other hand, as we have shown in the examples illustrated in the previous section, the fine-grained characteristics of the robot can be effectively shaped through an adaptive process. In particular, robots able to display the required behavioural skills can be developed through an evolutionary technique (Nolfi and Floreano, 2000) in which the fine-grained characteristics of the robots' control system (and eventually of the robot body) are varied randomly and variations are retained or discarded on the basis of the their effect at the level of the global behaviour exhibited by the robot (or by the robots). This type of adaptive process in fact is able to discover and retain the characteristics of the robot or robots that, in interaction with their other characteristics and with the environment, can lead to useful emergent properties through a trial and error process that does not require computing the relation between the fine-grained characteristics and the higher-level properties that emerge from their interactions.

From an engineering point of view these considerations imply that the development of effective embodied agents (i.e. agents able to exploit the properties emerging from the interactions) require new methods. One possibility consists in *artificial adaptive* methods, as the evolutionary method illustrated in section 2, in which the fine-grained characteristics of the agents are modified through a trial and error process and in which variations are retained or discarded on the basis of their effects at the level of the global behaviour exhibited by the agents in the environment. A second possibility might consist in the development of new *design for emergence* methods that could provide a way to handcraft agents able to exploit emergent properties. Whether effective methods of this kind can be developed or not represents an open question at the moment. For a preliminary attempt to formulate design for emergence principles see Pfeifer and Scheier (1999).

The complex system nature of behaviour and cognition and the consequent difficulty of identifying the characteristics that the fine-grained interaction between the robot and the environment should have in order to produce a desired global behaviour do not imply that the behavioural and cognitive skills exhibited by adaptive agents are intrinsically unpredictable. It only implies that solutions analogous to those that can be generated through an adaptive process cannot be generated by human designers. In other words, adaptive solutions tend to be qualitatively different from hand-crafted solutions. Moreover, the complex system nature of behaviour and cognition does not imply that adapted solutions are inscrutable or inexplicable as we have also demonstrated in the example reported above. Indeed, the detailed analysis of adapted natural and artificial solutions represents the main instrument that we have to identify the general principles which characterize embodied intelligence.

From a modelling point of view, the complex system nature of behaviour and cognition implies that the characteristics of living organisms crucially depends on the characteristics of the processes that determine how they change, phylogenetically and ontogenetically, as they adapt to their environment. In other words these considerations suggest that adaptivity should be considered an equally fundamental property of behavioural systems as their embodied and situated nature.

4. Discussion and conclusion

In this paper we illustrated how the behavioural and cognitive skills displayed by embodied and situated agents can be properly characterized as a complex system with multi-level and multi-scale organization and involving both bottom-up and top-down influences that emerge from several fine-grained interactions between the agent (or the agents) and the environment.

The complex systems nature of adaptive agents that are embodied and situated has important implications that constrain the organization of these systems and the dynamics of the adaptive process through which they develop their skills.

With respect to the organization of these systems, this complexity implies that agents' behavioural and/or cognitive skills (at any stage of the adaptive process) cannot be traced back to anyone of the three foundational elements (i.e. the body of the agents, the control system of the agents, and the environment) in isolation but should rather be characterized as properties which emerge from the interactions between these three elements and the interaction between behavioural and cognitive properties emerging from the former interactions at different levels of organizations. Moreover, it implies that 'complex' behavioural or cognitive skills might emerge from the interaction between simple properties or processes. A similar argument can be made for what concerns the course of the adaptive process which cannot be traced back to the three foundational elements in isolation but rather depends on the interactions between these elements and on the interaction between the higher-level behavioural and cognitive processes that emerge from lower-level interactions. Indeed new behavioural skills might originate during the adaptive process both as a result of variations of the characteristics of the robots and as a result of variations of the physical or social environment.

With respect to agents' adaptive process, the development of new 'complex' skills does not necessarily require the development of new morphological features or new dedicated control mechanisms. Indeed, behavioural or cognitive skills might arise spontaneously as a result of the interaction between properties serving different functions and/or as a result of simple additional characteristics thanks to the possibility to exploit the emergent results of the interaction between these new characteristics with the other pre-existing characteristics and skills.

The study of adaptive behaviour in artificial agents that has been reviewed in this paper has important implications both from an engineering point of view (i.e. for progressing in our ability to develop effective machines) and from a modelling point of view (i.e. for understanding the characteristics of biological organisms).

In particular, from an engineering point of view, progress in our ability to develop complex adaptive behavioural and cognitive systems can lead to development of new artefacts playing useful functionalities.

From a modelling point of view, progress in our ability to model and analyze behavioural and cognitive skills in adaptive embodied agents can improve our understanding of the general mechanisms behind animal and human intelligence.

More specifically, the analysis of these systems can allow us to identify which are the fundamental properties of natural intelligence that characterize a large variety of species independently from their specific morphological, neuro-physiological, and ecological characteristics.

Moreover, this type of research can help us to formulate new theoretical concepts and terms which can allow us to model and describe the key characteristics of natural intelligence. For example, it can contribute to better define two fundamental aspects of natural (and artificial) intelligence: *morphological computation* and *sensory-motor coordination*. The former concept refers to the fact that the characteristics of the body of an agent (from the overall morphological structure to the fine-grained characteristics of the body such us the exact position of the receptors or the degree of elasticity of different body parts) strongly determine the skills that an agent might exhibit and the complexity of the control mechanisms which are required to produce such skills. The latter concept refers to the fact that the sensory states experienced by an agent are determined not only by the characteristics of the environment and by the agent/environmental relation but also by the motor actions previously performed by the agent itself. Indeed, behavioural and cognitive skills might emerge from the dynamical process arising from the agent/environmental interactions without the need of dedicated control mechanisms provided that the rules that regulate how the agent reacts to sensory states have been shaped to appropriately exploit the properties emerging from the agent/environmental interactions.

Finally, the comprehension of the complex system nature of behavioural and cognitive skills illustrated in this paper can allow us to better define the notion of embodiment and situatedness, that are universally recognized as central aspects in the study of natural and artificial intelligence but that still lack a clear and uncontroversial definition. Possessing a body and being in a physical environment certainly represent a pre-requisite for considering an agent embodied and situated. However, a more useful definition of embodiment (or of true degree of embodiment) can be given in terms of the extent to which a given agent exploits its body characteristics to solve its adaptive problem (i.e. the extent to which its body structure can be adapted to the problems to be solved, or in other words, the extent to which its body performs morphological computation). Similarly, a more useful definition of situatedness (or of true degree of situatedness) can be given in terms of the extent to which an agent exploits its interaction with the physical and social environment and the properties originating from this interaction to solve its adaptive problems. For the sake of clarity we can refer to the former definition of the terms (i.e. merely possessing a physical body and being situated in a physical environment) as embodiment and situatedness in the weak sense, and to the latter definitions as embodiment and situatedness in a strong sense.

References

Baldassarre G., Parisi D., & Nolfi S. (2006). Distributed coordination of simulated robots based on self-organisation, Artificial Life, (12) 3: 289-311.

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.

Bongard J.C., & Paul C. (2001). Making evolution an offer it can't refuse: Morphology and the extradimensional bypass. In J. Keleman & P. Sosik (Eds.), Proceedings of the Sixth European Conference on Artificial Life, Lecture Notes in Artificial Intelligence, 2159. Berlin: Springer-Verlag.

Brooks R.A. (1991). Intelligence without reason. In J. Mylopoulos & R. Reiter (Eds.), Proceedings of 12th International Joint Conference on Artificial Intelligence. San Mateo, CA: Morgan Kaufmann.

Chiel H.J. & Beer R.D. (1997). The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. Trends in Neurosciences 20:553-557.

Endo I., Yamasaki F., Maeno T. & Kitano, H. (2002). A method for co-evolving morphology and walking patterns of biped humanoid robot. In Proceedings of the IEEE Conference on Robotics and Automation.

Funes P., Orme B. & Bonabeau E. (2003). Evolving emergent group behaviours for simple humans agents. In P. Dittrich & J. T. Kim (Eds.) Proceedings of the Seven European Conference on Artificial Life .

Gigliotta O. & Nolfi S. (2008). On the coupling between agent internal and agent/environmental dynamics: Development of spatial representations in evolving autonomous robots. Connection Science, 16: 148 – 165.

Goldenberg E., Garcowski J. & Beer R.D. (2004). May we have your attention: Analysis of a selective attention task. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam & J.-A. Meyer (Eds.), From Animals to Animats 8: Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior. Cambridge, MA: MIT Press.

Keijzer F. (2001). Representation and behavior. London, UK: MIT Press.

Nolfi S., Marocco D. (2002). Active perception: A sensorimotor account of object categorization. In From Animals to Animats 7. In B. Hallam, D. Floreano, J. Hallam, G. Hayes, J-A. Meyer (eds.) Proceedings of the VII International Conference on Simulation of Adaptive Behavior. Cambridge, MA: MIT Press, pp. 266-271.

Marocco D. & Nolfi S. (2007). Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. Connection Science, 19 (1): 53-74.

McGeer, T. (1990). Passive walking with knees. In Proceedings of the IEEE Conference on Robotics and Automation, volume 2:1640–1645.

Mondada F., Franzi E., & Ienne P. (1993). Mobile Robot miniaturisation: A tool for investigation in control algorithms. In: Proceedings of the Third International Symposium on Experimental Robotics, Kyoto, Japan.

Mondada F., Pettinaro G., Guigrard A., Kwee I., Floreano D., Denebourg J-L, Nolfi S., Gambardella L.M., Dorigo M. (2004). Swarm-bot: A new distributed robotic concept. Autonomous Robots, 17 (2-3): 193-221.

Nolfi S. & Floreano D. (2000). Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. Cambridge, MA: MIT Press/Bradford Books.

Nolfi S. (2002). Power and limits of reactive agents. Neurocomputing, 49:119-145.

Nolfi S. (2005). Behaviour as a complex adaptive system: On the role of self-organization in the development of individual and collective behaviour. Complexus, 2 (3-4):195-203.

Pfeifer R., Iida F. & Gómez, G. (2006), Morphological computation for adaptive behavior and cognition. In: International Congress Series, volume 1291:22-29.

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

Scheier C., Pfeifer R. & Kunyioshi Y. (1998) Embedded neural networks: exploiting constraints. Neural Networks (11):1551-1596.

Schmitz A., Gómez G., Iida F. & Pfeifer R. (2007). On the robustness of simple speed control for a quadruped robot. In: Proceeding of the International Conference on Morphological Computation.

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 and S. Wilson (Eds.), From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior. Cambridge, MA: MIT Press.

Sugita Y. & J. Tani (2005). Learning semantic combinatoriality from the interaction between linguistic and behavioral processes. Adaptive Behavior, (13) 1:33-52.

Tani J. & Fukumura N. (1997). Self-organizing internal representation in learning of navigation: A physical experiment by the mobile robot Yamabico. Neural Networks, 10 (1): 153-159.

Vaughan, E., Di Paolo, E.A. & Harvey, I., (2004). The evolution of control and adaptation in a 3D powered passive dynamic walker. In J. Pollack, M. Bedau, P. Husband, T. Ikegami & R. Watson (Eds). Proceedings of the Ninth International Conference on the Simulation and Synthesis of Living Systems. Cambridge, MA: MIT Press