Institute of Psychology
C.N.R. - Rome

Evolutionary Robotics: Exploiting the full power of self-organization

Stefano Nolfi
Institute of Psychology,
Division of Neural Systems and Artificial Life
National Research Council
Viale Marx 15, Roma, Italy
e-mail: stefano@kant.irmkant.rm.cnr.it
http://kant.irmkant.rm.cnr.it/nolfi.html

June 1998

To appear on: *Connection Science*, 10 (3-4), special issue on BioRobotics.

Evolutionary Robotics: Exploiting the full power of self-organization

Stefano Nolfi
Institute of Psychology, National Research Council
Viale Marx 15, 00137, Roma, Italy
Voice:++39-6-86090231, Fax:++39-6-824737
e-mail:stefano@kant.irmkant.rm.cnr.it
http://kant.irmkant.rm.cnr.it/nolfi.html

Abstract

In this paper I claim that one of the main characteristics that makes the Evolutionary Robotics approach suitable for the study of adaptive behavior in natural and artificial agents is the possibility to rely largely on a self-organization process. Indeed by using Artificial Evolution the role of the designer may be limited to the specification of a fitness function which measures the ability of a given robot to perform a desired task. From an engineering point of view the main advantage of relying on self-organization is the fact that the designer does not need to divide the desired behavior into simple basic behaviors to be implemented into separate layers (or modules) of the robot control system. By selecting individuals for their ability to perform the desired behavior as a whole, simple basic behaviors can emerge from the interaction between several processes in the control system and from the interaction between the robot and the environment. From the point of view of the study of natural systems, the possibility of evolving robots that are free to select their way to solve a task by interacting with their environment may help us to understand how natural organisms produce adaptive behavior. Finally, the attempt to scale up to more complex tasks may help us to identify what the critical features of Natural Evolution are which allowed the emergence of the extraordinary variety of highly adapted life forms present on the planet.

1. Introduction

Evolutionary Robotics is the attempt to develop robots and their sensorimotor control systems through an automatic design process involving Artificial Evolution (Cliff, Harvey & Husband, 1993, Nolfi, Floreano, Miglino & Mondada, 1994).

Evolutionary Robotics approaches are based on genetic algorithms (Holland, 1975). An initial population of different "genotypes", each coding the control system (and possibly the morphology) of a robot, are created randomly. Each robot is evaluated in the environment and to each robot is assigned a score (fitness) that measures its ability to perform a desired task. Those robots that have obtained higher fitness values are allowed to reproduce (sexually or non-sexually) by generating copies of their genotypes with the addition of changes introduced by some genetic operators (e.g. mutations, crossover, duplication, etc.). By repeating this process for several generations one can observe progress in the fitness values of the population.

Generally speaking we can say that this approach shares many of its characteristics with other approaches such as Behavior-Based Robotics, Robot Learning, and Artificial Life.

The *Behavior-Based Robotics* approach is based upon the idea of providing the robot with a range of simple basic behaviors and letting the environment determine which basic behavior should have control at any given time (Brooks, 1986). Basic behaviors are implemented in separate sub-parts of the control system and a coordination mechanism is responsible for determining which basic behavior should be activated or inhibited at any particular time. In this approach, as in Evolutionary Robotics, the environment plays a central role by determining which basic behavior is active at any given time. Moreover, behavioral modules and the coordination mechanism are usually designed through a trial and error process in which the designer progressively changes them by testing the resulting behavior in the environment. However, Evolutionary Robotics, by relying on an automatic evaluation process, usually makes a larger use of the trial and error process described above. Moreover, while in the Behavior-Based approach the breakdown of the desired behavior into simpler basic behaviors is accomplished intuitively by the designer, in the Evolutionary Robotics approach this is often the result of a self-organizing process (Nolfi, 1997b). Indeed, the entire organization of the evolving system, including its organization into sub-components, is the result of an adaptation process that usually involves a large number of evaluations of the interactions between the system and the environment (I will return to this issue in section 2).

Robot Learning is based on the idea that a control system (typically a neural network) can be trained using incomplete data and then allowed to rely on its ability to generalise the acquired knowledge to novel circumstances. The general motivation behind this approach is that it may produce better results than approaches based on explicit design, given the well-known difficulties of engineering behavioral systems (see below). In some cases the neural control system learns to perform a mapping between sensory inputs and motor states while in other cases learning is used to develop subsystems of the controller. Different learning algorithms can be used for this purpose: i.e. back-propagation learning (Rumelhart, Hinton & Williams, 1986); reinforcement learning (Barto, Bradtke & Singh, 1995); classifier systems (Booker, Goldberg & Holland, 1989); Kohonen Self-Organized Maps (Kohonen, 1982), etc. These algorithms impose different constraints on the type of architecture that can be used and on the quantity and quality of the supervision required from the designer. For example, if used to learn a mapping from sensors to motors, back-propagation learning requires that the designer provides an explicit indication of the correct activation of each motor in each cycle. Reinforcement learning instead only needs an evaluation of how good or bad the robot is doing at the time[1]. Evolutionary Robotics shares with these approaches the emphasis on self-organization. Indeed, Artificial Evolution may be described as a form of learning. However, Evolutionary Robotics differs from Robot Learning in two respects. First, the amount of supervision required by evolution is generally much lower - only a general evaluation of how well an individual accomplishes the desired task is required. Second, the evolutionary method does not introduce any constraint on what can be part of the self-organization process. Indeed, the characteristics of the sensors and of the

[1] Unsupervised learning algorithms, such as Hebbian learning or Kohonen Self-Organized Maps, do not need any supervision at all. However, although these algorithms can be used to train subparts of the control systems, they cannot map sensory inputs into motor states by themselves unless the architecture is constrained in such a way that the system converges into certain solutions. For an example in which they can be successfully used to learn goal directed behavior in conjunction with evolution see (Floreano and Mondada, 1996).

actuators (Cliff & Miller, 1996), the shape of the robot (Lund, Hallam & Lee, 1997), and the architecture of the control system (see Kodjabachian & Meyer, 1998) can be included into the evolutionary process.

Artificial Life represents an attempt to understand all life phenomena through their reproduction in artificial systems (typically through their simulation on a computer). More specifically, Artificial Life provides a unique framework for studying how entities at different levels of organization (molecules, organs, organisms and populations) interact among themselves (Parisi, 1997) although, of course, at the cost of introducing crude simplifications. To attain this ambitious goal, Artificial Life relies on the theory of Complex Systems and, from the experimental point of view, on the power of computers. A complex system is a system that can be described at different levels, in which global properties at one level emerge from the interaction of a number of simple elements at lower levels. Global properties are emergent in the sense that, even if they result from nothing else but local interactions among the elements, they cannot be predicted or inferred from a knowledge of the elements or of the rules by which the elements locally interact, given the high nonlinearity of these interactions. Evolutionary Robotics shares most of these characteristics with Artificial Life, but it also stresses the importance of using a physical device (e.g. a robot) instead of a simulated agent. By using real robots, several additional factors due to the physical properties of the robot and of the environment must be taken into account (e.g. friction, inertia, ambient light, noise, etc.) (Brooks, 1992). Moreover, only realistic types of sensors and actuators (instead of idealized ones that may not respect all the physical constraints or may have infinite precision) can be used. Similarly, the sensory inputs and the motor outputs should necessarily correspond to physical measures or forces (i.e. they are grounded representation, see Harnad, 1990) and cannot include any abstract information provided by the experimenter, even unconsciously. Finally, only information truly available in the environment can be used for training.

In the following three sections, I will try to show the implications of Evolutionary Robotics for other disciplines. Although I think that Evolutionary Robotics may be relevant for many different fields, I will restrict my analysis to Engineering, Ethology, and Biology. By doing so, it will become clear that the key characteristic of this approach is the possibility to rely largely on a self-organization process (Floreano, 1997).

2. An engineering perspective

That behavioral systems such as mobile robots are difficult to design is undisputed. As one can see in everyday life, there are efficient computer programs that can play chess or solve formal problems but there are no intelligent mobile robots in our homes or towns. The main reason why mobile robots are difficult to design is that their behavior is an emergent property of their motor interaction with the environment. The robot and the environment can be described as a dynamical system because the sensory state of the robot at any given time is a function of both the environment and of the robot previous actions. The fact that behavior is an emergent property of the interaction between the robot and the environment has the nice consequence that simple robots can produce complex behavior (see Braitenberg, 1984). However it also has the consequence that, as in all dynamical systems, the properties of the emergent behavior cannot easily be predicted or inferred from a knowledge of the rules governing the interactions. The reverse it also true: it is difficult to predict which rules will produce a given behavior, since behavior is the emergent result of the dynamical interaction between the robot and the environment.

The main strategy followed in order to overcome these difficulties has been that of *Divide and Conquer*: i.e. divide the problem into a list of hopefully simpler sub-problems (Harvey *et al*, 1997). Classical approaches to robotics have often assumed a primary breakdown into Perception, Planning, and Action. However, this way of dividing the problem has produced limited results and has been criticized by a number of researchers. Brooks (1986) proposed a radically different approach in which the division is accomplished at the level of behavior. The desired behavior is broken down into a set of simpler basic behaviors, which are activated or suppressed through a coordination mechanism. In this latter approach the control system is built up incrementally layer by layer where each layer is responsible for a single basic behavior by directly linking sensors to motors. Simple basic behaviors are implemented first, then new layers implementing other basic behaviors are added one at a time after intensive testing and debugging. This approach has proven to be more successful than the classical approach. Moreover it has been shown how both the layers (modules) responsible for simple basic behavior and the coordination mechanism can be obtained through a self-organizing process rather than by explicit design (see Urzelai *et al,* 1998; Maes, 1992; Mahadevan & Connell, 1992; Dorigo & Schnepf, 1993).

The approaches based on behavioral decomposition, however, still leave the decision of how to break the desired behavior down into simple basic behaviors to the designer. Unfortunately, it is not clear how a desired behavior should be decomposed and it is very difficult to perform such divisions by hand. Even researchers who successfully adopted the behavioral decomposition and integration approach feel that this is a crucial problem. Brooks for example, notes "Conversely, because of the many behaviors present in a behavior-based system, and their individual dynamics of interaction with the world, it is often hard to say that a particular series of actions was produced by a particular behavior. Sometimes many behaviors are operating simultaneously, or are switching rapidly" (Brooks, 1991, pp. 584-585). Colombetti, Dorigo, and Borghi note at the end of their paper: "learning techniques might be extended to other aspects of robot development, like the architecture of the controller. This means that the structure of behavioral modules should emerge from the learning process, instead of being pre-designed" (Colombetti, Dorigo & Borghi, 1996).

To better understand why it is difficult to break down a global behavior into a set of simpler basic behaviors we have to distinguish two ways of describing behavior: a description from the observer's point of view and a description from the robot's point of view (Sharkey & Heemskerk; 1997). The *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 sensorimotor loops. The *proximal description of behavior* is a description from the point of view of the agent's sensorimotor system that describes how the agent reacts in different sensory situations (see Figure 1).

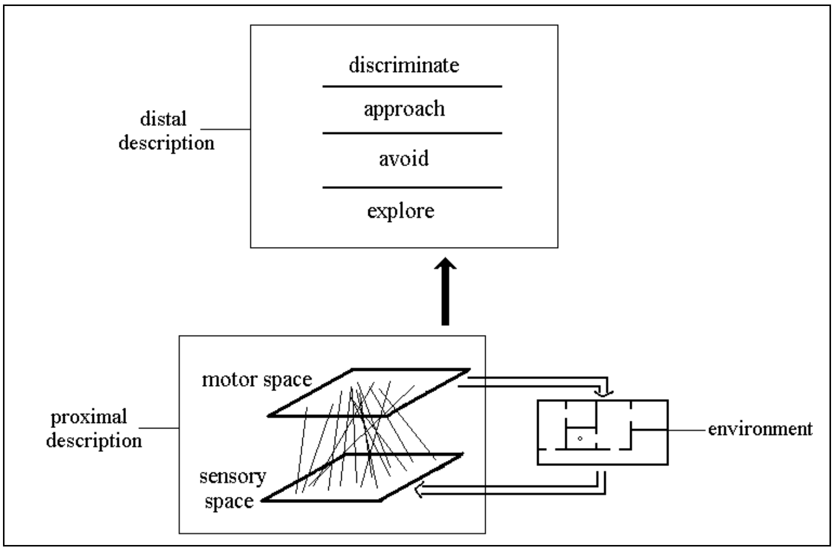

Figure 1. Proximal and distal descriptions of behavior. The distal description is from the observer's point of view and is based on words or sentences in our own language which are used to describe a sequence of sensory-motor loops. The proximal description is from the robot's point of view and is based on a description of how the robot reacts to each sensory state. The empty arrows indicate the reciprocal interactions between the sensory-motor mapping and the environment (the robot actions modify the environment and/or the relation between the robot and the environment which, in turn, modify the type of sensory pattern that the robot receives from the environment). The full arrow indicates that the behavior from the point of view of a distal description (top part of the figure) results from the dynamical interaction between the behavior (from the point of view of the proximal description) and the environment (bottom part of the figure).

It should be noted that behavior from the point of view of a distal description is the result not only of the behavior from a proximal description point of view (i.e. the sensory-motor mapping) but also of the environment. More precisely behavior, from a point of view of distal description, is the result of the dynamical interaction between the agent and the environment. In fact, 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 (we will come back to this point in section 3).

The fact that behavior, from a point of view of a distal description, is the result of the dynamical interaction between the agent and the environment can explain why it is difficult to break down a global behavior into a set of basic behaviors which are simple from the point of view of the proximal description. In general, the breakdown is accomplished intuitively by the researcher on the basis of a description of the global behavior from his point of view (i.e. from the point of view of a distal description). However, because the desired behavior is the emergent result of the dynamical interaction between the agent's control system and the environment it is difficult to predict which type of behavior will be produced by a given control system. The reverse is also true -- it is difficult to predict which control system will produce a desired behavior.

If this is true, if there is not a one-to-one mapping between sub-components of the agent's control system and sub-components of the corresponding behavior (from a point of view of the distal description) it becomes clear that it is questionable whether a distal description of behavior may be used to determine how the agent's control system may be structured as hypothesised in the decomposition and integration approach. In other words, the fact that a global behavior can be divided into a set of simpler basic behaviors from a point of view of the distal description does not imply that such basic behaviors can be implemented into separate and relatively simple layers (modules) of the agent's control system.

Evolutionary Robotics, by relying on an evaluation of the system as a whole and of its global behavior, releases the designer from the burden of deciding how to break the desired behavior down into simple basic behaviors.

It can be shown, for example, that it is impossible to train a controller for a Khepera robot (Mondada, Franzi & Ienne, 1993) placed in a rectangular arena surrounded by walls and containing a target object (a small cylinder) to find and stay close to the target by decomposing the desired behavior into basic behaviors. If we divide this global behavior into four basic behaviors (1 - explore the environment; 2 - avoid walls; 3 - approach and remaining close to the target; 4 - discriminate the target from the wall) and we train 4 separate neural networks to perform each basic behavior we fail[2]. On the contrary, we can easily evolve a single network which is able to produce the 4 different basic behaviors described above if we select individuals for their ability to perform the global behavior[3]. Interestingly, in this latter case, it can be seen how a simple uniform control system (a fully connected perceptron with 8 sensors and 2 outputs controlling the speed of robot's wheels) may be able to produce a global behavior which can be described (from the point of view of a distal description) as structured into the four different basic behaviors described above (Nolfi, 1997c). For an example in which evolved controllers exhibit an emergent task decomposition in a number of different basic behaviors (from the point of view of a distal description) which are organized in a subsumption-architecture-like fashion see Biro & Ziemke (1998).

In another experiment involving a Khepera robot used to clean an arena of objects by carrying them outside the arena with its gripper, simple uniform controllers have been compared with modular controllers (i.e. control systems divided into layers or modules which are used to control the robot in different environmental circumstances). The results have shown how, in this more complex case, the modular systems outperform the non-modular one. However, better performance is obtained only if what each module stands for is the results of a self-organization process and not of a decision of the experimenter. Moreover, it has been shown how, after training, it is not possible to find a simple correlation between module switching and basic behaviors (from the point of view of a distal description), rather, module switching is correlated with low-level sensory-motor mappings (Nolfi, 1997a).

3. An ethological perspective

[2] The authors showed how neural networks trained by back-propagation to discriminate between targets and walls on the basis of the state of the infrared sensors failed in most of the cases regardless of the type of architecture used.

[3] The comparison has been made between control systems relying on the same type of sensory information (the state of the 8 infrared sensors of the robot). However, as shown by (Scheier and Pfeifer, 1995), the task can be solved by using a control systems organized into separate layers (modules) if additional sensory information (the current and the previous state of the motors) is provided.

Biological molecules, cells, organs, organisms, and ecosystems are all entities at different levels which are potentially relevant for the study of behavior. As a consequence the study of behavior is being conducted within different disciplines. In particular two groups of disciplines can be identified which are separated by a total heterogeneity of the concepts they use to describe, analyze, and explain their object of study (molecular biology, cellular biology, developmental biology, genetics, and neuroscience, on one side; psychology, ecology, and evolutionary biology on the other side) (see Parisi, 1997). On one side neurosciences, for example, use concepts that are appropriate for describing a physical object (i.e. the nervous system). On the other side psychology uses concepts which do not normally refer to physical structures or processes. The concepts used by psychologists to talk about behavior and cognition are derived from philosophy and from the concepts we use in everyday life to describe, predict, and explain our own behavior and the behavior of other living creatures.

Given the existence of two different vocabularies, it is only possible to look a posteriori for correlations between physical and anatomo-physiological phenomena on one side and psychological phenomena on the other. This is the task of such disciplines such as psychophysiology and neuropsychology which serve as bridges. However, it is impossible to trace back the observations concerning the nervous system and the observations concerning behavior to the structure and way of functioning of a single entity that can be described and analyzed using a single theoretical vocabulary (Parisi, 1997).

In the last decade, new research fields which try to overcome this epistemological gap have been developed: Connectionism (Rumelhart & McClelland, 1986) and Embodied Cognition (Brooks, 1991; Varela *et al*., 1991; Pfeifer & Scheier, in press). Connectionism proposes neural networks as a unified theoretical framework for studying both behavior and cognition, on one side, and the nervous system, on the other[4]. Therefore, connectionism can be viewed as an attempt at overcoming the traditional dualism embraced by psychology (Parisi, 1997). Embodiment stresses the importance of the physical aspects of a system (physical shape, gravity, friction, inertia, idiosyncratic characteristics of each sensor and actuator etc.). Moreover, it stresses the importance of the interaction with the environment (Arbib, 1989; Meyer & Guillot, 1991; Wilson, 1991). Embodied agents are not just passively exposed to a continuously changing stream of sensory stimulation. By exploiting their interaction with the environment (i.e. by sensory-motor coordination), they can partially modify the stimuli they receive from the environment. From the point of view of Embodied Cognition, behavior cannot be considered a product of an agent isolated from the world, but can only emerge from a strong coupling of the agent with its environment. In other words, a given behavior cannot be explained on the basis of internal mechanisms only. Similarly, the structure of a nervous system producing a given behavior cannot be considered in isolation from the environment. The function of an agent's nervous system, in fact, is that of coordinating perception and action in order to generate adaptive behavior (Cliff, 1991; Chiel & Beer, 1997).

In most robotics research, however, the power of the interaction with the environment is largely unexplored. Few notable exceptions should be mentioned.

[4] This does not imply that connectionism is the only possible way for studying behavior and cognition with an uniform conceptual framework.

Braitenberg's vehicles are probably the first clear demonstration that a simple agent can produce complex behavior by exploiting the complexity of the environment (Braitenberg, 1984). Pfeifer and Scheier showed how the problem of perception could be greatly simplified if the agent's own movement and interaction with the environment is taken into account (Scheier & Pfeifer, 1995). They described how a Khepera robot, by relying on three simple pre-designed basic behaviors (i.e. move-forward, avoid-object, and turn-toward object) can learn to discriminate between small and large cylinders. This is possible because sensory-motor coordination makes the agent circle around objects which in turn significantly affect the type of sensory information which the robot receives from the environment.

The reason why the power of the interaction with the environment is still largely unexplored in robotics is that, as I pointed out in previous sections, adaptive behavior is difficult to obtain through design. In particular, it is hard to design systems which exploit sensory-motor coordination. For agents which interact with an external environment, in fact, each motor action has two different effects: (a) it partially determines how well the agent perform with respect to a given task; (b) it partially determines the next sensory patterns the agent will receive from the environment that in turn may determine if the agent will be able to solve its task or not. The problem is that determining which motor action the agent should perform each time by taking into account both (a) and (b) is very difficult given that: each motor action can have long term consequences; the effect of each motor action is a function also of the preceding and successive actions; the effect of each action is a function of the interaction between the agent and the environment.

The Evolutionary Robotics approach, by largely relying on self-organization, is not affected by these problems and, as a consequence, is an ideal framework for studying adaptive behavior[5]. Indeed, many of the evolved robots exploit an active interaction with the environment in order to maximize the selection criterion. For example, in trying to evolve the control system for a Khepera robot which is asked to discriminate between walls and cylinders, it was found that all evolved individuals solved the task by moving back and forth in front of the perceived objects (Nolfi, 1997c). Similarly, in trying to evolve robots able to reach the upper right hand or bottom left hand corner of a rectangular box starting from eight different positions (an experimental task studied by Gallistel with rats, see Gallistel, 1990), it was found that evolved robots were able to solve the task by carefully selecting the speed of the two wheels in the absence of stimulation. This simple strategy ensured that, given the shape and the dimension of the environment, long and short walls were encountered with significantly different angles. This in turn allowed the robots to easily reach the two target corners by following or avoiding the walls depending on the angle with which they were approached (Nolfi & Miglino, in press). Another clear example can be taken from Harvey, Husband, and Cliff (1994) who evolved the control system for a specially designed robot which was trained to visually discriminate between a triangular and a rectangular target. By looking at the architecture of one successful evolved individual (the architecture was subjected to the evolutionary process), they found that it was able to solve the task by relying only on the

[5] This does not imply that evolutionary robotics is the only possible methodology for studying adaptive behavior. Pfeifer and Scheier (1997), for example, proposed a list of design principles which should help researchers both to design artificial agents and to understand natural organisms. Another possibility, of course, is to use evidence from neuro-physiology when it is available at a sufficient level of detail (see, for example, Franceschini, 1997). We will return to this issue in the last section.

sensory information coming from two pixels of the camera. Clearly the two targets cannot be discriminated by using only two pixels without relying on sensory-motor coordination.

4. A biological perspective

Evolutionary Robotics and Biology share an interest in the following question: what are the key characteristics of natural evolution that make it so successful in producing the extraordinary variety of highly adapted life forms present on the planet? Producing better answers to this question may significantly increase both our understanding of biological systems and our ability to design artificial systems. From the point of view of Evolutionary Robotics this question may be re-stated as follows: in which conditions is an artificial evolutionary process likely to select individuals which develop complex competencies in order to adapt to their artificial environment? Possible answers to this question may be categorised into three main issues which will be described in the next three sections. As we will see it is possible, at least in principle, to develop complex forms of behavior without increasing the amount of supervision. This may be accomplished: (a) by generating incremental evolution through competitions between or within species; (b) by leaving the system free to decide how to extract supervision from the environment; (c) by including the genotype-to-phenotype mapping within the evolutionary process.

4.1 Incremental Evolution

From the point of view of Evolutionary Robotics, a key question is how Artificial Evolution can select individuals which have competencies that allow them to solve complex tasks (e.g. navigating in a complex environment). A related problem in Biology is to understand how and in which circumstances natural evolution may discover new competencies (e.g. the ability to fly or build a nest). To avoid confusion, I should clarify here that competencies and selection criteria are different entities. In natural evolution selection is based on a simple and general criterion: the ability to reproduce. In spite of this, natural evolution has been able to produce individuals with sophisticated competencies such as the ability to swim, fly, and communicate through natural language.

If one wishes to select individuals able to solve a task that requires a specific competence through Artificial Evolution, the easiest thing to do is to select the individuals for their ability to solve that task (i.e. to design a fitness criterion that scores individuals according to their ability to solve that task). However, it is easy to show that this simple strategy can only work for simple tasks. As the complexity of the task increases, the probability that some individuals of the first generations are able to accomplish, at least in part, the task is inversely proportional to the complexity of the task itself. For complex tasks, all individuals of the first generations are scored with the same null value, and as a consequence the selection process cannot operate (Nolfi, 1997a). I will refer to this problem as the *bootstrap problem*.

One possible solution to this problem is to increase the amount of supervision. The insights of the experimenter can be used to include in the selection criterion also rewards for sub-parts of the desired task (see Nolfi, 1997a). Another possibility is to start the evolutionary process with a simplified version of the task and then progressively increase its complexity by modifying the selection criterion (the latter technique is usually referred to as *incremental evolution*, see Harvey *et al*., 1997). This idea was previously proposed in the context of classifier systems by Dorigo & Colombetti (1994, 1997) who named it *shaping,* borrowing the term from the experimental psychology techniques used to train animals to produce predefined responses (Skinner, 1938). These techniques, by requiring more supervision, increase the risk of introducing inappropriate constraints (see section 2). However, from an engineering perspective, it is easier to use the insight of the experimenter for shaping the selection criterion than for designing the control system itself. Introducing more constraints in the selection criterion still leaves the evolving robots free to select their way to adapt to the environment, although the additional constraints may channel the evolutionary process in wrong directions.

A more desirable solution to the bootstrap problem, however, would be a self-organized process capable of producing incremental evolution that does not require any human supervision. This ideal situation spontaneously arises in co-evolving populations (i.e. in the case of competing populations with coupled fitness such as predator and prey). Each co-evolving population may progressively produce more complex challenges for the other population. As a consequence, as discussed by Dawkins and Krebs (1979), competing populations may reciprocally drive one another to increasing levels of complexity by producing an evolutionary “arms race”.

For these reasons, co-evolution is actively studied in Evolutionary Robotics. Nolfi & Floreano (in press), for example, have investigated the conditions in which co-evolution leads to progressively more complex behavior in co-evolving robots. In particular, they have been able to show that in some cases co-evolution can solve tasks that evolution of a single population cannot. The authors considered the case of two competing species of predator and prey robots which were respectively selected for their ability to catch prey and escape predators. At the beginning of the evolutionary process, predators should be able to catch prey which have a very simple behavior and are therefore easy to catch; likewise, prey should be able to escape simple predators. However, later on, both populations and their evolving challenges will become progressively more and more complex. Therefore, even if the selection criterion remains the same, the adaptation task becomes progressively more complex. In certain conditions this may lead to the selection of more and more complex forms of behavior for both species. To verify if co-evolution can really produce solutions to problems that evolution alone is unable to solve the authors tried to evolve a single population of predators able to catch the best prey obtained using artificial co-evolution. These experiments failed because of the bootstrap problem mentioned above (i.e. in most of the cases performance did not increase at all) showing that co-evolution may produce solutions to problems that evolution of a single population cannot solve.

4.2 Extracting supervision from the environment through lifetime learning

From the point of view of a natural or artificial organism the external environment does not provide any direct cue on how the agent should act to attain a given goal. However agents receive a large amount of information from the environment through the sensors. Such information (which is a function of both of the environmental structure and of the motor actions of the agent) may be used not only to determine how to react in different environmental circumstances but also to adapt to the current environment through lifetime learning. For example, a robot may learn the consequences of different actions in different environmental contexts or it may learn to classify sensory states not only on the basis of the current perceived sensory pattern but also on the basis of the preceding and following sensory patterns.

In principle, in an evolving population, any ability which can be acquired through lifetime learning can also be genetically acquired through evolution. However these two ways of adapting to the environment differ in one important respect: ontogenetic adaptation can rely on a very rich, although less explicit, amount of supervision. From the point of view of phylogenetic adaptation, individuals are evaluated only once on the basis of a single value which codifies how well they were adapted to their environment throughout all their lifetime (i.e. the number of offspring in the case of natural evolution and the fitness value in the case of Artificial Evolution). Instead, from the point of view of ontogenetic adaptation, individuals receive information from the environment through their sensors throughout all their lifetime. However, such huge amounts of information encode only very indirectly how well they did in different moments of their lifetime or how they should modify their behavior in order to increase their fitness. The problem is how such information can be transformed in an indication of what the agent should to do or how well it is doing.

Because of the same problems I discussed in Sections 2 and 3 it is probably hard to design a system capable or performing a good transformation. On the other hand, we may expect that evolution can be able to solve this type of problem by producing better and better transformations throughout generations. That this is possible has been proven in two experimental studies made by: Ackley & Littman (1991) and Nolfi & Parisi (1997). In both cases a fixed architecture divided into two sub-networks (of which the former has the function of determining how to react to the current sensory state and a latter has the function of generating a teaching signal for the first) was used. In Ackley and Littman sensory states were transformed into reinforcement signals while in Nolfi and Parisi sensory states were transformed into self-generated teaching inputs. By subjecting the weights of the two sub-networks to an evolutionary process, the authors showed how individuals which learn during their lifetime to adapt to their environment through self-generated teaching signals evolve and how evolved individuals are able to transform the sensory information into useful reinforcement signal or teaching inputs. A similar result can be obtained by evolving neural networks with a topology which may vary evolutionarily and which learn throughout lifetime by using unsupervised learning (Miller & Todd, 1990) or an architecture in which evolution selects different learning rules for different connections (Floreano and Mondada, 1996). In these cases the constraints on the architecture, which are evolved, channel the changes driven by the sensory states in the right directions.

As I said above, what can be obtained with evolution and learning can also be obtained with evolution alone. At a high level of description, for example, an individual that is born with a general strategy which is able to produce a behavior which is effective in a set of different environments is equivalent to another individual that is able to adapt to each environment through lifetime learning. On the other hand, at a lower level of description it is clear that these two individuals are organized in different ways. The individuals that do not start with a general strategy but adapt throughout lifetime should be able to detect the environment in which they are located, and should be able to modify their strategy accordingly (let us call this individual ‘plastic-generals’). On the contrary, the individuals which start with a general strategy which is already suitable to different environments do not need to do so (let us call these individual ‘full-generals’). From this point of view the full-generals will be more effective than plastic-general individuals given that they do not need to adapt throughout lifetime and may provide the appropriate motor answers to the current environmental situation from the beginning. On the other hand, it may be that in certain conditions a full-general individual cannot be selected because a full-general strategy does not exist (or because it is too complex and therefore the probability that it will be selected is low). If this is the case, the plastic-general solution is the only available option (Nolfi & Floreano, in press).

Preliminary evidence that this may be the case even in relatively simple environments can be found in Nolfi & Floreano (in press). In the experiments with predator and prey robots which I already mentioned in the previous section, they tried to co-evolve predator and prey robots by testing them against all the best competitors of previous generations. By doing so the authors observed that, in most of the cases, it was impossible to obtain predators or prey that were able to defeat a large number of competitors although it was always possible to easily find a set of different individuals with different strategies able to do so. In other words it appears that in this case, a full-general strategy does not exist or is too difficult to find while a collection of simple strategies appropriate in different circumstances may be easily found[6].

It should be made clear that, up to now, there has been no clear experimental evidence within Evolutionary Robotics that evolution and learning can evolve more complex levels of competence than evolution alone. However, the study of learning in an evolutionary perspective, is still in its infancy. I believe that the study of learning in interaction with evolution will produce in the next years an enormous impact on our understanding of what learning is.

4.3 Development and the evolution of evolvability

One of the aspects of Natural Evolution which is more crudely simplified in Evolutionary Robotics is development. This can explain why recently the importance of development is one of the more debated issues in Evolutionary Robotics. Interestingly, the importance and the role of development in Natural Evolution is a controversial issue in Evolutionary Biology too.

From an Evolutionary Robotics perspective, this issue is usually referred to as the ‘genotype-to-phenotype’ mapping problem. As claimed by Wagner and Altenberg in fact, “For adaptation to occur, these systems must possess *evolvability*, i.e. the ability of random variations to sometimes produce improvement. It was found that evolvability critically depends on the way genetic variation maps onto phenotypic variation, an issue known as the representation problem” (Wagner & Altenberg, 1996).

The simplest ‘genotype-to-phenotype’ mapping is the one-to-one mapping in which each gene codifies for a single character of the robot. However, in most of the experiments conducted in Evolutionary Robotics, the mapping is more complex than

[6] Preliminary evidence that evolving individuals that are also allowed to adapt through lifetime learning may be able to adapt their strategy to the current competitor are described in (Floreano and Nolfi, 1997).

that. It may involve: (a) several levels of organization (e.g. genotype, nervous system, and behavior) which are hierarchically organized and involve non-linear interactions between different levels (i.e. one character at one level is usually responsible for several characters at higher levels; see for example Floreano & Mondada, 1994); (b) growing instructions that are recursive in the sense that they are applied to their own results in a way that resembles the process of cell duplication and differentiation (see Kodjabachian & Meyer; this special issue); (c) plasticity (i.e. the likelihood of being influenced by the external environment (see the previous Section)); (d) genotypes which vary in length (Harvey, 1992).

The opportunity to study these features of development in isolation, to manipulate the way they are modelled, to generate huge amounts of data easily, may allow Evolutionary Robotics to help to understand the importance and the role of development from an evolutionary perspective.

Another important difference between Natural Evolution and Artificial Evolution is that in the former case the mapping itself is subjected to the evolutionary process while in the latter the mapping is generally designed by the experimenter (although it is designed by taking inspiration from how it is accomplished in natural organisms). Unfortunately, also in this case, it is difficult to design a good mapping. Similarly, it is difficult to shape the mapping just by imitating nature given that not everything is known and, for practical reasons, only some aspects of the natural process can be modelled.

All these problems and the fact that it is still not clear how the mapping itself may be subjected to the evolutionary process may explain why only limited results have been obtained in trying to show how more realistic mappings may allow the development of more complex behavior.

5. Discussion

In this paper I claimed that one of the main characteristics that makes the Evolutionary Robotics approach suitable for the study of adaptive behavior in natural agents and for the synthesis of adaptive behavior in artificial agents is the possibility to rely largely on a self-organization process. In principle, in fact, the role of the designer may be limited to the specification of a selection criterion. However, we saw how in real experiments the role of the designer is much greater than that: (a) in most of the cases the genotype-to-phenotype mapping is designed by the experimenter; (b) several parameters (e.g. the number of individuals in the population, the mutation and crossover rate, the length of the lifetime of each individual etc.) are determined by the experimenter; (c) in some cases the architecture of the controller is also handcrafted. In theory, all these parameters may be subjected to the evolutionary process; however in practice they are not. Should we conclude that the amount of design involved in these experiments is similar to that required in other approaches?

I do not think that this is the case given that most of these parameters are designed independently from the task. This may be more clearly shown by one example. By only varying the selection criterion, all other conditions being the same (i.e. number of individuals, mutation rate, lifetime length, architecture of control system, genotype-to-phenotype mapping, coding of the sensory-motor information, robot body) quite different forms of behavior can be obtained: (a) obstacle avoidance (Floreano & Mondada, 1994); (b) exploration (Lund & Hallam, 1996); (c) navigation (Nolfi &

Miglino, in press); (d) discrimination of different objects (Nolfi, 1997c); (e) escaping predators (Nolfi & Floreano, in press).

Of course, some of the characteristics which could be included in the evolutionary process but are often set by the experimenter (in particular the shape of the body, the structure of the sensory-motor system, and the architecture of the controller) can strongly affect the obtained results. On the other hand, in other approaches, also other characteristics should be handcrafted.

Another question we may want to raise is if Evolutionary Robotics is the best experimental tool for studying adaptive behavior. I do not think that this can be claimed. Evolutionary Robotics has also a number of drawbacks, the most serious being the time needed to conduct the evolutionary process on the real robot (see Mataric & Cliff, 1996). These drawbacks may prevent this approach from scaling up to the study of more complex forms of behavior. Moreover different approaches, by having different constraints and different characteristics, may be more effective in different circumstances.

One interesting approach, for example, is based on the identification of a set of designed principles which should guide the researcher in designing behavioral systems and, in doing so, to understand natural systems (Pfeifer & Scheier 1998, in press). The authors point out that models and experimental tools often rely on the implicit intuition of the designer and on her or his view of intelligence. The purpose of the design principles is to make this knowledge explicit and to take inspiration from systematic studies of biological systems instead of from implicit assumptions. Notice that in Pfeifer and Scheier's view, design principles are heuristics for the design and the study of autonomous agents and not detailed models. This general level of description leaves room for emergence. As authors point out, "Emergence is required because behavior cannot be reduced to an internal mechanism only…. For example, if we want to achieve wall-following behavior we should not design a module for wall-following within the agent, but we should define basic processes which together, in the interaction with the environment, lead to this desired behavior. This is called design for emergence" (Pfeifer & Scheier, 1998, p.15). Evolutionary Robotics fully shares this view of intelligence. Indeed design principles may provide a theoretical background for Evolutionary Robotics and help the researcher to understand the results of the evolutionary process. Understanding how evolved agents solve their task is difficult given that: (1) behavior is an emergent property of the interaction between the agent and the environment; (2) the way in which evolved agents solve their tasks is usually very different from the way in which we would design systems which should solve the same tasks. Design principles, by providing new ways to describe behavioral systems, may solve the second of these two problems. They can serve as hypotheses which might help us to understand. On the other hand, our understanding of emergent behavior in evolving agents (natural and artificial) may help us to discover new design principles and/or discard existing ones. Finally, design principles may help us to design the characteristics which are not subjected to the evolutionary process (e.g. an hypothetical design principles might be that although both the weights and the architecture of evolving agents should be subjected to the evolutionary process, the genotype-to-phenotype mapping should ensure that the latter changes less frequently than the former).

Last but not least, behavioral systems can be successfully understood and developed by collecting neuro-physiological and behavioral data on natural organisms. When this is possible at a sufficient level of detail, this is certainly the best way to proceed. One clear example is the work of Franceschini on the ability of the fly to navigate in a complex environment. By carefully studying the vision system of the fly he came up with a model so detailed that it was easy for him to implement it in a mobile robot able to navigate in an environment by visuo-motor coordination (Franceschini, 1997). For another example in the case of rat navigation see Burgess, Donnett, & O'Keefe (1998). In cases like this, there is no need of self-organization processes at all. We just need to understand what emerged from natural evolution. Unfortunately however, this deep comprehension of both behavioral and neuro-physiological data has been (and maybe can be) reached only in some cases.

One last point which we should discuss is the relevance of Evolutionary Robotics experiments for the understanding of biological systems. As most models of biological systems, evolutionary robotic models are extremely simplified with respect to the natural processes they are supposed to capture. Moreover, by being models based on self-organization they cannot model, even in principle, all the details of a single natural system. If we could re-start the natural evolution process, we would probably get something very different from what we have. Similarly, even if we could replicate roughly the initial conditions of an evolutionary phase, we might obtain a very different outcome. From this we might conclude that Evolutionary Robotics may be useful to understand general principles about natural evolution and adaptive behavior or to test theoretical hypothesis (e.g. if all relevant aspects of the phenomena are included in the experiment). However, when we move from general issues (e.g. the role of co-evolution between competing species in the emergence of complex behavior, or the effect of the interaction between learning and evolution) to more specific issues (e.g. how rats navigate in a certain environment) it becomes more and more questionable if the results obtained are of any relevance for explaining the object of study.

In the latter case (for example, in the study of an individual species) the solutions found by using Artificial Evolution can be used at most as counter-examples. In modeling a given form of behavior biologists or ethologists often make implicit assumptions. For example, they may assume that an observed behavior performed by an individual organism necessarily implies a certain type of strategy. By evolving an artificial agent for the ability to perform the same behavior one can show that different types of solutions exist (for an example obtained by comparing rats and evolved robots performing a navigation task see Nolfi & Miglino, in press). Although this does not prove that the strategy selected through Artificial Evolution is that adopted by the natural organism, it is useful to discover that other hypotheses are viable and that the assumptions which were made were wrong.

Acknowledgement

I thank Dario Floreano, Christian Scheier, and Tom Ziemke for useful comments on the first draft of the paper.

References

Ackley, D. H. & Littman, M.L. (1991) Interactions between learning and evolution, in: C. G. Langton, J. D. Farmer, S. Rasmussen, C. E. Taylor (eds.), *Artificial Life II*. Reading, MA: Addison-Wesley.

Arbib, M.A. (1989) *The Methaphorical Brain 2*. New York: Wiley.
Barto, A. G., Bradtke, S.J. & Singh, S.P. (1995) Learning to act using real-time dynamic programming. *Artificial Intelligence*, 6, 105-122.
Biro, Z. & Ziemke, T. (in press). Evolution of visually-guided approach behaviour in recurrent artificial neural network robot controllers. In *Proceedings of the Fifth International Conference of the Society for Adaptive Behavior (SAB98)*, University of Zürich (Switzerland), 17 - 21 August 1998, MIT Press
Booker, L., Godberg, D.E., & Holland, J.H. (1989) Classifier systems and genetic algorithms, *Artificial Intelligence*, 40, 235-282.
Braitenberg, V. (1984) *Vehicles*. Cambridge, MA: MIT Press.
Brooks, R.A. (1986) A robust layered control system for a mobile robot, *IEEE Journal of Robotics and Autonomation,* 2, 14-23.
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.
Brooks, R.A. (1992) Artificial life and real robots, In F. J. Varela & P. Bourgine (Eds.), *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life.* Cambridge, MA: MIT Press/Bradford Books.
Burgess, N., Donnett, J.G. & O'Keefe, J. (1998). Using a mobile robot to test a model of the rat hippocampus. *Connection Science*, 10 (3-4), special issue on BioRobotics.
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.
Cliff, D.T. (1991) Computational neuroethology: A provisional manifesto, In J-A. Meyer & S.W. Wilson (eds), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. Cambridge MA: MIT Press/Bradford Books.
Cliff, D.T., Harvey, I. & Husbands, P. (1993) Explorations in Evolutionary Robotics. *Adaptive Behavior,* 2, 73-110.
Cliff, D.T. & Miller, G.F. (1996) Co-evolution of pursuit and evasion II: Simulation Methods and results, In P. Maes, M. Mataric, J-A. Meyer, J. Pollack, H. Roitblat & S. Wilson (Eds.), *From Animals to Animats IV: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press-Bradford Books
Colombetti, M., Dorigo, M. & Borghi, G. (1996) Behavior analysis and training: A methodology for behavior engineering. *IEEE Transactions on Systems, Man, and Cybernetics* - Part B, 3, 365-380.
Dawkins, R., & Krebs, J.R. (1979) Arms races between and within species. *Proceedings of the Royal Society of London B*, 205, 489-511.
Dorigo, M. & Schnepf, U. (1993) Genetic-based machine learning and behavior based robotics: A new synthesis. *IEEE Transaction on Systems, Man, and Cybernetics*, 1, 141-154.
Dorigo, M. & Colombetti, M. (1994) Robot shaping: Developing autonomous agents through learning, *Artificial Intelligence*, 71, 321-370.
Dorigo, M. & Colombetti, M. (1997) *Robot Shaping: An Experiment in Behavior Engineering*. Cambridge, MA: MIT Press/Bradford Books.
Floreano, D. (1997)Reducing human design and increasing adaptability in Evolutionary Robotics, In T. Gomi (Ed.), *Evolutionary Robotics*: *From Intelligent Robots to Artificial Life*. Kanata, Canada: AAI Books.
Floreano, D. & Mondada, F. (1994) Automatic Creation of an Autonomous Agent: Genetic Evolution of a Neural-Network Driven Robot, In D. Cliff, P. Husbands, J. Meyer & S. W. Wilson (Eds.), *From Animals to Animats 3: Proceedings of Third Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press/Bradford Books.

Floreano, D. & Mondada, F. (1996) Evolution of plastic neurocontrollers for situated agents, In P. Maes, M. Mataric, J-A. Meyer, J. Pollack & S. Wilson. (Eds.), *From Animals to Animats IV*. Cambridge, MA: MIT Press.

Franceschini, N. (1997) Combined optical, neuroanatomical, electrophysiological and behavioral studies on signal processing in the fly compound eye, In C. Taddei-Ferretti (Ed.), *Biocybernetics of Vision: Integrative Mechanisms and Cognitive Processes*. London: World Scientific.

Gallistel, C. R. (1990) *The Organization of Learning*. Cambridge, MA: MIT Press.

Harnad, S. (1990) Symbol grounding problem, *Physica D*, 42, 335-464.

Harvey, I. (1992) The SAGA cross: The mechanics of recombination for species with variable length genotypes, In R. Manner & B. Manderick, (Eds.), *Parallel Problem Solving from Nature*. North-Holland.

Harvey, I., Husband, I. & Cliff, D. (1994) Seeing the light: Artificial evolution, real vision, In D. Cliff, P. Husband, J-A Meyer & S. Wilson (Eds.), *From Animals to Animats 3*, Proceedings of the Third International Conference on Simulation of Adaptive Behavior. Cambridge, MA: MIT Press/Bradford Books.

Harvey, I., Husbands, P. & Cliff, D., Thompson A., Jakobi N. (1997) Evolutionary Robotics: The Sussex approach, *Robotics and Autonomous Systems*, 20, 205-224.

Holland, J. H. (1975) *Adaptation in Natural and Artificial Systems*. Ann Arbor, Michigan: University of Michigan Press.

Kodjabachian, J. & Meyer, J.A. (1998) Evolution and development of modular control architectures for 1-D locomotion in six-legged animats, *Connection Science*, 10 (3-4), special issue on BioRobotics.

Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43, 59-69.

Lund, H. H. & Hallam, J. (1996) Sufficient neurocontrollers can be surprisingly simple. *Technical Report n. 824*. Department of Artificial Intelligence, University of Edinburgh.

Lund, H. H., Hallam, J. & Lee, W.P. (1997) Evolving robot morphology, In *Proceedings of IEEE 4th International Conference on Evolutionary Computation*. IEEE Press.

Maes, P. (1992) Learning behavior networks from experience, In F. J. Varela, P. Bourgine (Eds.), *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*. Cambridge, MA: MIT Press/Bradford Books.

Mahadevan, S. & Connell, J. (1992) Automatic programming of behavior-based robots using reinforcement learning. *Artificial Intelligence*, 55, 311-365.

Mataric, M. J. & Cliff, D. (1996) Challenges in evolving controllers for physical robots, *Robotics and Autonomous Systems*, 1, 67-83.

Miller, G.F. & Todd, P.M. (1990) Exploring adaptive agency I: Theory and methods for simulating the evolution of learning. In D.S. Touretzky, J.L. Elman, T.J. Sejnowski & G.E. Hinton (Eds.), *Proceedings of the 1990 Connectionist Models Summer School*. San Matteo, CA: Morgan Kaufmann

Meyer, J-A. & Guillot, A. (1991) Simulation of adaptive behavior in animats: Review and prospect, In J-A. Meyer & S.W. Wilson (Eds), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. Cambridge MA: MIT Press/Bradford Books.

Mondada, R., Franzi, E. & Ienne, P. (1993). Mobile robot miniaturization: A tool for investigation in control algorithms, In *Proceedings of the Third International Symposium on Experimental Robots*, Kyoto, Japan.

Nolfi, S. (1997a) Using emergent modularity to develop control system for mobile robots, *Adaptive Behavior*, Special issue on Environment Structure and Behavior, 3-4, 343-364.

Nolfi, S. (1997b) Evolving non-trivial behaviors on real robots: A garbage collecting robot, *Robotics and Autonomous System*, 22, 187-198

Nolfi, S. (1997c) Evolving non-trivial behavior on autonomous robots: Adaptation is more powerful than decomposition and integration, In T. Gomi (Ed.), *Evolutionary Robotics: From Intelligent Robots to Artificial Life*. Kanata, Canada: AAI Books.

Nolfi, S. & Floreano, D. (in press) Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution?, *Artificial Life*

Nolfi, S., Floreano, D., Miglino, O. & Mondada F. (1994) How to evolve autonomous robots: different approaches in evolutionary robotics, In R.A. Brooks & P. Maes (Eds.), *Proceedings of the Fourth International Conference on Artificial Life.* Cambridge, MA: MIT Press.

Nolfi, S. & Miglino, O. (in press) Studying the emergence of grounded representations: Exploiting the power and the limits of sensory-motor coordination, In G. Mulhauser (Ed.), *Evolving Consciousness*, Advances in Consciousness Research Series. Amsterdam : John Benjamins Press.

Nolfi, S. & Parisi, D. (1997) Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 1, 99-105

Parisi, D. (1997) Artificial life and higher level cognition. *Brain and Cognition,* 34, 160-184.

Pfeifer, R. & Scheier, C. (1997) Sensory-motor coordination: The metaphor and beyond. In R. Pfeifer, and R. Brooks (Eds.), *Robotics and Autonomous Systems*,

Pfeifer, R. & Scheier, C. (1998) Embodied cognitive science: A novel approach to the study of intelligence in natural and artificial systems, In T. Gomi (Ed.), *Evolutionary Robotics: From Intelligent Robots to Artificial Life*. Ontario, Canada: Applied AI Books.

Pfeifer, R. & Scheier, C. (in press) *Understanding Intelligence*. Cambridge, MA:MIT Press

Rumelhart, D.E. & McClelland, J.L. (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, MA:MIT Press

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986) Learning representation by Back-Propagation of errors. *Nature*, 323, 533-536.

Sharkey, N.E. & Heemskerk, J. (1997) The neural mind and the robot, In A. J. Browne (Ed.), *Neural Network Perspectives on Cognition and Adaptive Robotics,* Bristol U.K. :IOP press.

Scheier, C. & Pfeifer, R. (1995) Classification as sensorimotor coordination: A case study on autonomous agents, In F. Moran, A. Moreno, J.J. Merelo & P. Chacon (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*. Springer Verlag.

Skinner, B.F. (1938) *The Behavior of Organisms: An Experimental Analysis*, D. New York: Appleton Century.

Urzelai, J., Floreano, D., Dorigo, M. & Colombetti, M. (1998) Incremental robot shaping, *Connection Science*, 10 (3-4), special issue on BioRobotics.

Varela, F. J., Thompson, E. & Rosch E. (1991) *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, MA: MIT Press.

Wagner, G.P. & Altenberg, L. (1996) Complex adaptations and the evolution of evolvability, *Evolution,* 50, 967-976.

Wilson, S.W. (1991) The animat path to AI, In J-A. Meyer & S.W. Wilson (Eds), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. Cambridge MA: MIT Press/Bradford Books.