

1 The role of self-organization for the synthesis and the understanding of behavioral systems

1.1 Introduction

The basic idea behind evolutionary robotics goes as follows (see figure 1.1). An initial population of different artificial chromosomes, each encoding the control system (and sometimes the morphology) of a robot, are randomly created and put in the environment. Each robot (physical or simulated) is then let free to act (move, look around, manipulate) according to a genetically specified controller while its performance on various tasks is automatically evaluated. The fittest robots are allowed to reproduce (sexually or asexually) by generating copies of their genotypes with the addition of changes introduced by some genetic operators (e.g., mutations, crossover, duplication). This process is repeated for a number of generations until an individual is born which satisfies the performance criterion (*fitness function*) set by the experimenter.

Evolutionary robotics shares many of characteristics with other approaches, such as behavior-based robotics, robot learning, and artificial life.

Behavior-Based Robotics

The *behavior-based robotics* approach is based upon the idea of providing the robot with a collection of simple basic behaviors. The global behavior of the robot emerges through the interaction between those basic behaviors and the environment in which the robot finds itself (Brooks 1986, 1999; Arkin 1998). Basic behaviors are implemented in separate sub-parts of the control system and a coordination mechanism is responsible for determining the relative strength of each behavior in a particular moment. Coordination may be accomplished by means of competitive or cooperative methods. In competitive methods only one behavior affects the motor output of the robot in a particular moment (see, for example, the subsumption based method proposed by Brooks 1986). In cooperative methods different behaviors may contribute to a single motor action although with different strength (see, for example, the method based on behavioral fusion via vector summation [Arkin 1989]).

In this approach, as in evolutionary robotics, the environment plays a central role by determining the role of each basic behavior at any given time. Moreover, these systems are usually designed through a trial and error process in which the designer modifies the current behaviors and progressively increase the number of basic behaviors while testing the resulting global 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,

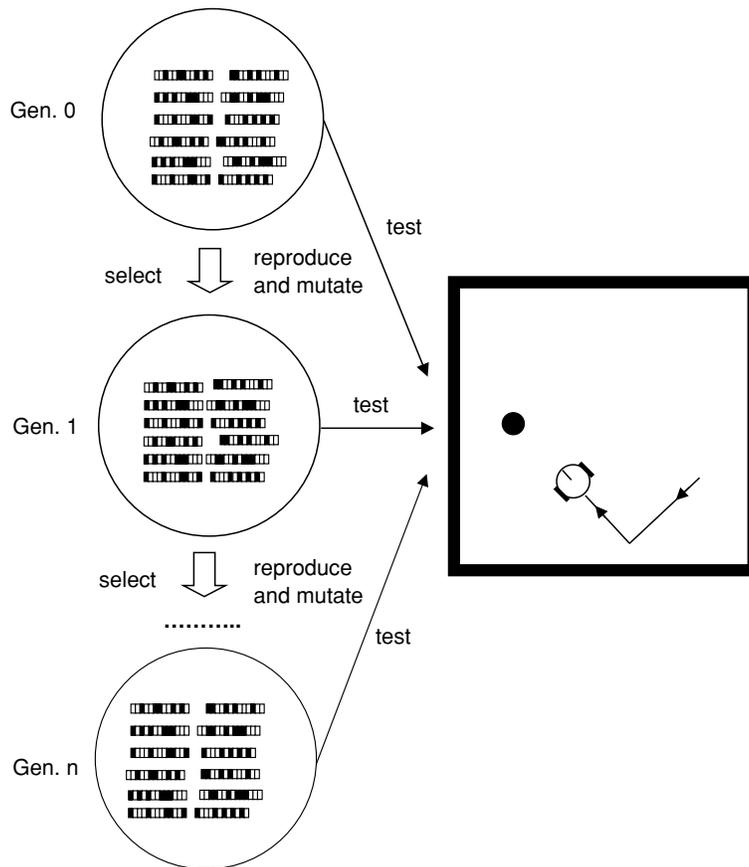


Figure 1.1

Basic Evolutionary Robotics methodology. A population of artificial chromosomes is decoded into corresponding controllers which are tested one at a time in the environment. The fittest controllers are allowed to reproduce and generate the next population of chromosomes.

in evolutionary robotics this is often the result of a self-organizing process (figure 1.2). Indeed, the entire organization of the evolving system, including its organization into subcomponents, is the result of an adaptation process that usually involves a large number of evaluations of the interactions between the system and the environment (we will return to this issue in section 3 and in chapter 5 and 6).

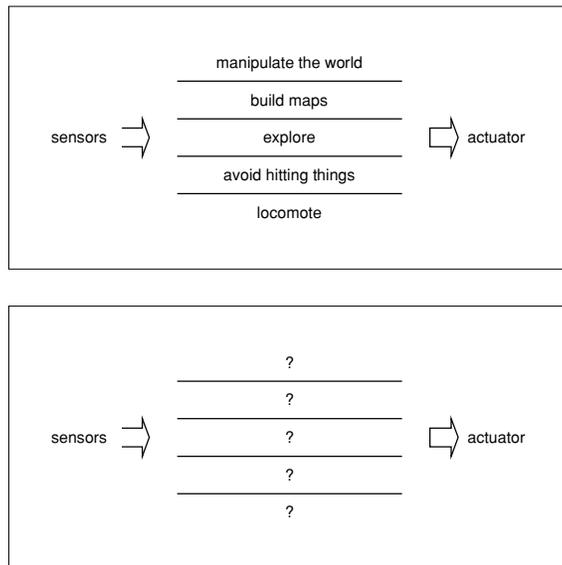


Figure 1.2

In the behavior-based approach the desired behavior is broken down by the designer into a set of basic behaviors which are implemented into separate sub-parts (layers) of the robot's control system (top box). In Evolutionary Robotics the designer does not need to decide how to divide the desired behavior into basic behaviors (bottom box). The way in which a desired behavior is broken down into modules is the result of a self-organization process.

Robot learning

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 generalize 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: back-propagation learning (Rumelhart et al. 1986); reinforcement learning (Barto et al. 1995); classifier systems (Booker et al. 1989); 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

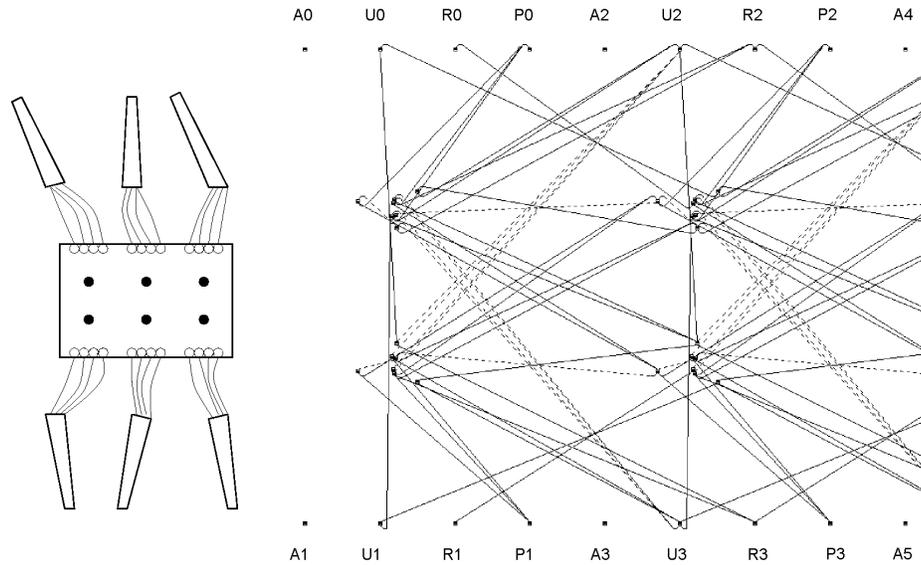


Figure 1.3

Example of evolved architecture for a tripod robot. **Left:** The robot body. Black circles represent original cells that develop into sub-networks connecting sensory and motor neurons (white circles) **Right:** The corresponding developed locomotion network. Labels indicate sensory and motor neurons corresponding to the white circles in the figure on the left. (From Kodjabachian and Meyer 1998a. Reprinted by permission of Taylor and Francis Ltd, <http://www.tandf.co.uk/journals>.)

correct values for each motor at each time step. 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 in principle does not introduce any constraint on what can be part of the self-organization process. Indeed, the characteristics of the sensors and of the actuators (Cliff and Miller 1996), the shape of the robot (Lund et al. 1997), and the architecture of the control system (Kodjabachian and Meyer 1998a, 1998b) can be included into the evolutionary process (we will return to this issue in section 1.4 and in chapter 9). Figure 1.3, for example, shows the results of an evolutionary experiment which will be further explored in chapter 9 where the architecture of the neural network controlling a tripod robot has undergone an evolutionary process.

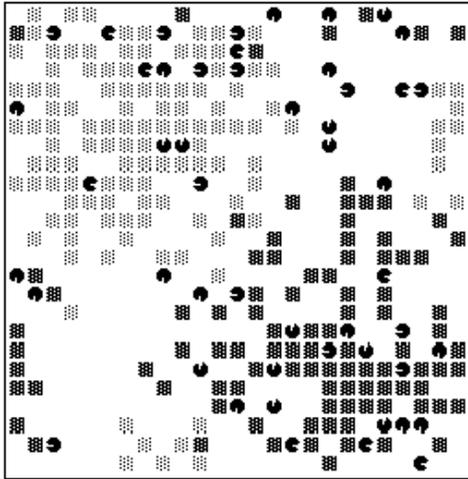


Figure 1.4

An example of agent for exploring Artificial Life. The environment is represented by a 2-dimensional grid of cells and the agent has sensors and actuators with infinite precision. (Reprinted with permission from Menczer and Belew 1997.)

Artificial life

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 dynamical systems and, from an experimental point of view, on the power of computers. A complex dynamical 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 physical devices (robots) instead of simulated agents (figure 1.4).

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; that is, they are grounded representations (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 sections, we will try to show the implications of evolutionary robotics for other disciplines. Although we think that evolutionary robotics may be relevant for many different fields, we will restrict our analysis to engineering, ethology, and biology. By doing so, it will become clear that the key characteristics of this approach is the possibility to rely largely on a self-organization process (Floreano 1997).

1.2 An engineering perspective

It is undisputed that behavioral systems such as autonomous mobile robots are difficult to design. 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 is 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 conventional strategy followed in order to overcome these difficulties has been *divide and conquer*: find a solution by dividing the problem into a list of hopefully simpler subproblems. 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 modulated 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 that both the layers (modules) responsible for simple basic behaviors and the coordination mechanism can be obtained through a self-organizing process rather than by explicit design (see Maes 1992; Mahadevan and Connell 1992; Dorigo and Schnepf 1993; Ram et al. 1994; Urzelai et al. 1998).

The approaches based on behavioral decomposition, however, still leave to the designer the decision of how to break the desired behavior down into simple basic behaviors. Unfortunately, it is not clear how a desired behavior should be decomposed and it is very difficult to perform such decomposition by hand. Even researchers who successfully adopted the behavioral decomposition and integration approach feel that this is a crucial problem. Rodney 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 1991a, pp. 584–585). Colombetti et al. (1996) note at the end of their article: “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.”

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 behaviors: a description from the observer’s point of view and a description from the robot’s point of view (Sharkey and 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 1.5).

It should be noted that the distal description of behavior is the result not only of the proximal description of behavior (i.e., the sensory-motor mapping), but also of the environment. More precisely, the distal description of behavior is the result of a 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, modify the environment or the relative position of the agent in the environment and therefore partially determine the type of sensory patterns that the agent will receive from the environment (we will come back to this point in the next section).

This dynamical interaction can explain why it is difficult to break down a global

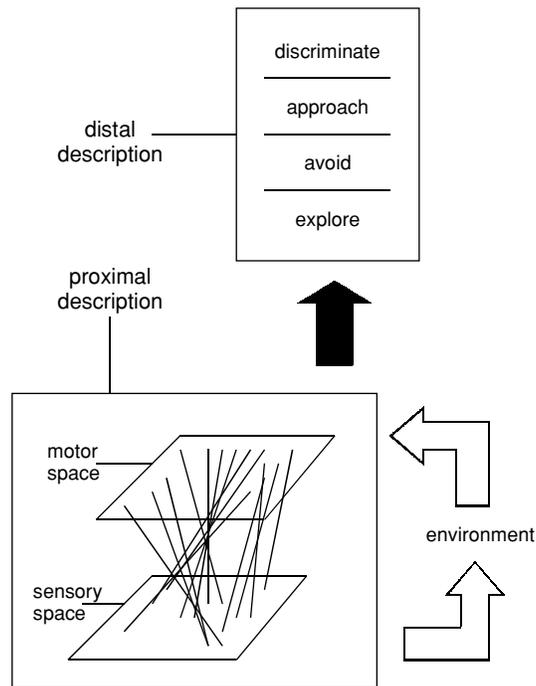


Figure 1.5

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 a distal description of behavior (top part of the figure) results from the dynamical interaction between a proximal description of behavior and the environment (bottom part of the figure).

behavior into a set of basic behaviors that 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 distal description of the global behavior. However, since 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 there is not a one-to-one mapping between subcomponents of the agent's control system and subcomponents of the corresponding behavior from the point of view of the

distal description, it is questionable whether a distal description of behavior can effectively be used to determine how the agent's control system should be structured, as hypothesized 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 the 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.

As we will show in section 4 of chapter 5, for example, it is extremely difficult to train a controller for a Khepera robot (Mondada et al. 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, the controller will not function properly. On the contrary, we can easily evolve a single network to produce the 4 different basic behaviors described above if we select individuals for their ability to perform the global behavior. Interestingly, in this latter case, it can be seen that a simple and homogeneous 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. 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 and Ziemke 1998.

In another experiment involving a Khepera robot evolved to clean an arena by carrying objects outside the walls with its gripper, simple uniform controllers have been compared with modular controllers (i.e., control systems divided into layers, or modules, used to control the robot in different environmental circumstances). As we will see in section 6.2, 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 taken by the experimenter. Moreover, we will see that, 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 particular characteristics of sensory-motor mapping, that is with a proximal description of behavior.

1.3 An ethological perspective

Molecules, cells, organs, organisms, and ecosystems are all entities at different levels that 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 the 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 that we use in everyday life to describe, predict, and explain our own behavior and the behavior of other living creatures (see figure 1.6).

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 psychophysics and neuropsychology which serve as bridges. However, it is very difficult to trace back observations concerning the nervous system and observations concerning behavior to the structure of a single entity that could 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 and McClelland 1986) and embodied cognition (Brooks 1991a; Varela et al. 1991; Pfeifer and Scheier 1999). 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.² Therefore, connectionism can be viewed as an attempt to overcome 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 and Guillot 1991; Wilson 1991).

Perception is often viewed as a disembodied process in which the agent is passively exposed to a continuously changing stream of sensory stimulation without consideration of what the agent needs to do. In contrast, a strong coupling between action and perception is a central issue in embodied cognition which tends to view perception on a need to know basis (Arkin 1990). By exploiting their interaction with the environment (i.e., by sensory-motor coordination), agents can partially modify the stimuli they receive from the environment.

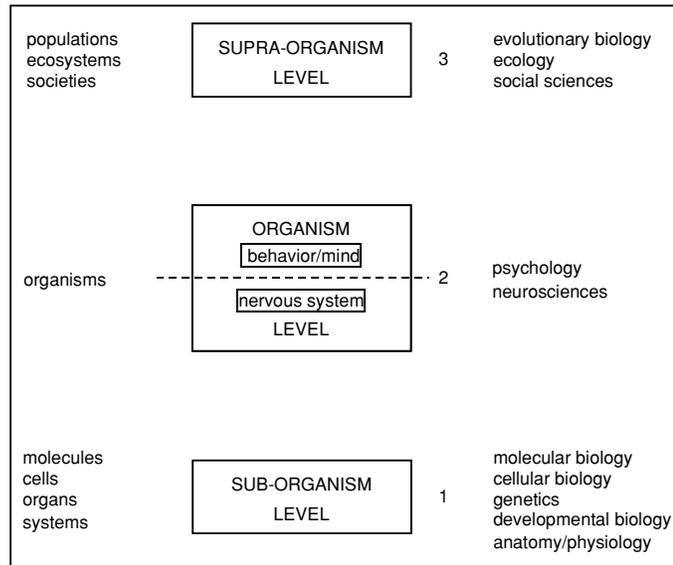


Figure 1.6

Three levels of entities relevant for the study of behavior and the scientific disciplines that study them. The radical separation between the concepts used by the neurosciences to describe the nervous system and the concepts used by psychologists to describe behavior and mental life creates a discontinuity within the intermediate level of the hierarchy. Courtesy of Domenico Parisi.

Therefore 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 own 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 and Beer 1997).

In most robotics research, however, the power of the interaction with the environment is largely unexplored. Few notable exceptions should be mentioned. 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 that the problem of perception could be greatly simplified if the agent's own movement and interaction with the environment is taken into account (Scheier and Pfeifer 1995). They described a Khepera robot that can learn to discriminate between small and large cylinders by relying on three simple predesigned basic behaviors (i.e., move forward,

avoid object, and turn toward object). 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 we pointed out in previous sections, adaptive behavior is difficult to obtain through design. In particular, it is hard to design systems that 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 performs with respect to a given task; (b) it partially determines the next sensory pattern that the agent will receive from the environment, which in turn may determine if the agent will be able to solve its task or not (we will return to this issue in chapter 5). 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.

Evolutionary robotics, by largely relying on self-organization, is not affected by these problems and, as a consequence, is an ideal framework for studying adaptive behavior.³ Indeed, many of the evolved robots exploit active interaction with the environment in order to maximize the selection criterion. For example, in section 5.4 we will see that individuals evolved for their ability to discriminate between walls and cylinders solved the task by moving back and forth in front of the perceived objects. 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 any sensory stimulation. Given a certain shape and dimension of the environment, this simple strategy ensured that long and short walls are encountered at significantly different angles. This in turn allows the robots to easily reach the two target corners by following or avoiding the walls depending on the angle at which they were approached (this experiment will be reported in section 5.5).

1.4 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 restated 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.

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, we 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 specific task. This amounts to design a fitness criterion that scores individuals according to their ability to solve the 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. We will refer to this fact 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 rewards for subparts of the desired task (an example of this technique will be described in section 6.2). 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*). This idea was previously

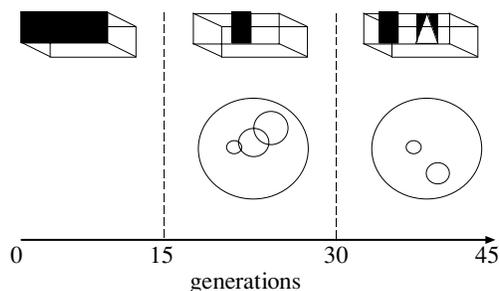


Figure 1.7

Incremental evolution of shape discrimination and sensory morphology (adapted from Harvey et al., 1994). The desired behavior is to approach a rectangle and avoid a triangle. The top part of the illustration shows the modification of the environment across generations. The lower part shows a schematic representation of the position and width of the receptive fields of the vision units used by the best evolved controllers (data not available for generation 15). The fitness function is significantly changed at generation 30.

proposed in the context of classifier systems by Dorigo and Colombetti (1994, 1997) who named it *shaping*, borrowing the term from 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 1.2). 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.

For example, Harvey et al. (1994) have incrementally evolved visually guided behaviors and sensory morphologies for a mobile robot expected to discriminate shapes and navigate towards a rectangle while avoiding a triangle (figure 1.7). The authors started with a simple version of the task that selected controllers for their ability to navigate towards a wall covered by color paper. Successively, the authors narrowed the area of the color paper on the wall and resumed evolution from the last evolved generation. Finally, the rectangular shape was displaced and a triangular shape was placed nearby. The fitness function was modified to encourage avoidance of the triangle and approach of the rectangle, and the previously evolved population was incrementally evolved in the new environment. The final evolved sensory morphology used only two pixels to detect whether the perceived shape was rectangular or triangular.

This experiment shows how manipulating the learning experiences throughout generations and introducing more constraints in the selection criterion one can solve tasks that cannot be solved otherwise, but there is a danger that additional constraints might 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 competing co-evolving populations with coupled fitness, such as predator and prey scenarios.

In chapter 8, we will investigate the conditions in which co-evolution leads to progressively more complex behavior in competing robots. We will consider the case of two competing species of predator and prey robots which are selected for their ability to catch prey and escape predators, respectively. 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. A consequence discussed by Dawkins and Krebs (1979) is that competing populations might reciprocally drive one another to increasing levels of complexity by producing an evolutionary “arms race.” For this reason, as we will see in section 8.7, in some cases co-evolution can solve tasks that evolution of a single population cannot.

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.

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 their whole lifetime. However, this huge amount of information encodes only very indirectly how well an individual did in different moments of its own lifetime or how it should modify behavior in order to increase its own fitness. The problem is how such

information can be transformed into an indication of what the agent should do or how well it is doing.

Because of the problems discussed in sections 1.2 and 1.3, it is probably hard to design a system capable of performing a good transformation. On the other hand, we may expect that evolution solve this type of problem by producing subsystems capable of autonomously extracting supervision information that can be used for fast lifetime learning. This has been shown in two computational experiments carried out by Ackley and Littman (1991) and Nolfi and Parisi (1997) that will be described in section 7.6. In both cases the controller architecture was divided into two submodules of which the first has the function of determining how to react to the current sensory state and the latter has the function of generating a teaching signal for the former. In Ackley and Littman sensory states were transformed into reinforcement signals while in Nolfi and Parisi sensory states were transformed into self generated teaching signals. By subjecting the weights of the two sub-networks to an evolutionary process, the authors reported evolution of individuals which learn during their lifetime to adapt to the environment through self generated teaching signals by transforming sensory information into useful reinforcement signal or teaching signals. As shown by Miller and Todd (1990) and Floreano and Mondada (1996a), a similar result can be obtained by evolving neural networks with topologies that may vary evolutionarily and that learn throughout lifetime with unsupervised learning (see section 7.5). In these cases the constraints on the architecture channel the changes driven by the sensory states in the right directions.

As we 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 capable of producing a behavior which is effective in a set of different environments is equivalent to another individual capable of adapting 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. 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 these individuals *plastic general*). On the other hand, individuals that have a general strategy already suitable for different environments do not need to change (let us call these individuals *full general*). From this point of view full general individuals will be more effective than plastic general individuals because they do not have to undergo an adaptation process throughout lifetime. However, it may happen that in certain conditions full general individuals 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 very low). If this is the case, a plastic general solution is the only available option.

Evidences that this may be the case even in relatively simple environments will be described in chapter 8. As we will see, if we try to co-evolve full general predator and prey robots by testing them against all the best competitors of previous generations we may fail. In fact, in most of the cases it is impossible to obtain predators or prey that are able to defeat a large number of competitors although it is always possible to easily find a set of different individuals with different strategies able to do so. In other words it appears that in some cases, full general strategies do not exist or are too difficult to find while a collection of simple strategies appropriate in different circumstances may be easily found. 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 will be described in section 8.8.

It should be made clear that, up to now, there has been only little experimental evidence within evolutionary robotics that evolution and learning can evolve more complex levels of competence than evolution alone (Floreano and Urzelai in press). However, the study of learning in an evolutionary perspective, is still in its infancy. We 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.

Development and evolution of evolvability

One of the aspects of natural evolution that is more crudely simplified in evolutionary robotics is development. This can explain why recently the importance of development is one of the most debated issues in evolutionary robotics. Interestingly, the importance and 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, “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 and Altenberg 1996).

The simplest genotype-to-phenotype transformation is a *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 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 (most of the experiments that will be described in this book involve different levels of organization); (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 (some examples will be described in section 9.2); (c) plasticity (see chapter 7); (d) genotypes that vary in length (see chapter 9).

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 genotype-to-phenotype mapping. Similarly, it is difficult to shape the mapping just by imitating nature given that not everything is known and that, for practical reasons, only some aspects of the natural process can be modeled.

All these problems and the fact that it is still not clear how the mapping itself may be subjected to the evolutionary process explain why only limited results have been obtained so far in the attempt to evolve more complex behaviors by using more biologically plausible genotypes.

1.5 Conclusions

In this chapter we claimed 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 measuring 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, complex behaviors can emerge from the interaction between several simple 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 make hypothesis about the critical features of natural evolution that allowed the emergence of the extraordinary variety of highly adapted life forms present on the planet.