

Toward Open-Ended Evolutionary Robotics: Evolving Elementary Robotic Units able to Self- Assemble and Self-Reproduce

Raffaele Bianco & Stefano Nolfi
Institute of Cognitive Sciences and Technologies, CNR
Viale Marx, 15
00137 Rome, Italy
e-mail: rbianco@ip.rm.cnr.it s.nolfi@istc.cnr.it

Abstract

In this paper we discuss the limitations of current evolutionary robotics models and we propose a new framework that might solve some of these problems and lead to an open-ended evolutionary process in hardware. More specifically, the paper describes a novel approach, where the usual concepts of population, generations and fitness are made implicit in the system. Individuals co-evolve embedded in their environment. Exploiting the self-assembling capabilities of the (simulated) robots, the genotype of a successful individual can spread in the population. In this way, interesting behaviours spontaneously emerge, resulting in chasing and evading other individuals, collective obstacle avoidance, coordinated motion of self-assembled structures.

1. Introduction

During the last ten years, researchers working in Evolutionary Robotics (ER) have shown how the attempt to develop robots through a self-organisation process based on artificial evolution can be used to solve problems that can hardly be solved with other methods (Floreano and Mondada, 1996; Nolfi, 1997; Hornby et al., 1999; Nolfi & Floreano, 1999). However, most of the researchers in the field would admit that ER is still unable to develop systems with the complexity and the robustness of even the simplest forms of natural life.

The limited power of current ER models with respect to natural evolution is probably due to several reasons. One possible reason, for instance, is that in ER the characteristics of the evolutionary process and the characteristics of the systems under evolution are crudely simplified with respect to nature. This might imply that ER experiments miss some of the crucial ingredients that prevent them from developing far more complex and interesting results. Another possibility is the fact that current ER experiments are below a crucial complexity thresholds with respect, for instance, to the richness of the environment and of the possible robot/environment interaction. One important limitation, however, is certainly our inability to identify the initial conditions that might lead to a truly open-ended evolutionary process able to lead to the development of a large variety of different individuals with different capabilities and to the development of new traits that constitute significant innovations.

In this paper we will discuss how some of the current limitations can be overcome and we will present a new ER framework that, within the current state of the art of technology, might potentially lead to an open-ended evolutionary process. In section 3, we will discuss the limitations of the current ER models and how these limitations can be overcome. In section 4 we will present our experimental setup. In section 5 we will describe the obtained results. Finally, in section 6 we will present our conclusions and our future research plan.

2. Open-ended evolution

By open-ended evolution we mean an evolutionary process that leads to a large variety of qualitatively different solutions and to the development of novelties, that is new traits that tend to be retained for long evolutionary periods and to constitute important building blocks for further evolutionary stages. Examples of major novelties discovered by natural evolution are: multi-cellular individuals, new cell types (e.g. the neural cells), new organs and systems (e.g. the central nervous system). For a more detailed account on how open-ended evolution and creativity can be characterized in the context of evolving embodied agents see Taylor (2001).

In the following three sub-sections we will discuss the three most important factors that, in our opinion, might promote open-ended evolution: (1) implicit and general selection criteria, (2) favourable organisation of the evolving individuals, and (3) changing environmental conditions. Previous researches addressing these three topics will be reviewed and discussed in section 3.

2.1 Implicit and general selection criteria.

Implicit selection criteria leave evolving individuals free to select their ways (among several alternatives) of solving their adaptive problems. In the context of artificial evolution, implicit selection criteria might be contrasted with explicit selection criteria that typically include several if-then conditions and specify in detail how evolving individuals should solve their adaptive problems. For example, a fitness function that scores evolving robots on the basis of the ability to not run out of energy is more implicit and admits more solutions than a fitness function that scores robots for the ability to always take the shortest path and periodically return to the recharging station. Evolutionary processes based on explicit selection criteria work like optimisation techniques in which the goal is to find the optimal value of a set of parameters and the solution of the problem is indicated by the experimenter beforehand. Evolutionary processes based on implicit selection criteria, instead, might lead to a sequence of qualitatively different solutions, and to solutions that are different from the expectations of the experimenter.

Moreover, the attempt to evolve robots able to solve general problems might more likely lead to open-ended evolutionary processes than the attempt to evolve robots that should solve specific problems. For example, evolving robots for the ability to walk might less likely lead to an open-ended evolutionary process than the attempt to evolve robots able to escape predators: a problem that might lead to the development of several abilities including the ability to walk. Similarly, the attempt to evolve robots able to navigate in an ever changing environment, by potentially requiring the development of continuously new abilities, might more likely lead to an open-ended evolutionary process than the attempt to evolve a robot in a static environment.

In the case of natural evolution, there are no selection criteria that determine whether individuals can or cannot reproduce aside from the ability to reproduce itself. This simple criterion, however, coupled with a rich and dynamically changing environment, indirectly introduces a cascade of adaptive needs such as: the ability to survive long enough to reproduce; the ability to compete with other individuals for limited resources; the ability to move in the environment, to build a nest etc. The open-ended character of natural evolution therefore can be ascribed, at least in part, to a simple selection criterion that combined with a dynamically changing social context has created the adaptive pressure for a progressively larger and larger set of abilities.

2.2 Favourable organisation of evolving individuals

The chance that variations at the genotypic level might generate significant innovations or brand new traits at the phenotypic level is not only a function of the adaptive needs resulting directly or indirectly from the selection criterion but also a function of the current organisation of evolving individuals. This aspect emphasizes the historical nature of the evolutionary process. The evolution of multi-cellularity, for example, has been a crucial prerequisite for many successive innovations. Similarly, the evolution of homeobox genes probably has been an important prerequisite for the synthesis of a large variety of different bauplans during the so-called Cambrian explosion (Gould, 2002).

The transition from single to multicellular individuals, in particular, has led to the emergence of individuals with different levels of organisations (i.e. cells, tissues, organs). The appearance of individuals that are much larger and more complex than their predecessors created the conditions for the emergence of new abilities and new adaptive traits (e.g. new ways to locomote, the possibility to exploit new source of energies and new environmental niches etc.).

An organizational properties that deserve special attention is metabolism, that is the organisms' capability to maintain their own organization (their identity) in the face of environmental perturbations that tend to alter such organization and in the face of the need to continuously regenerate their components. Organisms are not constituted by fixed material but are autopoietic systems, that is dynamic entities organized as a network of processes in which components and processes maintain their identity by regenerating themselves (Varela, 1997).

2.3 Dynamically changing environmental conditions.

The structure of the environment, including the social environment, and more specifically the richness and dynamical nature of environmental changes is another important factor that might condition the emergence of an open-ended evolutionary process. A certain level of stability in environmental conditions is necessary to allow evolving individuals to survive and adapt but, on the other hand, changing environmental conditions introduce new adaptive pressures that might promote the emergence of novelties. Indeed dramatic changes resulting from catastrophic events might have been one the major causes of the emergence of novelties in natural evolution (Gould, 2002).

Moreover, environmental changes occurring in the social context (either as an indirect result of the modification of the impact on the environment of the other individuals, or as a direct result of changes occurring in other evolving individuals) might lead to a process in which the struggle for survival becomes progressively more fierce. Interestingly these situations might lead to a sort of incremental evolutionary process in which the advantage of possible adaptations tend to become progressively stronger through out generations, thus progressively increasing the pressure to retain adaptive changes (Dawkins and Krebs, 1979).

Finally, the presence of a social environment consisting of other evolving individuals is obviously a necessary prerequisite for the development of cooperation (including social learning, cultural transmission etc.) between individuals. Cooperation might allow collections of individuals to solve problems that single individuals cannot solve. In some cases, cooperation might even lead to the emerge of radical new forms of individuals consisting of strongly cooperating entities who lose their individual autonomy. Examples of strongly cooperating elements formed by entities that lost (partially or totally) their individual autonomy are: (1) cells formed by elements, such us mitochondria, that previously existed as autonomous individuals, (2) multicellular organisms, and (3) social insects such as wasps and termites.

3. State of the art in ER

Although nobody has yet succeed in the attempt to trigger a truly open-ended evolutionary process, several experiments within ER have directly or indirectly addressed this issue. More specifically, several models described in the literature provided important insights into how the above three factors, that might promote open-ended evolution could be realized into an ER setting and on how they could affect the evolutionary process. In the following three subsections we will review these contributions and we will describe how all these factors can be integrated into a single model.

3.1 Implicit and general selection criteria

Nolfi and Floreano (2000) claimed that selection criteria (fitness functions) can be characterized as points in a fitness space with three dimensions: (1) a functional-behavioural dimension that indicates whether the fitness function rates specific functioning modes of the controller or whether it rates the behavioural outcome of the controller, (2) an explicit-implicit dimension that indicates the amount of variables and constraints included in the function, and (3) an external-internal dimension that indicates whether the variables and constraints included in the fitness function are computed using information available to the evolving agent. They also claimed that FEE functions (i.e. fitness functions that describe how the controller should work [functional], rate the system on the basis of several variables and constraints [explicit], and employ precise external measuring devices [external]), are appropriate to optimise a set of parameters for complex but well defined control problem in a well-controlled environment. On the contrary, BII functions (i.e. fitness functions that: rate only the behavioural outcome of an evolutionary controller [behavioural], rely on few variables and constraints [implicit], and rely on information that can be accessed on-board [internal]) are more suitable for developing adaptive robots capable of autonomous operations in partially unknown and unpredictable environments without human intervention.

Obviously BII fitness functions are also the most interesting from the point of view of promoting open-ended evolutionary processes. Indeed, the analysis of some experiments conducted by using BII fitness functions have shown that significant novelties do actually emerge during evolution, although the evolutionary process reaches a stable state after a limited number of generations in which no further innovations can be observed. For example, evolving the control system of a Khepera robot (Mondada et al., 1993) that is asked to move as fast and as straight as possible in an environment provided with a recharging area, Floreano and Mondada (1996) observed the emergence of robots able to periodically return to the charging station. Similarly, evolving control systems of a Khepera robot provided with a simple gripper for the ability to pick up objects and move objects outside the arena, Nolfi (1997) observed the emergence of several abilities not directly rewarded in the fitness function such as the ability to discriminate between walls and cylinders, the ability to approach cylinders and avoid walls etc. Moreover, by comparing the result of experiments with different fitness functions, Nolfi also showed that fitness functions with more variables and constraints did not speed up the evolution of the requested ability but, on the contrary, slowed it down (Nolfi, 1997).

One of the reasons that might explain why these experiments end rather quickly in a stable state and do not lead to the continuous development of innovations in an open-ended fashion is certainly the fact that the chosen selection criteria are not general enough. The problems chosen only admit a limited number of solutions and, once one of these solutions are found, there is no more room for further innovations. Unfortunately, however, the use of general and implicit selection criteria could cause a “bootstrap problem” (Nolfi, 1997), a problem that arises when all individuals of the first generations are completely unable to solve the problem on the basis of which they are selected, and are consequently all scored with the same null values.

One way to use a general and implicit selection criterion while avoiding the bootstrap problem is competitive co-evolution, that is the co-evolution of two competing species with coupled fitness (such as the co-evolution of two population of predators and prey). In this case, in fact, a certain

level of generality is assured by the fact that each population should be able to exhibit an effective behaviour in varying environmental circumstances (prey for example should be able to escape different type of predators). Moreover, the problem that being able to face hard competitors might be too hard for individuals with randomly selected genotypes, is solved by the fact that competitors are not smart from the beginning, given that they also are initially provided with randomly selected genotypes.

Indeed, experimental results obtained by co-evolving predator and prey robots indicate that the evolutionary process leads to a large variety of different solutions and does not tend to converge on a stable state after few generations as in the case of standard evolutionary experiments (Cliff and Miller, 1995, 1996; Nolfi and Floreano, 1999). Moreover, at least in some cases, competitive co-evolution can solve problems that standard co-evolution cannot solve (Nolfi and Floreano, 1999).

The fact that the solutions of a given selection criterion crucially depend on the structure and organisation of other evolving individuals who compete directly or indirectly for common resources, also explains why the absence of any selection criterion beside the ability to survive and replicate, has led to an open ended evolutionary process in nature. To our knowledge, nobody has yet tried to set up an experimental setting in which physical artefacts self-replicate autonomously and are not subjected to any selection criterion beside the ability to self-replicate effectively and compete with others for external limited resources. In this paper we will present the first attempt in this direction.

3.2 Favourable organisations of evolving individuals

The organisation of evolving individuals is determined by the current genome of the population and by the current rules that determine how genetic information affects the corresponding phenotypic structure (the genotype-to-phenotype mapping). Given that in most ER experiments the genome of the initial population is randomly generated and the genotype-to-phenotype mapping is determined by the experimenter, the problem of identifying favourable initial organisation corresponds to the problem of identifying a good genotype-to-phenotype mapping.

A good mapping should exhibit the following four properties (Grau, 1994; Nolfi and Floreano, 2000). A first requirement is *expressive power*, i.e. the possibility to encode many different phenotypic characteristics such as the architecture of the controller, the morphology of the robot, the rules that control the plasticity of the individual. Only the traits that are encoded in the genotype, in fact, can adapt evolutionarily and co-evolve with other traits.

A second related requirement is *completeness*, that is the possibility to potentially encode any possible phenotype into a corresponding genotype. As an example of limited completeness consider the case of genotype-to-phenotype mapping in which genetic variations might only lead to a limited number of different morphologies and control systems. A limited degree of completeness, might prevent the possibility to select effective solutions and more generally the probability to discover novelties.

A third requirement is *compactness*, that is genotype-to-phenotype mappings where the length of the genotype only weakly reflects the complexity of the corresponding phenotype. The length of the genotype, in fact, affects the dimensionality of the space to be searched by the genetic algorithm that, in turn, might significantly affect the result of the evolutionary process. Compactness could be achieved by means of different mechanisms. A first important mechanism consists in exploiting emergent properties that are not directly encoded in the genotype but emerge from the interactions occurring between genetic encoded properties, the phenotypic effects of these properties and the environment. A second important mechanism consists in the possibility of encoding repeated structures, such as similar sub-components of the controller, using a single set of genetic instructions.

Finally, a fourth requirement is *autonomy*, the ability of an individual to autonomously develop and reproduce without human intervention. Autonomy is not only important because it might make

possible a spontaneous and potentially never ending evolutionary process in hardware. It is also necessary in order to allow the evolutionary adaptation of the reproduction and developmental processes themselves.

The most advanced approaches in ER in this respect are those that, by following the seminal work of Sims (1995), allow the co-adaptation of the control system and of the morphology of the evolving robot. This is achieved by: (1) selecting a set of elementary parts including, solid parts (e.g. cylindrical objects), sensors, motorized joints, neurons, and neural connections, and (2) encoding in the genotype a set of growing rules that determine how the phenotype is constructed through the progressive additions of elementary parts. Models within this family differ mainly with respect to the grammar used to encode the growing rules (see for example Funes and Pollack, 1998; Hornby and Pollack, 2002; Bongard and Pfeifer, 2003).

Overall, the family of genotype-to-phenotype mapping used in these works has a good level of expressive power and of compactness but a rather limited level of completeness and completely lacks autonomy. Completeness is probably limited by several factors such as the characteristics of rigid materials, the difficulty of controlling complex structures made by several elementary elements and including several degrees of freedom. The lack of autonomy can be explained by considering the lack of artificial materials that can grow smoothly, change their shape continuously, and differentiate into different components such as sensors and actuators.

A possible alternative approach, that will be investigated in this paper, consists in the attempt to evolve robots made of elementary body elements able to self-assemble and self-reproduce autonomously, without human intervention. A similar attempt has been recently carried by Mytilinaios et al. (in press) who designed modular cubic units in which machines can construct and be constructed by other machine made of the same elementary units.

Although the idea of creating self-assembling and self-replicating physical systems was already proposed by Jon von Neumann in 1948, it has been investigated mostly in abstract systems such as cellular automata and artificial chemistries (Langton, 1984; Lohn and Reggia, 1997, Moore, 1970; Sipper and Reggia, 2001). Recent advances in technology and science of self-organizing and self-assembling machines (Pfeifer and Scheier, 1999; Bonabeau et al. 1999; Nolfi and Floreano, 2000; Camazine et al. 2001, Baldassarre et al., 2003a; Mondada et al., 2004; Guo et al., 2004), however, provide new theoretical and experimental tools that makes this ambitious goal reachable in the near future.

In principle experimental settings (like the one presented in this paper) in which evolving individuals are constituted by collection of interacting physical entities able to self-assemble and self-organize and in which individuals need to maintain their organization in order to self-replicate might also lead to the emergence of forms of metabolism or more generally to autopoietic systems (i.e. systems that actively try to maintain an organization which assure their own durability despite environmental perturbations). For examples, elementary units might self-assemble forming full-formed individuals with a given shape and size capable to actively maintain their full-formed organization by seeking and recruiting new elementary units to compensate the loss of body parts resulting from collisions with other individuals. Although this aspect will not be investigated in the experiments that will be described in the following sections, this issue might represent an interesting research line. For attempts to evolve robots able to display forms of autopoietic organization through neural plasticity, the reader may see Beer (1995) and Di Paolo (2000).

3.3 Dynamically changing environmental conditions

Most of the research conducted within ER involves single robots interacting with a static and simple environments. In these cases the robots forming the evolving population interact only indirectly through the fitness function. On the importance of considering the design of the environment and of the interactions between individuals, as well as the design of the individuals themselves, see Taylor

(2001). However, in addition to the notable exception constituted by the co-evolutionary experiments reported in section 3.1, some recent work involves the evolution of several robots sharing the same environment.

A pioneering work was conducted by Martinoli (1999) who evolved the control system of a group of simulated Khepera robots that were asked to find “food items” randomly distributed in an arena. The robots were provided with 10 infrared sensors and two motors controlling the two corresponding wheels. In some cases the evolved individuals displayed interesting collective behaviours such as exploring the arena in couples. Quinn et al. (2002) evolved teams of 4 mobile robots provided with infrared sensors and two motors controlling the two corresponding wheels necessary for the ability to move while keeping close together. Baldassarre et al. (2002a, 2003b) evolved teams of 4 mobile robots provided with infrared sensors, sound sensors and a simulated sound organs for the ability to aggregate and to move together toward a light target. The same authors (Baldassarre et al., 2003a, 2004) evolved the control system of a team of simulated robots provided with two motors, a traction sensor, and light sensors able to display coordinated movements, explore an environment by avoiding walls, and cooperatively push an object toward a light target.

Watson et al. (2002) evolved the control system of a population of eight robots provided with two motorized wheels, infrared and light sensors that were selected for the ability to approach light sources. Differently from the case of the experiments described above, the eight robots are not homogeneous (i.e. they do not have identical neural controllers) but have different genotypes and therefore potentially different behaviours. Evolution is accomplished through periodically exchanges of gene materials between robots located nearby in the environment. Selection of adaptive traits is assured by the fact that the probability of sending genes to other robots is proportional to the fitness of the robot itself (i.e. the ability to reach the light source).

In the model presented in this paper, individuals reproduce and grow by using the bodies of other individuals. Reproduction is accomplished by the possibility that individuals have to physically interact with a social environment in order to propagate their genotype in the body of other individuals. Growth is accomplished by self-assembling, that is the ability to physically connect to other individuals. Growth also allows: (a) the formation of individuals organized hierarchically in two levels (at least): the level of the individual and the level of the elementary units constituting the individual, and (b) the emergence of entities with adapted morphologies. Indeed, the way in which elementary units physically assemble determines the morphology of the corresponding individual. Finally, complex individuals formed by collection of self-assembled elementary units behave according to a distributed control system in which each elementary unit acts independently on the basis of local sensory information and in which elementary units coordinate so to allow the entire individual to display a coherent effective behaviour.

4. Experimental Setup

4.1 Vision

Imagine an experiment in which several elementary robot units able to move, to sense the environment, and grasp external objects (to connect/disconnect with other elementary robot units) are placed in a bounded environment. Each elementary unit is provided with a genotype (initially randomly generated) that encodes the properties of the corresponding control system. When an elementary robot unit connects to another elementary unit, it copies (with mutation) its own genotype into the new unit. Finally, each time step elementary units have a certain probability to die. Dead elementary units are recycled by being supplied with a mutated version of their original genotype.

If we leave the elementary robot units free to move and interact among themselves and with the external environment we should expect that, after some time, some of the units will happen to

connect to other elementary units thus propagating mutated copies of their genotype in the population. Since the elementary units that will connect to other elementary units first will be provided with genotype and control systems that are better than the average, we should expect that the ability of the units to reach and connect with other units will increase in the population as time goes by. Moreover, given that units that are connected receive mutated version of the genotype of the connecting units and given that adaptive mutations tend to spread in the population, we should expect the triggering of an evolutionary process in which progressively better behavioural strategies are observed in the population.

Once two elementary units self-assemble, higher order individuals composed by two connected elementary units are formed. We will call individuals formed by at least two connected elementary units, type-2 individuals, that is individuals in which one can recognize two hierarchical levels of organisation (the level of the individual and the level of the individual's elementary constituents). Initially, the elementary units forming type-2 individuals will not be capable of cooperating. Therefore, type-2 individuals will not be able to survive for long period of times and to propagate their genotype (i.e. reaching and grasping other elementary units). However, type-2 individuals that will be slightly better in coordinating will have more chances to increase the frequency of their genes in the population. The continuation of this process will progressively lead to better and better type-2 individuals able to show coherent and effective behaviours.

Self-replication of type-2 individuals might occur through disconnection and self-assembling. For example, the death of an elementary unit assembled into a type-2 individual formed by three elementary units might cause the disconnection of the other two elementary units that, by self-assembling with other elementary units, might form two type-2 offspring. Offspring will be similar to their parent (i.e. they will have a similar genotype, a similar behaviour, a similar shape thanks to the genetically encoded propensity to assume a given shape). The tendency to assemble in order to form a genetically determined shape is accomplished through: (1) a genetically encoded preference for connecting with other elementary units along given relative directions, and (2) the possibility of connected elementary units to "migrate" within the body of the individual up to the tail (i.e. to the last assembled unit) and to assemble with a genetically determined angle.

In principle, the same process might lead to the further development of type-3 individuals, that is individuals in which an observer can recognize three hierarchical levels of organisation: the level of the individual, the level of the parts composing the individual, and the level of the elementary units composing each part. The investigation of this issue is part of our future research plan, given that it would require experiments involving a much larger number of elementary units than those used in the experiments described in this paper.

4.2 The elementary robotic units, the environment, and the rules governing the interactions

The elementary robotic units consists of a simplified version of s-bots (Mondada et al., 2004), mobile autonomous robots capable of self-assembling that we are currently manufacturing within an European R&D project called Swarbots (see also www.swarm-bots.org).

Each robot (Mondada et al., 2004) has a cylindrical body with a diameter of 116 mm and consists of a mobile base ("chassis") with two motors controlling a track and a teathed wheel (Figure 1, left), and a main body ("turret"). The turret is provided with two grippers, one rigid and one flexible, that allow the robots to self-assemble and to grasp objects. Each robot is provided with a number of different sensors (Mondada et al., 2004), but only the infrared sensors and the traction sensor described below have been used in the experiments reported in this paper.

To test the feasibility of the model in a simple experimental setting, and given that these robots are still under fabrication, we simulated the robots and the environment. The simulator was based on the SDK VortexTM toolkit (Critical Mass Labs, Canada), which allows programming realistic simulations of dynamics and collision of rigid bodies in 3D.

Each simulated robot (Figure 1, right) is composed by a chassis with two passive and two motorized wheels and a cylindrical turret connected to the chassis by a hinge joint that can rotate along the vertical axis. Each robot also has a physical link through which it can physically connect to another robot along the perimeter of its turret. Connections between robots are simulated by creating a joint between the turrets of the corresponding robots.

Each robot is provided with a traction sensor placed at the turret-chassis junction that detects the direction (with respect to the direction of the chassis) and the intensity of the traction force that the turret exerts on the chassis. Concerning the actuators, each robot can control the two motorized wheels independently.




Figure 1. **Left:** An s-bot. **Right:** A simulated simplified version of the robot shown in the left part of the figure. Grey cylinders and spheres represent the motorized and the passive wheels, respectively. The grey circular and the white rectangular parts represent the turret and the chassis, respectively.

The genotype of each robot encodes the connection weights of the robot neural controller that consists of a simple perceptron with a fixed architecture (see Figure 2). The neural controller has 14 sensory neurons directly connected to two motor neurons. Eight sensory neurons encode the activation state of eight corresponding infrared sensors distributed around the chassis that allow the detection of other robots' bodies up to a distance of 180 cm. Four sensory neurons encoded the intensity of the traction from four different preferential orientations with respect to the chassis's orientation (front, right, rear, left). Each sensory neuron has an activation proportional to the cosine of the angle between the sensor's preferential orientation and the traction direction when the angle is within $[-90, +90]$ degrees, and has an activation of 0.0 otherwise. One sensory neuron is activated when the robot is connected to another robot (when a robotic unit connects and later disconnects, this neuron linearly decreases its activation state during the succeeding 4 seconds, see section 5.1). One input neuron is a bias unit that is always fully activated. The activation state of the motors units is normalized between $[-5, +5]$ rad/s and is used to set the desired speed of the two corresponding wheels. Connection weights are encoded with 8 bits and normalized in the range $[-10.0, 10.0]$. The genotype of each robot also encodes the angular orientation with which caught robotic units should be connected to the tail of the winning individual. The angular orientation of the connection is encoded with 8 bits and normalized in the range $[-45, +45]$ degrees.

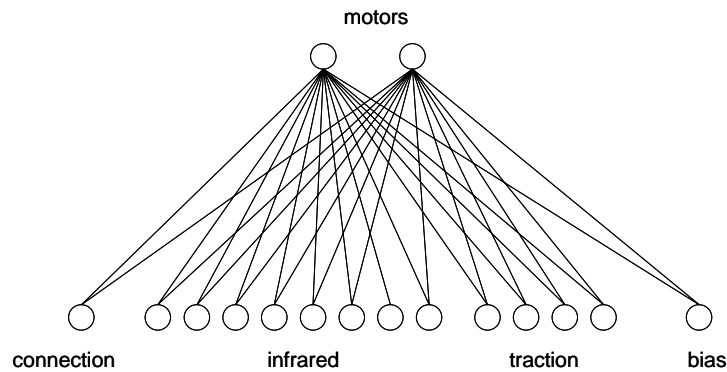


Figure 2. The architecture of the neural controller

When the distance between the elementary units of two individuals decreases below a given threshold, a connection link is created between the two elementary units and the genotype of the winning elementary unit is copied (eventually with mutation, see below), on the genotype of the losing unit. Winners and losers are determined by the potential physical impact of the movement of one individual on the other individual, and vice versa, on the basis of the current speed and of the relative direction of movement of the two individuals. In other words, the individual that moves faster and straighter toward the other individual has the possibility to connect to the other individual (or to one of the elementary unit of the other individual if this is formed by several elementary units) and to inject its genotype in the other individual. The connected elementary unit is then transported up to the tail of the winning individual and connected according to the preferential angular orientation encoded in the genotype of the elementary units corresponding to the tail.

When the losing individual is formed by two or more elementary units, the genotype injected into the connected elementary unit propagates on all other elementary units constituting the losing individual, the links between the elementary units of the losing individual are disconnected, and the unconnected losing elementary units are immobilized for the successive 50 seconds (during this time, they might receive connection from other units).

Every time step elementary units have a probability of 0.0133% to die. Dead units receive a mutated version of their genotype, if connected are disconnect from other units, and are moved in a new randomly selected location. During the mutation process each bit of the genotype have a probability to be inverted of 0.6%. Elementary units also die when they are caught by other individuals. When the genotype of a winning unit is injected into the body of another elementary unit, each bit of the genotype have a probability to be inverted of 0.0% or 0.01% (i.e. all experiments were replicated in the two conditions, see below).

A population of 64 elementary robotic units with randomly generated genotypes is placed in a square environment surrounded by walls (Figure 3) and is left free to move and interact autonomously. Experiments are stopped after 15 million time steps (each time step lasts 100 ms). At the beginning of each experiment, the elementary robotic units are placed with a uniform distribution and a randomly selected orientation (see Figure 3).




Figure 3. The environment including 64 elementary robotic units at the beginning of an experiment.

5 Results

In this section we will describe the results of three sets of experiments. The first two, consist of a simplified version of the experimental setup described in the previous section. The third experiment, instead, includes all aspects of the experimental setup described above. Each of the three experiments were replicated 5 times.

5.1 Emergence of an ability to survive and reproduce

In the first experiment robots are left free to interact but are not allowed to self-assemble. At the beginning of the experiment, robots are provided with randomly generated genotypes and therefore tend to display trivial behaviours such as moving by producing circular trajectories or moving straight and crashing into walls. However, after some time, some of the robots reach other individual robots and propagate their genotype in the population. Given that the robots that succeed in propagating their genotype are those that, initially by chance, display behaviours that are slightly better than the others (from the ability to survive and reproduction point of view), the genotypes of better individuals tend to propagate in the population and the genotypes of worse individuals tend to be extinguished. This differential reproduction process combined with the production of new traits (resulting from mutations) triggers an evolutionary process in which new adaptive traits tend to be retained and in which individual robots tend to display progressively better behaviours (from the ability to survive and reproduction point of view).

Indeed, by analysing the behaviour of individual robots during the evolutionary process in five replications of the experiment, we observed that after approximately 2 million time steps, some of the robots of the population display forms of behaviours that maximize their ability to survive and reproduce. Survival is typically accomplished by moving as straight and as fast as possible, in order to minimize the risk of being caught by other robots, and by quickly avoiding walls. Reproduction is maximized by the ability to catch other robots by chasing them and by trying to anticipate their movements.

Figure 4 shows a typical evolved behaviours. During the first 50 cycles displayed in the figure robot C chases and reaches the robot D thus propagating its genotype into the robot D (Figure 4, see white arrows). During the next 50 cycles, the individual D reaches the robot F thus further propagating the genotype received by the individual C (Figure 4, black arrows). Notice that by exploiting the sensory neuron that becomes active when a robot connects to another robot, injects its genotype, and then immediately disconnects, robots do not connect to the same robot twice. They

avoid the robots placed nearby during the 4 seconds in which the activation of this sensor linearly decreases in order to propagate their genotype to other robots that are less genetically related.

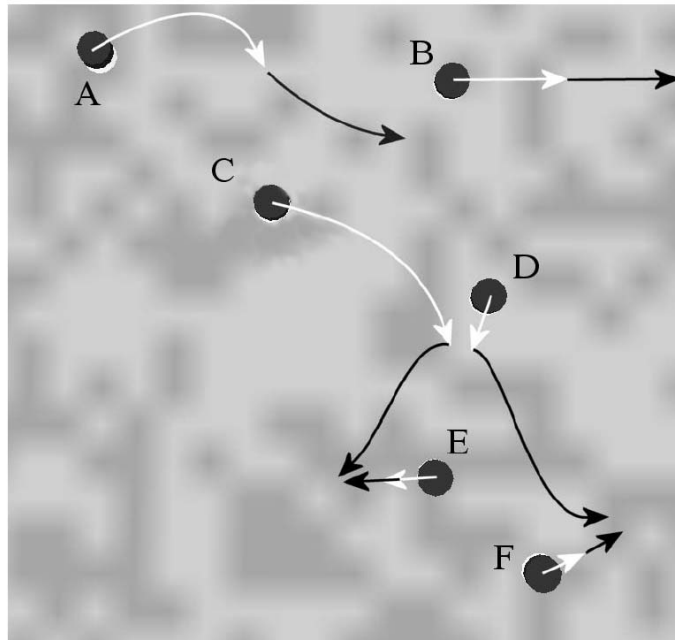


Figure 4. A typical evolved behaviour after 5 million time steps in one of the replications of the experiment. Grey circles represent some elementary units moving in a small portion of the environment. White and black arrows indicate the trajectory displayed by the units during two successive phases of 50 time steps, respectively.




Figure 5. **Left:** the number of robots caught by the best evolving individuals (with most elementary units caught). **Right:** average lifetime of the individuals able to survive longer. Both graphs refer to the same replication of the experiment described in Figure 4 and indicate the moving average of 200 time steps.

Performance measures taken during the evolutionary process unfortunately provide little indication on what is going on. This can be explained by considering that the absolute performance of individuals (for example with respect to the ability to survive and reproduce) strongly depends on the behaviour of the rest of the population that also varies evolutionarily. Indeed, the analysis of the number of elementary units caught by individuals throughout evolution (Figure 5, left) and the average length of the lifetime of the individuals (Figure 5, right) do not display any clear trend.

By testing evolving individuals in environments populated by other individuals with randomly generated genotypes a progress in the ability to catch other robots can be observed (see Figure 6, left). However, this is only a very rough description of what really is going on during evolution. In fact, individuals that have the same absolute performance when placed in a population formed by

individuals with randomly selected genotypes might have rather different performances when placed in a population of evolved individuals. The analysis of the ability to survive of evolving individuals tested in a population composed by individuals with highly effective phenotypes (considering the ability to catch other robots), does not indicate any clear trend (see Figure 6, right). Whether techniques successfully used to analyse competitive co-evolving population (Cliff and Miller, 1995; Nolfi and Floreano, 1999; Rosin and Belew, 1997) might be useful in this case or whether new techniques should be developed is part of our future research plan.

Similar results were observed in the five replications of the experiment and in both experimental conditions (i.e. with and without mutation during the propagation of genotypes into caught elementary units).




Figure 6. **Left:** the average number of robots caught by evolving individuals. Data obtained by testing 6 evolving individuals in a population formed by 58 individuals with randomly generated genotypes. **Right:** the average length of the lifetime of individuals. Data obtained by testing the performance of 6 evolving individuals placed in a population formed by 58 individuals with the genotype of the individual that scored best with respect to the test described in the left side of the Figure. Both graphs are obtained by computing the average performance over 10 trials lasting 3000 cycles. Performances are computed by testing evolving individuals every 500.000 time steps. Both data of Figure 5 and 6 refer to one of the best replications of the experiment in the condition in which mutation rate during genotype injection is 0.0%.

5.2 Emergence of type-2 individuals able to coordinate

In this second experiment robots are allowed to self-assemble by connecting together but assembled elementary units are not allowed to migrate and to assemble with the genetically encoded angular orientations. Therefore the shape of type-2 individuals cannot evolve. The shape simply results from the way in which individuals get in contact and assemble.

As in the case of the experiments described in the previous section, the visual analysis of the behaviour of the evolving robots indicates that they tend to display progressively better behaviours from the point of view of survival (i.e. avoiding to be caught by other robots) and reproduction (i.e. propagating their genotypes by chasing and reaching other robotic units). In this case however, winning robots do not only propagate their genotypes on the body of losing robots but also physically connect to them. This self-assembling leads to an evolutionary process in which two type of individuals (type-1 and type-2 individuals) at two different hierarchical levels of organisation are subjected to the evolutionary process.

It is important to notice that although both type-1 and type-2 individuals tend to be selected for their ability to survive and reproduce, the formation of type-2 individuals introduces new challenges and new opportunities for evolving individuals. The new challenges, for example, are due to the fact that the elementary robotics units constituting the type-2 individual should be able to cooperate in order to survive and reproduce. Assembled individual units that are not able to display coordinated movements, e.g. being unable to move when their chassis are oriented toward different directions, in fact, will not be able to reproduce and will be quickly captured by other individuals. The new opportunities are due to the fact that, by being larger and by having more sensory information than

type-1 individuals, type-2 individuals might display better survival and reproduction capabilities than type-1 individuals.

Indeed the analysis of evolved individuals showed how, in all replications of the experiments, evolved type-2 individuals display an ability to produce coordinated movements, coordinated obstacle avoidance, and coordinate hunting behaviours. By looking at a typical evolved behaviour (figure 7) one can observe how: (a) type-1 individuals are able to effectively reproduce by chasing and reaching other individuals, (b) new formed type-2 individuals are able to negotiate a common direction and to keep moving along such a direction until they detect an individual to chase, (c) type-2 individuals display an ability to modify their direction in a coordinated manner in order to avoid walls and effectively chase and reach other type-1 and type-2 individuals. Please notice that disassembling resulting from spontaneous death of elementary units prevent the formation of larger and larger clusters.

Figure 7 represents a typical evolved behaviour in which an elementary robotic units reaches and connects to another elementary unit. The two elementary units forming this new type-2, by producing an effective cooperative behaviour, reach and connect to a third elementary unit forming a type-2 individual constituted by three elementary units.

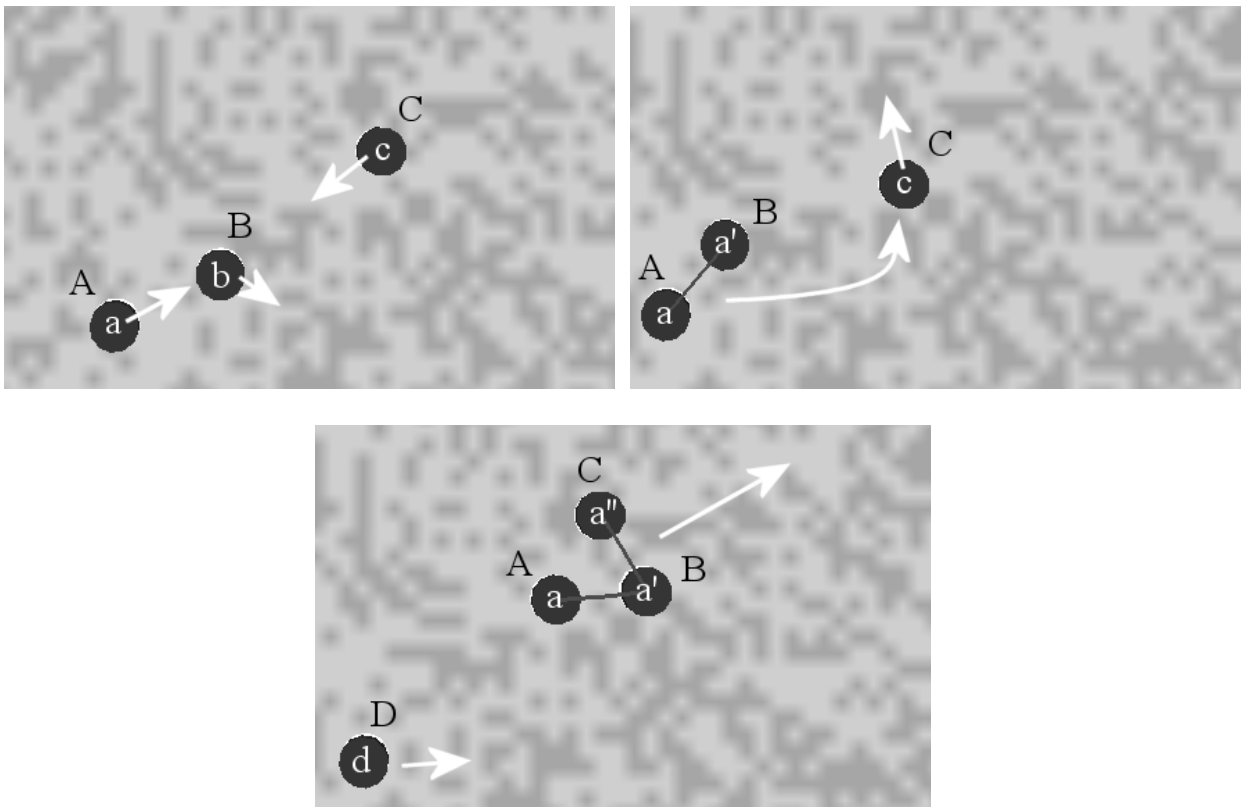


Figure 7. A typical evolved behaviour. **Top-left:** The elementary unit A reaches and connects to the elementary unit B and injects its genotype into the connected unit. **Top-right:** The type-2 individual formed by the elementary units A and B reaches the elementary unit C. The elementary unit B connects to the unit C and injects its genotype so that the type-2 individual is now formed by three elementary units with related genetic material. **Bottom-left:** The type-2 individual formed by units A, B and C moves in a coordinated fashion.

The negotiation of a common direction of movement in type-2 individuals is accomplished by each elementary robotic unit exploiting the information provided by the traction sensor that indirectly provides an indication of the average direction of movement of the type-2 individual overall (Baldassarre et al., 2003a, 2004). The information provided by the traction sensor, also plays a key role in the ability to collectively avoid walls. Indeed, collisions with obstacles, by producing a traction force toward the opposite direction with respect to the direction of movement of the robot,

allow colliding robots to turn thus avoiding obstacles. These turning behaviours produce a traction force in other connected robots that tend to turn accordingly so that the whole group produces a collective and coordinated obstacle avoidance behaviour (Baldassarre et al., 2003a). Similarly, the ability to collectively chase and reach other individuals results from a combination of the ability of the single elementary units to chase and catch other individuals that are in sight and the ability to coordinate through the traction sensor.




Figure 8. The average number of robots caught by evolving individuals. Data obtained by testing 6 evolving individuals in a population formed by 58 individuals with randomly generated genotypes. Performances are computed by testing evolving individuals every 500,000 time steps. Data refer to one of the best replications of the experiment in the condition in which mutation rate during genotype injection is 0.0%.

The average number of elementary units caught by evolving individuals tested in a population made of individuals with randomly selected genotypes only shows a progress during the first million time steps (Figure 8). As pointed out above, however, this test only provides a very rough indication of what is going on. In fact, individuals that have the same absolute performance when placed in a population formed by individuals with randomly selected genotypes might have rather different performances when placed in a population of evolved individuals.

Similar results were observed in the five replications of the experiment and in both experimental conditions (i.e. with and without mutation during the propagation of genotypes into caught elementary units).

5.3 Emergence of type-2 individuals with co-adapted body shapes

This third and last experiment involves individuals that, as in the case of the experiment described in the previous section, can self-assemble but that are also able to rearrange their shape according to their genetically encoded preference. More precisely “caught” elementary units are transported to the tail of the winning individual and connected according to the angle (ranging from -45 and 45 degrees) indicated in the genotype of the connecting elementary unit.

Analysing the evolutionary process we observed: (a) the emergence of individuals able to self-assemble, (b) the emergence of type-2 individuals formed by several connected elementary units able to coordinate and to exhibit coordinated behaviours, and (c) the emergence of type-2 individuals with well defined shapes such as straight, left-handed, and right-handed chains.

Remarkably, the visual analysis of evolving individuals (Figure 9) and the analysis of the performance obtained by testing evolving individuals in populations consisting of individuals with randomly generated genotypes (Figure 10) indicate that performances, with respect to the ability to reach and catch other individuals, is significantly better in this experiment than in the second experiment in which evolving type-2 individuals could not co-evolve their control system and their morphological shape.




Figure 9. A typical evolved behaviour in one of the replication of the experiment. **Top-left:** The elementary unit B reaches and connects to the elementary unit C and injects its genotype into the connected unit. **Top-right:** The type-2 individual formed by the elementary units B and C quickly turns and moves toward the elementary unit A. **Bottom-left:** The type-2 individual formed by units B and C reaches and connects to the unit A. **Bottom-right:** The type-2 individual formed by units B and C and A start to move in a coordinated fashion.

Figure 9 is an example of a typical evolved behaviour. The figure shows how the genotype of the elementary unit B propagates to the other two elementary units during self-assembling and how, the genetically encoded angular preference of unit B leads to the formation of a right-handed chain morphology.

Figure 10 indicates the performance of evolving individuals tested against individuals with randomly generated genotypes in the case of the best replication of the second and third experiment (thin and thick lines respectively). It shows that, the best individuals of the third experiment outperform those of the second experiment from the viewpoint of spreading their genotype by reaching and catching other individuals (both data refer to the best corresponding replications of the experiment in the condition in which mutation probability during genotype injection is 0.0%).

Similar results were observed in the five replications of the experiment and in both experimental conditions (i.e. with and without mutation during the propagation of genotypes into caught elementary units).




Figure 10. The average number of robots caught by evolving individuals. Data obtained by testing 6 evolving individuals with 58 other individuals with randomly generated genotypes. Both curves are obtained by computing the average performance over 10 trials lasting 3000 cycles. Performances are computed by testing evolving individuals every 500.000 time steps. Thin and thick lines represent the performance in the case of the second and third experiments, respectively. Both data refer to one of the best replications of the two experiments in the condition in which mutation rate during genotype injection is 0.0%

6. Conclusions

In this paper we discussed the limitations of current evolutionary robotics models and we proposed a new framework that might potentially solve some of these problems and lead to an open-ended evolutionary process in hardware.

The framework proposed involves a population of autonomous elementary robotic units that are left free to interact and to self-assemble in a simple environment. The possibility of self-assembling and propagating their genotype into the body of assembled units leads to a spontaneous evolutionary process without the need for explicit fitness functions. Moreover, the ability to self-assemble and the way in which self-assembly is affected by genetically encoded traits allows the emergence of individuals that: (a) are formed by several connected elementary units able to coordinate and cooperate to display a coherent behaviour, and (b) arrange themselves so as to form body shapes that are adapted to the behaviour exhibited by the individual.

The results of the experiments performed in simulation are rather encouraging and demonstrate how, by selecting a set of simple rules that determine the interactions between elementary units, one can trigger an evolutionary process that might spontaneously lead to the emergence of new adaptive needs (e.g. the need to produce coordinated movements) and to the development of progressively better solution to these needs.

The model proposed includes several important innovations with respect to the majority of the works conducted in evolutionary robotics. Namely:

- (1) As the Embodied Evolution Model proposed by Watson et al. (2002), it consists of an autonomous, fully distributed and asynchronous evolutionary algorithm that, by not requiring any centralized sorting procedure, potentially allows to trigger a completely autonomous and hand-free evolutionary process in hardware. This feature in turn potentially allows the parallel evaluation of a large number of evolving robots.
- (2) Based on elementary robotic units that are able to autonomously self-assemble, it potentially allows to co-evolve the morphology and the control system of the robots in hardware.
- (3) Not relying on an explicit fitness function and being based on roughly the same selection criterion that shaped natural evolution (e.g. simply the ability to survive and reproduce) it might potentially lead to a truly open-ended evolutionary process.

In future works we plan to replicate these experiments in hardware using the robots developed within the Swarm-bot project (Mondada et al., 2004) as basic elementary units and to investigate the possibility of developing better suited hardware elementary units. Moreover we plan to investigate more powerful ways to genetically encode the morphology of evolving type-2 individuals.

Acknowledgements

Stefano Nolfi thanks Gianluca Baldassarre, Dario Floreano, John Hallam, Auke Ijspeert, Luca Gambardella, and Gianluca Massera for useful comments on a preliminary version of the ideas

illustrated on this paper. The authors thanks Tim Taylor, Tom Ziemke, Vito Trianni, and Auke Ijspeert for their comments and suggestions of the first draft of this paper. This research has been supported by the SWARM-BOT project founded by the Future and Emerging Technologies program (IST-FET) of the European Community under grant IST-2000-31010. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that may be made of data appearing in this publication. The Swiss Government supported the Swiss project's participants (grant 01.0012).

References

- Baldassarre G., Nolfi S. & Parisi D. (2002). Evolving mobile robots able to display collective behaviours. In C.K. Hemelrijk (Ed.), *Proceedings of the International Workshop on Self-Organisation and Evolution of Social Behaviour*. Zurich, Switzerland: Swiss Federal Institute of Technology.
- Baldassarre G., Nolfi S. & Parisi D. (2003a). Evolution of collective behavior in a team of physically linked robots. In S. Cagnoni et al. (Eds.), *Applications in Evolutionary Computing, Lecture Notes in Computer Science*. Berlin: Springer Verlag.
- Baldassarre G., Nolfi S. & Parisi D. (2003b). Evolving mobile robots able to display collective behavior. *Artificial Life*, 9: 255-267.
- Baldassarre G., Parisi D. & Nolfi S. (2004). Coordination and behavior integration in cooperating simulated robots. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam & J-A. Meyer (Eds.), *From Animals to Animats 8: Proceedings of the VIII International Conference on Simulation of Adaptive Behavior*. Cambridge (MA): The MIT Press.
- Beer R.D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence* 72:173-215.
- Bonabeau E., Dorigo M. & Theraulaz G. (1999). *Swarm intelligence*. Oxford: Oxford University Press.
- Bongard, J.C. & Pfeifer R. (2003) Evolving complete agents using artificial ontogeny. In F. Hara & R. Pfeifer, (Eds.), *Morpho-functional Machines: The New Species (Designing Embodied Intelligence)*. Berlin: Springer-Verlag.
- Camazine S., Deneubourg J.L., Franks N.R., Sneyd J., Theraulaz G. & Bonabeau E. (2001). *Self-organization in biological systems*. Princeton, NJ.: Princeton University Press.
- Cliff D. & Miller G.F. (1995). Tracking the red queen: Measurement of adaptive progress in co-evolutionary simulations. In F. Moran, A. Moreno, J. J. Merelo & P. Chacon (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*. Berlin: Springer Verlag.
- Cliff D. & 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.
- Dawkins R. & Krebs J.R. (1979). Arms races between and within species. *Proceedings of the Royal Society of London B*, 205, 489-511.
- Di Paolo E.A. (2000). Homeostatic adaptation to inversion of the visual field and other sensory-motor disruptions. In J-A. Meyer, A. Berthoz, D. Floreano, H. Roitblat, & S. Wilson (Eds.) *From Animals to Animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behaviour*. Cambridge (MA): MIT Press.
- Hornby G.S., Fujita M., Takamura S, Yamamoto T. & Hanagata O. (1999). Autonomous evolution of gaits with the Sony quadruped robot. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela & R. E. Smith (Eds.), *GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference*. San Francisco, CA: Morgan Kaufmann.
- Floreano D. & Mondada F. (1996) Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics--Part B: Cybernetics*, 26(3):396-407.

- Funes, P. & Pollack, J. (1998) Evolutionary body building: Adaptive physical designs for robots *Artificial Life*, 4: 337-357.
- Gould S.J (2002). *The structure of Evolutionary Theory*. Harward University Press.
- Gruau F. (1994) Automatic definition of modular neural networks, *Adaptive Behavior*, 3:151-183.
- Guo Y., Poulton G.T., Valencia P. & James G.C. (2004). Designing self-assembly for 2-dimensional building blocks. In G. Di Marzo Serugendo et al. (Eds.), *Engineering Self-Organising Systems: Nature-Inspired Approaches to Software Engineering, Lecture Notes in Artificial Intelligence*. Berlin: Springer Verlag.
- Hornby G.S. & Pollack J.B. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial Life*, 8:3.
- Langton C.G. (1984). Self-reproduction in cellular automata. *Physica D*, 10:134-144.
- Lohn J.D. & Reggia J.A. (1997). Automatic discoveries of self-replicating structure in cellular automata. *IEEE Transactions on Evolutionary Computation*, 1 (3): 165-178.
- Martinoli A. (1999). Swarm intelligence in autonomous collective robotics: from tools to the analysis and synthesis of distributed control strategies. *PhD Thesis*. Lausanne: Computer Science Department, EPFL.
- Mytilinaios E., Marcus D., Desnoyer H. & Lipson H. (in press) Designed and Evolved Blueprints For Physical Self-Replicating Machines. *Proceedings of the Ninth International Conference on Artificial Life (ALIFE IX)*.
- Moore E.F. (1970). Machine models of self-reproduction. In Burks A.W. (Ed.) *Essays on Cellular Automata*. Urbana, IL: University of Illinois Press.
- Mondada F., Pettinaro G., Guignard A., Kwee I., Floreano D., Denebourg J-L, Nolfi S., Gambardella L.M. & Dorigo M. (2004). Swarm-bot: A new distributed robotic concept. *Autonomous Robots*, 17(2-3): 193-221.
- Nolfi S. (1997). Evolving non-trivial behaviors on real robots: a garbage collecting robot. *Robotics and Autonomous System*, 22: 187-198
- Nolfi S. & Floreano D. (1999). Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution? *Artificial Life*, 4 (4), 311-335
- Nolfi S. & Floreano D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books.
- Pfeifer R. & Scheier C. (1999). *Understanding Intelligence*. Cambridge, MA:MIT Press
- Quinn M., Smith L., Mayley G. & Husbands P. (2002). Evolving teamwork and role allocation with real robots. In R.K. Standish, M.A. Bedau & H.A. Abbass (Eds.), *Proceedings of the 8th International Conference on Artificial Life*. Cambridge, MA: MIT Press.
- Rosin C.D. & Belew R.D. (1997). New methods for competitive coevolution. *Evolutionary Computation*, 5 (1), 1-29.
- Sims K. (1995). Evolving 3D morphology and behavior by competition. *Artificial Life*, 1:353-372.
- Sipper M. & Reggia J.A. (2001). Go forth and replicate. *Scientific American*, 2:35-43.
- Varela F.J. (1997). Patterns of life: Intwining identity and cognition. *Brain and Cognition*, 34:72-87.
- Von Neumann J. (1966). *Theory of Self-reproducing Automata* (edited and completed by A. W. Burks). Urbana, IL: University of Illinois Press.
- Taylor T. (2001). Creativity in evolution: individuals, interactions and environments. In P.J. Bentley & D. W Corne (Eds.) *Creative Evolutionary Systems*, Morgan Kaufman.
- Watson R.A., Figici S.G. & Pollack J.B. (2002). Embodied evolution: distributing an evolutionary algorithm in a population of robots. Robotics and population of robots. *Robotics and Autonomous Systems*, 39:1-18.