

Evolution of Embodied Intelligence

Dario Floreano, Francesco Mondada, Andres Perez-Uribe, Daniel Roggen

Autonomous Systems Laboratory (ASL)
Institute of Systems Engineering (I2S)
Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland
<http://asl.epfl.ch>
dario.floreano@epfl.ch

Introduction

For hundreds of years mankind has been fascinated with machines that display life-like appearance and behaviour. The early robots of the 19th century were anthropomorphic mechanical devices composed of gears and springs that would precisely repeat a pre-determined sequence of movements. Although a dramatic improvement in robotics took place during the 20th century with the development of electronics, computer technology, and artificial sensors, most of today robots used in factory floors are not significantly different from ancient automatic devices because they are still programmed to precisely execute a pre-defined series of actions. Are these machines intelligent? In our opinion they are not; they simply reflect the intelligence of the engineers that designed and programmed them. In the early 90's, we and other researchers started to address this issue by letting robots evolve, self-organise, and adapt to their environment in order to survive and reproduce, just like all life forms on Earth have done and keep doing. The name *Evolutionary Robotics* was coined to define the collective effort of engineers, biologists, and cognitive scientists to develop artificial robotic life forms that display the ability to evolve and adapt autonomously to their environment.

Within this perspective, artificial intelligence is a continuous and open-ended process that capitalises on physical interactions between the agent and its environment without human intervention. Embodiment does not only provide realism and semantic grounding to intelligent artefacts. It also provides opportunities that are unconceivable for bodiless systems. Embodied systems can tap upon a virtually infinite range of sensory cues and actions available in the physical world. Given the limits of their processing and behavioural abilities, they can be opportunistic and select only those sensory cues and actions that are necessary to carry on with the business of survival and reproduction. For example, ants can build magnificent nests with differentiated space, climate control, and air filtering. They do so without resorting to a plan, but by executing an evolved set of simple, but highly specific, sensory-triggered actions. In embodied systems, computation, representation, and memory can be partially outsourced to the physical laws and material persistence of the world. Consider for example the task of goal-directed navigation. One option to achieve that task is to build, store, and use an internal model of the entire environment. Another option is to select

and associate simple sensory cues with sequences of actions that will lead from one cue to the next until the goal is reached. Whereas man-made intelligent systems tend to use the first option, there is mounting evidence that animals (at least simple ones) exploit the second option.

In this chapter we will give an overview of some milestone experiments in evolution of physical robots and describe some examples of the intelligence that these robots develop.

Evolutionary Robotics

The possibility of evolving artificial creatures through an evolutionary process had already been evoked in 1984 by neurophysiologist Valentino Braitenberg in his truly inspiring booklet “Vehicles. Experiments in Synthetic Psychology”. Braitenberg proposed a thought-experiment where one builds a number of simple wheeled robots with different sensors variously connected through electrical wires and other electronic paraphernalia to the motors driving the wheels. When these robots are put on the surface of a table, they will begin to display behaviours such as going straight, approaching light sources, pausing for some time and then rushing away, etc. Of course, some of these robots will fall off the table. All one needs to do is continuously pick a robot from the tabletop, build another robot just like one on the table, and add the new robot to the tabletop. If one wants to maintain a constant number of robots on the table, it is necessary to copy-build one robot for every robot that falls from the table. During the process of building a copy of robots, one will inevitably make some small mistakes, such as inverting the polarity of an electrical connection or using a different resistance. Those mistaken copies that are lucky enough to remain longer on the tabletop will have a high number of descendants, whereas those that fall off the table will disappear for ever from the population. Furthermore, some of the mistaken copies may display new behaviours and have higher chance of remaining for very long time on the tabletop. You will by now realise that the creation of new designs and improvements through a process of selective copy with random errors without the effort of a conscious designer was already proposed by Darwin to explain the evolution of biological life on Earth.

However, the dominant view by mainstream engineers that robots were mathematical machines designed and programmed for precise tasks, along with the technology available at that time, delayed the realisation of the first experiments in Evolutionary Robotics for almost ten years. In the spring of 1994 our team at EPFL (Swiss Federal Institute of Technology in Lausanne) [8] and a team at the University of Sussex in Brighton [15] reported the first successful cases where robots evolved with minimal human intervention and developed neural circuits allowing them to autonomously move in real environments. The two teams were driven by similar motivations. On the one hand, we felt that a designer approach to robotics was inadequate to cope with the complexity of the interactions between the robot and its physical environment as well as with the control circuitry required for such interactions. Therefore, we decided to tackle the problem by letting these complex interactions guide the evolutionary development of robot brains subjected to certain selection criteria (tech-

nically known as fitness functions), instead of attempting to formalise the interactions and then designing the robot brains. On the other hand, we thought that by letting robots autonomously interact with the environment, evolution would exploit the complexities of the physical interactions to develop much simpler neural circuits than those typically conceived by engineers who use formal analysis methods. We had plenty of examples from nature where simple neural circuits were responsible for apparently very complex behaviours. Ultimately, we thought that Evolutionary Robotics would not only discover new forms of autonomous intelligence, but also generate solutions and circuits that could be used by biologists as guiding hypotheses to understand adaptive behaviours and neural circuits found in nature.

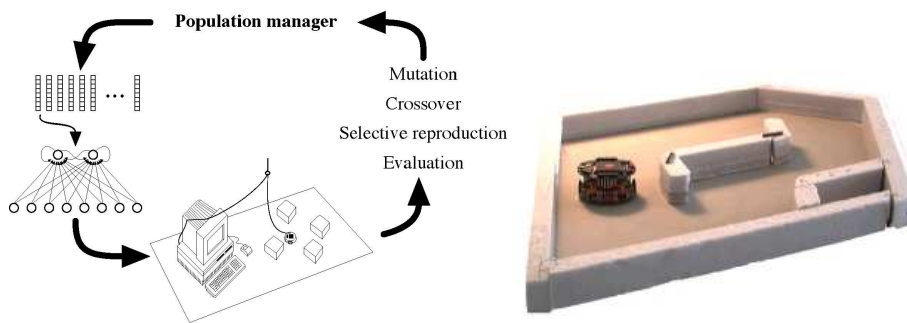


Fig. 1. Left: Artificial Evolution of neural circuits for a robot connected to a computer. Right: The miniature mobile robot *Khepera* in the looping maze used during an evolutionary experiment.

In order to carry out evolutionary experiments without human intervention, at EPFL we developed the miniature mobile robot *Khepera* [25] (6 cm of diameter for 70 grams) with eight simple light sensors distributed around its circular body (6 on one side and 2 on the other side) and two wheels (figure 1). Given its small size, the robot could be attached to a computer through a cable hanging from the ceiling and specially designed rotating contacts in order to continuously power the robot and let the computer keep a record of all its movements and neural circuit shapes during the evolutionary process, a sort of fossil record for later analysis. The computer generated an initial population of random artificial chromosomes composed of 0's and 1's that represented the properties of an artificial neural network. Each chromosome was then decoded, one at a time, into the corresponding neural network whose input neurons were attached to the sensors of the robots and the output unit activations were used to set the speeds of the wheels. The decoded neural circuit was tested on the robot for some minutes while the computer evaluated its performance (fitness). In these experiments, we wished to evolve the ability to move straight and avoid obstacles. Therefore, we instructed the computer to select for reproduction those individuals whose two wheels moved on the same direction (straight motion) and whose sensors had lower activation (far from obstacles). Once all the chromosomes of the population had been tested on the same physical robot, the chromosomes of selected individuals were organised in pairs and parts of their genes were exchanged (crossover) with small random errors (mutations) in order to generate a number of offspring. These

offspring formed a new generation that was again tested and reproduced several times. After 50 generations (corresponding to approximately two days of continuous operation), we found a robot capable of performing complete laps around the maze without ever hitting obstacles. The evolved circuit was rather simple, but still more complex than hand-designed circuits for similar behaviours because it exploited non-linear feedback connections among motor neurons in order to get away from some corners. Furthermore, the robot always moved in the direction corresponding to the higher number of sensors. Although the robot was perfectly circular and could move in both directions in the early generations, those individuals moving in the direction with fewer sensors tended to remain stuck in some corners because they could not perceive them properly, and thus disappeared from the population. This represented a first case of adaptation of neural circuits to the body shape of the robot in a specific environment.

The Sussex team instead developed a *Gantry robot* consisting of a suspended camera that could move in a small box along the x and y coordinates and also rotate on itself [15]. The image from the camera was fed into a computer and some of its pixels were used as input to an evolutionary neural circuit whose output was used to move the camera. The artificial chromosomes encoded both the architecture of the neural network and the size and position of the pixel groups used as input to the network. The team used a form of incremental evolution whereby the gantry robot was first evolved in a box with one painted wall and asked to go towards the wall. Then, the size of the painted area was reduced to a rectangle and the robot was incrementally evolved to go towards the rectangle. Finally, a triangle was put nearby the rectangle and the robot was asked to go towards the rectangle, but avoid the triangle. A remarkable result of these experiments was that evolved individuals used only two groups of pixels to recognise the shapes by moving the camera from right to left and using the time of pixel activation as an indicator of the shape being faced (for the triangle, both groups of pixels become active at the same time, whereas for the rectangle the top group of pixels becomes active before the lower group). This was compelling evidence that evolution could exploit the interaction between the robot and its environment to develop smart simple mechanisms that could solve apparently complex tasks.

The next question was whether more complex cognitive skills could be evolved by simply exposing robots to more challenging environments. To test this hypothesis, at EPFL we put the Khepera robot in an arena with a battery charger in one corner under a light source (figure 2) and let the robot move around as long as its batteries were discharged [9]. To accelerate the evolutionary process, the batteries were simulated and lasted only 20 seconds; the battery charger was a black painted area of the arena and when the robot happened to pass over it, the batteries were immediately recharged. The fitness criterion was the same used for the experiment on evolution of straight navigation (figure 1), that is keep moving as much as possible while staying away from obstacles. Those robots that managed to find the battery charger (initially by chance) could live longer and thus accumulate more fitness points. After 240 generations, that is 1 week of continuous operation, we found a robot that was capable of moving around the arena, go towards the charging station only 2 seconds before the battery was fully discharged, and then immediately returning in the open arena. The robot did not simply sit on the charging area because it was too close to the walls and its fitness was very low (remember from the previous experiment that robots had

higher fitness when its proximity sensors had lower activation). When we analysed the activity of the evolved neural circuit while the robot was freely moving in the arena, we discovered that the activation of one neuron depended on the position and orientation of the robot in the environment, but not on the level of battery charge (figure 2). In other words, this neuron encoded a spatial representation of the environment (sometimes referred to as “cognitive map” by psychologists), computationally similar to some neurons that neurophysiologists discovered in the hippocampus of rats exploring an environment.

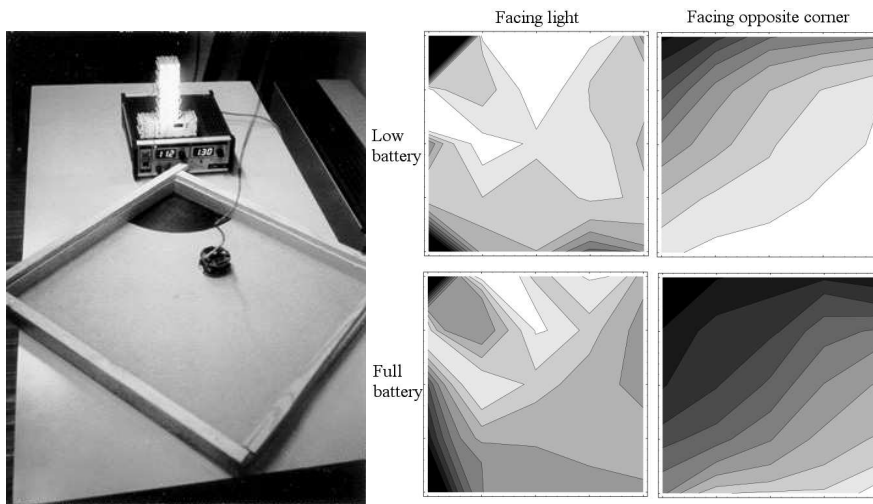


Fig. 2. Left: A Khepera robot is positioned in an arena with a simulated battery charger (the black-painted area on the floor). The light tower above the recharging station is the only source of illumination. Right: Activity levels of one neuron of the evolved individual. Each box shows the activity of the neuron (white = very active, black = inactive) while the robot moves in the arena (the recharging area is on the top left corner). The activity of the neuron reflects the orientation of the robot and its position in the environment, but is not affected by the level of battery charge.

Competitive Co-evolution

Encouraged by these experiments, we decided to make the environment even more challenging by co-evolving two robots in competition with each other. The Sussex team had begun investigating co-evolution of predator and prey agents in simulation to see whether increasingly more complex forms of intelligence emerged in the two species [24]. They showed that the evolutionary process changed dramatically when two populations co-evolved in competition with each other because the performance of each robot depends on the performance of the other robot. In the Sussex experiments the fitness of the prey species was proportional to the distance from the preda-

tor whereas the fitness of the predator species was inversely proportional to the distance from the prey. Although in some evolutionary runs they observed interesting pursuit-escape behaviours, often co-evolution did not produce interesting result.

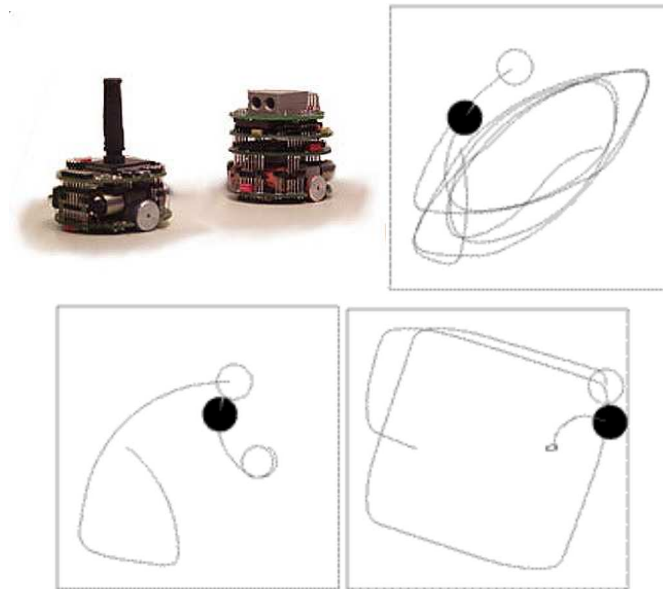


Fig. 3. Co-evolutionary prey (left) and predator (right) robots. Trajectories of the two robots (prey is white, predator is black) after 20, 45, and 70 generations.

At EPFL we wanted to use physical robots with different hardware for the two species and give them more freedom to evolve suitable strategies by using as fitness function the time of collision instead of the distance between the two competitors [10]. In other words we did not explicitly select predator robots for getting closer to the prey and prey robots for keeping at a distance from predators, but we let them choose the most suitable strategies to succeed the ultimate survival criterion: catch the prey and avoid the predator, respectively. We created a predator robot with a vision system spanning 36 degrees and a prey robot that had only simple sensors capable of detecting an object at 2 cm of distance, but that could move twice as fast as the predator (figure 3). These robots were co-evolved in a square arena and each pair of predator and prey robots were let free to move for 2 minutes (or less if the predator could catch the prey). The results were quite surprising. After 20 generations, the predators developed the ability to search for the prey and follow it while the prey escaped moving all around the arena. However, since the prey could go faster than the predator, this strategy did not always pay off for predators. 25 generations later we noticed that predators watched the prey from far and eventually attacked it anticipating its trajectory. As a consequence, the prey began to move so fast along the walls that often predators missed it and crashed into the wall. Again, 25 generations later we discovered that predators developed a “spider strategy”. Instead of attempting to go after the prey, they quietly moved towards a wall and waited there for the prey to

arrive. The prey moved so fast near the walls that it could not detect the predator early enough to avoid it!

However, when we let the two robot species co-evolve for more generations, we realised that the two species rediscovered older strategies that were effective against the current strategies used by the opponent. This was not surprising. Considering the simplicity of the environment, the number of possible strategies that can be effectively used by the two robot species is limited. Even in nature, there is evidence that co-evolutionary hosts and parasites (for example plants and insects) recycle old strategies over generations.

Stefano Nolfi, who worked with us on these experiments, noticed that by making the environment more complex (for example with the addition of objects in the arena) the variety of evolved strategies was much higher and it took much longer before the two species re-used earlier strategies [26]. We also noticed that the competing selection pressure on the two species generated much faster evolution and behavioural change than in robots evolved in isolation under an externally defined fitness function. These experiments never stopped surprising us and indeed turned out to be a source of inspiration for the best-selling novelist Michael Crichton in his latest science fiction book *Prey* [6]. We feel that this area of research has still much to deliver for the bootstrapping of machine intelligence.

Cooperative Co-evolution

Beside competition, living organisms display a sort of "collective intelligence", characterised by complex levels of cooperation that provide them with higher evolutionary advantage. For instance, it has been estimated that one-third of the animal biomass of the Amazon rain forest consists of social insects, like ants and termites [17]. The success of social insects might come from the fact that social interactions can compensate for limitations of the individual, both in terms of physical and cognitive capabilities.

A social insect colony is a complex system often characterised by division of labour and high genetic similarity among individuals [37]. Ants, bees, wasps, and termites provide some of the most remarkable examples of altruist behaviour with their worker caste, whose individuals forego their own reproduction to enhance reproduction of the queen. These and other examples of group harmony and cooperation show the colony as if it behaved as a "superorganism" where individual-level selection is muted, with the result that colony-level selection reigns.

Biologists agree that relatedness plays a major role in favouring the evolution of cooperation in social insects [19]. However, the concept of the colony as a superorganism has been challenged [19]. In collaboration with ant biologist Laurent Keller and robot designer Roland Siegwart, we are trying to determine whether the role of relatedness and the level of selection can be experimentally demonstrated using colonies of artificial ants implemented as small mobile robots with simple vision and communication abilities (figure 4). For this purpose, we have defined experimental settings where these robotic ants are supposed to look for food items randomly scattered in a foraging area. The robots are provided with artificial genomes that code for their behaviours in an indirect manner (i.e., the patterns of behaviour activation coded

by the same genetic code vary according to the phenotype frequencies in the colony). There are two kinds of food items. Small food items can be transported by single robots to the nest. Large food items require two cooperating robots to be pushed away. By varying the energetic value of the food items, we can put more or less pressure on the advantage of cooperative behaviours.




Fig. 4. Left: The sugarcube robot Alice equipped with vision system, distance sensors, communication sensors, and two frontal „mandibles“ to better grasp objects. Right: The arena with small and large objects. The nest is under the textured wall where a small gap let objects –but not robots– fall on the floor.

In a first set of experiments carried out in simulation, we investigated how colony performance evolved under different levels of selection (individual and colony level) and under high versus low “genetic” relatedness between robots of the same colony. We ran experiments using a minimalist simulator of the collective robotics evolutionary setup [28], and found that “genetically” homogeneous colonies of foraging simulated robots performed better than heterogeneous ones. Moreover, our experiments showed that altruistic behaviours have low probability of emerging in heterogeneous colonies evolving under individual-level selection. Our current work is aimed at running these experiments in colonies of 20 sugar-cube robots in order to better study the role of physical interactions.

Physical Interactions

Collaboration among animals can also take place at a pure physical level. For instance, a mother can help her kids by pushing, pulling, or transporting them on the back. Human acrobats can build towers with their bodies, ants can build bridges, rafts, pulling chains or doors, and bees can build curtains or balls, for instance. In all these

examples the group of individuals can achieve a task impossible for a single individual by dynamically aggregating into different and functional physical structures. To investigate this new research direction, in collaboration with other European partners [29], we are developing a new robotic concept, called *s-bot*, capable of physically interconnecting to other *s-bots* to form a *swarm-bot* (<http://www.swarm-bots.org>). Each *s-bot* is a fully autonomous mobile robot capable of performing basic tasks such as autonomous navigation, perception of the environment and grasping of objects (figure 5). Ants can lift each other and heavy objects with their mandibles and can establish flexible connections between each other with their legs. Similarly, each *s-bot* is equipped with a strong beak gripper that can lift heavy objects or another *s-bot* and with a flexible gripper that can grasp another *s-bot* on the belt to maintain physical contact. *S-bots* can organise in *swarm-bot* configuration by dynamically attaching to each other and form various shapes according to environmental constraints or task needs.



Fig. 5. Left: The prototype of the *s-bot* with the strong beak gripper and the flexible arm. Right: Several *s-bots* can self-connect to build a *swarm-bot* capable of passing obstacles one single *s-bot* cannot deal with.

In addition to these features, an *s-bot* is capable of communicating with other *s-bots* by emitting and receiving sounds. *S-bots* can also use body signals by changing the colour of their body belts to display their internal states. Other *s-bots*, with their vision system, can see this corporal expression and react, for instance helping the “red” robot, following the “blue” one, or connecting to the “green” one to form a *swarm-bot* configuration. Assembled in *swarm-bot* configuration, the robots are able to perform exploration, navigation and transport of heavy objects in very rough terrain, where a single *s-bot* could not possibly achieve the task.

The control of this hardware structure is very challenging and has implications on the whole design, from mechanics to software. In this project we resort to a combination of artificial evolution, behaviours inspired from the world of social insects, and standard engineering methodologies. Standard engineering methodologies are applied in all local sub-problems where classical approaches are well known, reliable and form a basic structure on top of which we can build the collective control. This is for instance the case of mechanical design, low-level motor control, sensor management (not processing), and low level communication procedures. Bio-inspired solutions are applied where natural mechanisms are well identified and can be translated into our

robot design and control. Examples of bio-inspired design elements are the shape of the grippers and the interactive synchronisation of the robots when grasping an object. Another bio-inspired element is the general concept to solve complex tasks with the combination of many simple mechanisms. On the top of these two approaches we apply artificial evolution to exploit in the best way the specific properties of each part for a given behaviour.

Artificial evolution generated a set of simple rules capable of coordinating the movement of a group of connected s-bots [1]. In this particular case, evolution exploited the property of a force sensor within the body of each s-bot to integrate the behaviour of the whole group without need of external communication or additional coordination layers. These results indicate that physical interactions alone can provide useful information for coordination. Still, it is the responsibility of the engineer to provide sensors and actuators that can be handled efficiently by evolution. This illustrates a big difference with respect to natural evolution, where the behaviours and the body of organisms co-evolve.

Active Vision and Feature Selection

Brains are characterised by limited bandwidth and computational resources. At any point in time, we can focus our attention only to a limited set of features or objects. One of the most remarkable —and often neglected— differences between machine vision and biological vision is that computers are often asked to process an entire image in one shot and produce an immediate answer whereas animals are free to explore the image over time searching for features and dynamically integrating information over time.

We thought that the computational complexity of vision-based behaviour could be greatly simplified if the processes of active vision and of feature selection are co-evolved while the robot interacts with the environment. Each of these two processes has been investigated and adopted in machine vision. Active vision is the sequential and interactive process of selecting and analysing parts of a visual scene [2]. Feature selection instead is the development of sensitivity to relevant features in the visual scene to which the system selectively responds [14]. However, the combination of active vision and feature selection is still largely unexplored.

To investigate that hypothesis, we devised a very simple neural architecture composed of only one layer of synaptic connections (figure 6, left) that link visual neurons to two sets of motor outputs. One set of output units controls the behaviour of the system (for example, the movements of a robot or the categorisation of an image discrimination system). The other set controls the behaviour of the vision system (movement over the visual field, zooming factor, pre-filtering strategy). The synaptic weights, which are genetically encoded and evolved using a simple genetic algorithm, are responsible both for the visual features to which the system responds to and for the actions of the vision system.

We carried out a series of experiments on co-evolution of active vision and feature selection for behavioural systems equipped with primitive retinal systems and deliberately simple neural architectures [7]. In a first set of experiments, we show that sen-

sitivity to very simple features is co-evolved with, and exploited by, active vision to perform complex shape discrimination [18]. We also show that such discrimination problem is very difficult for a similar vision system without active behaviour because the architecture must solve non-linear transformations (position and size invariance) of the image in order to solve the task. Instead, the co-evolved active vision and feature selection system rely on linear transformations of parts of the image (oriented edges and corners), which are actively searched and sequentially scanned in order to provide the correct answer. In a second set of experiments, we applied the same co-evolutionary method and architecture for driving a simulated car over roads in the Swiss Alps and show that active vision is exploited to locate and fixate the edge of the road while driving the car. In a third set of experiments, we used once again the same co-evolutionary method and architecture for an autonomous robot equipped with a pan/tilt camera (figure 6, right) that is asked to navigate in an arena located in an office environment [22]. Evolved robots exploit active vision and simple features to direct their gaze at invariant parts of the environment (horizontal edge between the floor and furniture) and perform collision-free navigation. In a fourth set of experiments, we apply this methodology to an all-terrain robot with a static, but large, field of view that must navigate in a rugged terrain. Here again, the system becomes sensitive to a set of simple visual features that are maintained within the retina by the active vision mechanisms.

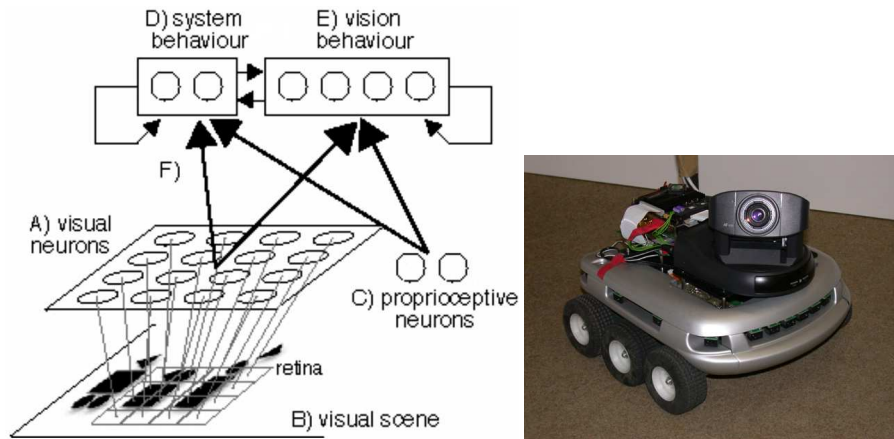


Fig. 6. Left: Architecture of the control system. The architecture is composed of A) a grid of visual neurons with non-overlapping receptive fields whose activation is given by B) the grey level of the corresponding pixels in the image; C) a set of proprioceptive neurons that provide information about the movement of the vision system; D) a set of output neurons that determine the behaviour of the system (pattern recognition, car driving, robot navigation); E) a set of output neurons that determine the behaviour of the vision system; F) a set of evolvable synaptic connections. The number of neurons in each sub-system can vary according to the experimental settings. Right: The Koala robot equipped with a mobile camera whose image is fed into the vision neurons of the neural architecture.

Evolution of Learning

Another interesting direction in Evolutionary Robotics is the evolution of learning. In a broad sense, learning is the ability to adapt during lifetime and we know that most living organisms with a nervous system display some type of adaptation during life. The ability to adapt quickly is crucial for autonomous robots that operate in dynamic and partially unpredictable environments, but the learning systems developed so far have many constraints that make them hardly applicable to robots interacting with an environment without human intervention. Of course, evolution is also a form of adaptation, but modifications occur only over several generations, and that may require too long time for a robotic system (for a comparative discussion of lifelong learning and evolution, see [27]). In order to compensate for the problems of both approaches, we decided to genetically encode and evolve the mechanisms of neural adaptation [11]. The idea was to exploit evolution to find good combinations of learning structures, rather than static controllers, and to evolve learning structures that without the constraints of off-the-shelf learning algorithms. The artificial chromosomes encoded a set of rules that were used to change the synaptic connections among the neurons while the robot moved in the environment. The results were very interesting.

A Khepera robot equipped with a vision system was put in an arena with a light bulb and a light switch (figure 7). The light switch is marked by a black stripe painted on the wall. The fitness is given by the amount of time spent by the robot under the light bulb when the light is on. Initially the light is off. Therefore, the robot must first go towards the black stripe to switch the light on (notice that the fitness function does not explicitly encourage this behaviour). The black and grey areas on the floor are used by the computer to detect through a sensor positioned under the robot when to switch the light on and when to accumulate fitness points, but this information is not given to the evolutionary controller. Evolved robot learned during their lifetime the sequence of behaviours necessary to increase their fitness. These included: wall avoidance, movement towards the stripe, movement towards the light, and resting under the light.

Not only the evolution of learning rules resulted in more complex skills, such as the ability to solve sequential tasks that simple insects cannot solve, but also the number of generations required was much smaller. However, the most important result was that evolved robots were capable of adapting during their lifetime to several types of environmental change that were never seen during the evolutionary process, such as different light conditions, environmental layouts, and even a different robotic body. Very recently, Akio Ishiguro and his team at the University of Nagoya used a similar approach for a simulated humanoid robot and showed that the evolved nervous system was capable of adapting the walking style to different terrain conditions that were never presented during evolution [13]. The learning abilities that these evolved robots display are still very simple, but current research is aimed at understanding under which conditions more complex learning skills could evolve in autonomous evolutionary robots.

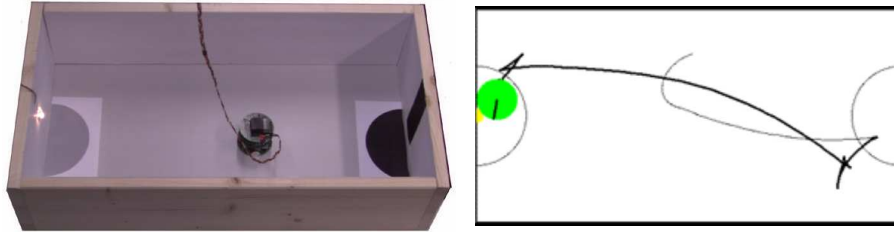


Fig. 7. Left: A Khepera robot with a vision system is positioned in an arena with a light bulb and a light switch (black stripe on the wall). At the beginning of the robot life, the light bulb is off. The robot must develop from random synaptic connection using genetically determined learning rules how to switch the light on and stay under the light bulb. Right: Trajectory of an evolved robot with enabled synaptic adaptation.

Evolvable Hardware

In the experiment described so far, the evolutionary process operated on the features of the software that controlled the robot (in most cases, in the form of an artificial neural network). The distinction between software and hardware is quite arbitrary and in fact one could build a variety of electronic circuits that display interesting behaviours without any software. A few years ago, some researchers realised that the methods used by electronic engineers to build circuits represent only a minor part of all possible circuits that could be built out of a given number of components. Furthermore, electronic engineers tend to avoid circuits that display complex and highly non-linear dynamics, and more in general those which are hard-to-predict, which may be just the type of circuits that a behavioural machine requires. Adrian Thompson at the University of Sussex suggested the evolution of electronic circuits without imposing any design constraints [35]. Thompson used a type of electronic circuit, known as Field Programmable Gate Array (FPGA), whose internal wiring can be entirely modified in a few nanoseconds. Since the circuit configuration is a chain of 0's and 1's, he used this chain as the chromosome of the circuit and let it evolve for a variety of tasks, such as sound discrimination and even robot control. Some evolved circuits used 100 times less components than circuits conceived for similar tasks with conventional electronic design, and displayed novel types of wiring. Interestingly, evolved circuits were sensitive to environmental features, such as temperature, which is usually a drawback in electronic design practice, but is a common feature of all living organisms.

The field of Evolutionary Electronics was born and these days several researchers around the world use artificial evolution to discover new types of circuits or let circuits evolve to new operating conditions. For example, Adrian Stoica and his colleagues at NASA/JPL are designing evolvable circuits for robotics and space application [34], while Tetsuya Higuchi, another pioneer of this field, at the Electro-

Technical Laboratory near Tokyo in Japan is already bringing to the market mobile phones and prosthetic implants with evolvable circuits [16].

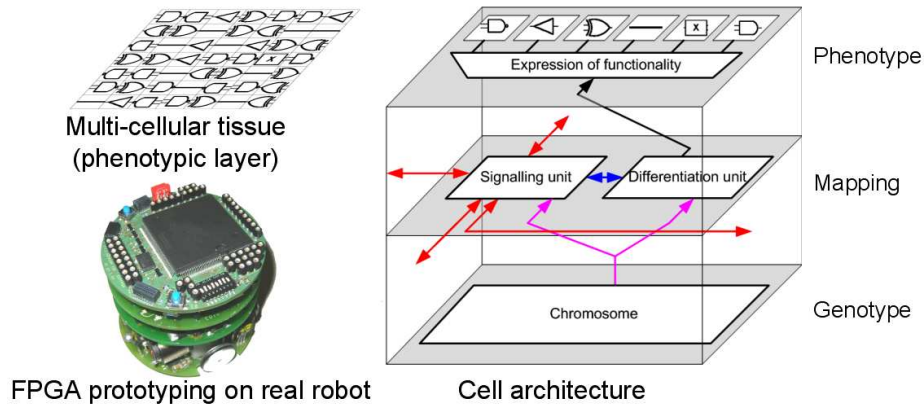


Fig. 8. A schematic representation of the electronic tissue. Each cell of the tissue is composed of three layers, a genotype layer to store the artificial genome of the entire tissue, a phenotype layer to express the functionality of the cell, and an intervening mapping layer to dynamically express the genes into functionalities according to gene expression and cell signalling processes. In addition, each cell of the circuit has input and output connections with the environment. Cells can be dynamically added or removed from the circuit at runtime. A prototype of the electronic tissue has been added on top of the Khepera robot and evolved to generate tissues of spiking neural controllers.

At EPFL, in collaboration with other European partners [36], we are pushing even further the analogy between silicon devices and biological cells in the attempt to create an electronic tissue capable of evolution, self-organisation, and self-repair (<http://www.poetic tissue.org>). The electronic tissue is multi-cellular surface composed of several tiny re-configurable electronic circuits that can be attached or detached while the tissue is in operation. Similarly to a biological cell, each electronic cell is composed of three layers (figure 8). The genotype layer stores the artificial genome of the entire tissue. The phenotype layer expresses the functionality of the cell such as a neuron, a hair cell, a photoreceptor, a motor cell, etc. Finally, the mapping layer regulates the gene expression mechanisms depending on inter-cellular electronic signals. In addition, each electronic cell or group of cells can be attached to a sensor (a phototransistor, a whisker, a microphone, etc.) and/or to an actuator (a servomotor or an artificial muscle). An artificial genome is sent to a mother cell that sends it to all available cells, mimicking a process of cell duplication. As a cell receives a genome, a process of gene expression starts. The gene expression mechanism is affected by intercellular signals so that the functional property expressed by a cell partially depends on the type and intensity of received signals, on its position in the tissue, on the time of genome reception, and on environmental stimulation. For example, cells connected to photoreceptors may have a higher likelihood to process photons. Early prototypes of the system have been interfaced to a robot by connecting the sensors and actuators

to cells. The tissue has been subjected to an evolutionary process where different genomes are sequentially tested, reproduced, crossed over and mutated until the robot displayed suitable navigation in a maze [32].

Evolutionary Morphologies

In the early experiments on evolution of navigation and obstacle avoidance (figure 1), the neural circuits adapted over generations to the distribution of sensors of the Khepera robot. However, in Nature also the body shape and sensory-motor configuration is subjected to an evolutionary process. Therefore, one may imagine a situation where the sensor distribution of the robot must adapt to a fixed and relatively simple neural circuit. The team of Rolf Pfeifer at the A.I. laboratory in Zurich developed Eyebot, a robot with an evolvable eye configuration, to study the interaction between morphology and computation for autonomous robots [20]. The vision system of Eyebot is similar to that of houseflies and is composed of several directional light receptors whose angle can be adjusted by motors. The authors evolved the relative position of the light sensors while using a simple and fixed neural circuit in a situation where the robot was asked to estimate distance from an obstacle while moving along a track. The experimental results confirmed the theoretical predictions: The evolved distribution of the light receptors displayed higher density of receptors toward in the frontal direction than on the sides of the robot. The messages of this experiment are quite important: on the one hand the body shape plays an important role in the behaviour of an autonomous system and should be co-evolved with other aspects of the robot; on the other hand, computational complexity can be traded with a morphology adapted to the environment.

Back in 1997, when quadruped robots were still an affair of research laboratories, we used a co-evolutionary approach to investigate the balance between morphology and control of a four leg robot [12] (figure 9). More specifically, we were interested in finding a good ratio between leg and body size as well as minimise the number of motorised degrees of freedom provided by a behaviour-based control system with a number of evolvable parameters. We carried out co-evolution of body and control in 3D simulations, but constrained the genetic representation of the robot morphology to a number of primitives that could be built using available technology. Evolved robots were capable of walking forward and turning very smoothly to avoid obstacles using an infrared sensor positioned in front of the robot. These robots used rotating joints only on the front legs. We then built a physical robot according to the dimensions found by the co-evolutionary process (figure 9, right) and downloaded the evolved control system for autonomous navigation. The physical robot displayed the same walking behaviour shown in simulation, although it had a noticeable trembling (which looked as if it was affected by the mad-cow disease) caused by the differences between simulations and physical reality. Since our purpose was to study the interactions between body and control co-evolution, we did not attempt to improve the walking behaviour of the physical robot. However, a possible strategy would be to evolve the learning rules (as described in a section above) and have the “newborn” physical robot adapt online to its own physical characteristics. Also adding some

noise to the sensors and actuator while simulating the robot may help bridge the gap to reality [23] by avoiding that the controller over-specialises to the simulation.



Fig. 9. An evolved 4 legged robot. The control system of the robot, its body size, and length of legs have been evolved in 3D simulations (left). The physical robot (right) has been built according to the evolved genetic specifications. The evolved control system is transferred from the simulated to the physical robot. Such evolved robot can walk and avoid obstacles. The robot is approximately 20 cm long and less than 1kg without batteries. Leg control performed by a set of HC11 microcontrollers.

Co-evolution of the body and controller has also been applied to biped robots [4]. The results showed better walking characteristics than when only the controller was evolved. The idea of co-evolving the body and the neural circuit of autonomous robots had already been investigated in simulations by Karl Sims [33], but only recently this has been achieved in hardware. Jordan Pollack and his team at Brandeis University have co-evolved the body shape and the neurons controlling the motors of robots composed of variable-length sticks whose fitness criterion is to move forward as far as possible [21]. The chromosomes of these robots include specifications for a 3D printer that builds the bodies out of thermoplastic material. These bodies are then fitted with motors and let free to move while their fitness is measured. Artificial evolution generated quite innovative body shapes that resemble biological morphologies such as those of fishes.

A Look Ahead

Over the last 10 years, the role of embodiment and behavioural interaction has been increasingly recognised as a cornerstone of natural and artificial intelligence. New research initiatives in information technologies, neuroscience, and cognitive science sponsored by the European Commission, U.S. National Science Foundation, and a number of national programs explicitly emphasise these two aspects.

Many more examples of evolutionary robots exhibiting intelligent behaviours are available out there, too many to be covered in this short document. However, we are just scratching the surface of a radically new way of understanding how intelligent life emerged on this planet and could evolve in machines. There are a number of conceptual and technological challenges ahead. For example, evolution does not automatically lead to intelligent behaviours. A lot of prior knowledge and experience is

still required to select appropriate parameters, such as the genetic encoding, the neural network architecture, the mapping of sensors and actuator to the network or even the fitness function. Developing better methodologies to select those parameters is an important aspect that needs to be tackled for evolving more complex systems. Also, we are facing what is called the “bootstrap problem”. If the environment or the fitness function is too harsh for the evolving individual during the initial generations (so that all the individuals of the first generation have zero fitness), evolution cannot select good individuals and make any progress. A possible solution (and by far not the only one) is to start with environments and fitness functions that become increasingly more complex over time. However, this means that we must put more effort in developing methods for performing incremental evolution that, to some extent, preserve and capitalise upon previously discovered solutions. In turn, this implies that we should understand what are suitable primitives and genetic encoding upon which artificial evolution can generate more complex structures. A key aspect will most likely be the emergence of modular and hierarchical structures through mechanisms of genetic regulatory networks, cell differentiation, and inter-cellular signalling. Another challenge is hardware technology. Despite the encouraging results obtained in the area of evolvable hardware, many of us feel that we should drastically reconsider the hardware upon which artificial evolution operates. This means that maybe we should put more effort in self-assembling materials that give less constraints to the evolving system, facilitate the evolutionary process, and may eventually lead to truly self-reproducing machines.

Acknowledgements

This paper is an updated and extended version of *A.I. 101: Machine Self-evolution* by the same authors. The experiments carried out at EPFL were possible thanks to the collaboration with Stefano Nolfi, Joseba Urzelai, Jean-Daniel Nicoud, Andre Guignard, Laurent Keller, Roland Siegwart, Eduardo Sanchez, Marco Dorigo, Luca Gambardella, and Jean-Luis Deneubourg. The authors thank the Swiss National Science Foundation (grant nr. 620-58049.99) and the Future Emergent Technologies division of the European Commission for continuous support of this exciting research field (Swiss OFES grants 00.05291-1 and 01.0012-1).

References

1. Baldassarre, G., Nolfi, S. and Parisi, D. (2002) Evolving Mobile Robots Able to Display Collective Behaviours. In C. K. Hemelrijk and E. Bonabeau, (eds.) *Proceedings of the International Workshop on Self-Organisation and Evolution of Social Behaviour*, 11-22.
2. Bajcsy, R. (1988) Active Perception. *Proceedings of the IEEE*, 76, 996-1005.
3. Ballard, D.H. (1991) Animate Vision. *Artificial Intelligence*, 48, 57-86.
4. Bongard, J.C. and Paul, C. (2001). Making Evolution an Offer It Can't Refuse: Morphology and the Extradimensional Bypass. In Keleman, J. and Sosik, P. (eds.) *Proc. of the 6th European Conf. on Artificial Life*, Prague, CZ, 401-412.

5. Braitenberg, V. (1984). *Vehicles. Experiments in Synthetic Psychology*. Cambridge, MA: MIT Press.
6. Crichton, M. (2002) *Prey*. New York: Harper Collins.
7. Floreano, D., Kato, T., Marocco, D. and Sauser, E. (2003) Co-evolution of Active Vision and Feature Selection. *Biological Cybernetics*. Submitted.
8. Floreano, D. and Mondada, F. (1994) Automatic Creation of an Autonomous Agent: Genetic Evolution of a Neural Network Driven Robot. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson (eds.), *From Animals to Animats III*, Cambridge, MA: MIT Press, 421-430.
9. Floreano, D. and 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.
10. Floreano, D., Nolfi, S., and Mondada, F. (2001) Co-Evolution and Ontogenetic Change in Competing Robots. In M. Patel, V. Honavar, and K. Balakrishnan (eds.), *Advances in the Evolutionary Synthesis of Intelligent Agents*, Cambridge (MA): MIT Press.
11. Floreano, D. and Urzelai, J. (2000a) Evolutionary Robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13, 431-443.
12. Floreano, D. and Urzelai, J. (2000b) Artificial Evolution of Adaptive Software: An Application to Autonomous Robots. *3D The Journal of Three dimensional images* (in japanese), 14(4), 64-69
13. Fujii, A., Ishiguro, A., Aoki, T. and Eggenberger, P. (2001) Evolving Bipedal Locomotion with a Dynamically-Rearranging Neural Network. In J. Kelemen and P. Sosik (eds.) *Advances in Artificial Life (ECAL 2001)*, Berlin: Springer Verlag, pp. 509-518.
14. Hancock, P.J., Baddeley, R.J. and Smith L.S. (1992) The principal components of natural images. *Network*, 3, 61-70.
15. Harvey, I., Husbands, Ph., and Cliff, D. (1994) Seeing the Light: Artificial Evolution, Real Vision. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson (eds.), *From Animals to Animats III*, Cambridge, MA: MIT Press, 392-401
16. Higuchi, T., Iwata, M., Keymeulen, D., Sakanashi, H., Murakawa, M., Kajitani, I., Takahashi, E., Toda, K., Salami, M., Kajihara, N. and Otsu, N. (1999) Real-World Applications of Analog and Digital Evolvable Hardware. *IEEE Transactions on Evolutionary Computation*, 3(3), 220-235.
17. Holldobler, B. and Wilson, E.O. (1990) *The Ants*. Berlin: Springer-Verlag.
18. Kato, T. and Floreano, D. (2001) An Evolutionary Active-Vision System. In *Proceedings of the Congress on Evolutionary Computation (CEC'01)*, Piscataway: IEEE Press.
19. Keller, L. and Reeve, H.K. (1999) Dynamics of conflicts within insect societies. In L. Keller (Ed.), *Levels of Selection in Evolution*, Princeton University Press.
20. Lichtensteiger, L. and Eggenberger, P. (1999). Evolving the morphology of a compound eye on a robot. In *Proceedings of the Third European Workshop on Advanced Mobile Robots (Eurobot '99)*. New York: IEEE Press, 127-134.
21. Lipson, H. and Pollack, J.B. (2000). Automatic design and manufacture of robotic life-forms. *Nature*, 406 (6799): 974-978.
22. Marocco, D. and Floreano, D. (2002) Active Vision and Feature Selection in Evolutionary Behavioral Systems. In Hallam, J. and Floreano, D. and Hayes, G. and Meyer, J. (eds.) *From Animals to Animats 7: Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, 247-255.
23. Miglino, O., Lund, H.H. and Nolfi, S. (1995). Evolving Mobile Robots in Simulated and Real Environments. *Artificial Life* 2(4), 417-434.
24. Miller, G.F. and Cliff, D. (1994). Protean Behavior in Dynamic Games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D. and Husbands, P. and Meyer, J. and Wilson, S. W. (eds.) *From Animals to Animats III*. Cambridge, MA: MIT Press, 411-420.
25. Mondada, F., Franzi, E. and lenne, P. (1993). Mobile Robot Miniaturization: A Tool for Investigation in Control Algorithms. In Yoshikawa, T. and Miyazaki, F (eds.) *Proceedings of*

- the Third International Symposium on Experimental Robotics*, Tokyo: Springer Verlag, 501-513.
26. Nolfi, S. and Floreano, D. (1998) Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution? *Artificial Life*, 4(4), 311-335.
 27. Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.
 28. Perez-Uribe, A., Floreano, D. and Keller, L. (2003) Effects of group composition and level of selection in the evolution of cooperation in artificial ants. In W. Banzhaf et al. (eds.) *Proceedings of the European Conference on Artificial Life*, Berlin: Springer Verlag, 128-137.
 29. Pettinaro, G.C., Kwee, I., Gambardella, L.M., Mondada, F., Floreano, D., Nolfi, S., Deneubourg, J.-L. and Dorigo, M. (2002) SWARM Robotics: A Different Approach to Service Robotics. In *Proceedings of the 33rd International Symposium on Robotics*, Stockholm, Sweden, October 7-11, 2002. International Federation of Robotics.
 30. Pfeifer R. (2000). On the relation among morphology and materials in adaptive behavior. In: J-A Meyer, A. Berthoz, D. Floreano, H.L. Roitblat, and S.W. Wilson (eds.) *From Animals to Animats 6. Proceedings of the VI International Conference on Simulation of Adaptive Behavior*. Cambridge, MA: MIT Press, 413-419.
 31. Pfeifer, R. and Scheier, C. (1999) *Understanding Intelligence*. Cambridge, MA: MIT Press.
 32. Roggen, D., Floreano, D. and Mattiussi, C. (2003) A Morphogenetic Evolutionary System: Phylogenesis of the POEtic Circuit . In Tyrrell, Haddow, Torresen (eds.) *Proceedings of the Fifth International Conference on Evolvable Systems*. Berlin: Springer Verlag, 153-164.
 33. Sims, K. (1994) Evolving 3D Morphology and Behavior by Competition. In R. A. Brooks and P. Maes (eds.) *Artificial Life IV*, Cambridge, MA: MIT Press, 28-39.
 34. Stoica, A., Zebulum, R., Keymeulen, D., Tawel, R., Daud, T. and Thakoor, A. (2001), Reconfigurable VLSI Architectures for Evolvable Hardware: from Experimental Field Programmable Transistor Arrays to Evolution-Oriented Chips. *IEEE Transactions on VLSI Systems, Special Issue on Reconfigurable and Adaptive VLSI Systems*, 9(1), 227-232.
 35. Thompson, A. (1998) *Hardware Evolution: Automatic design of electronic circuits in reconfigurable hardware by artificial evolution*. Berlin: Springer Verlag.
 36. Tyrrell, A.M., Sanchez, E., Floreano, D., Tempesti, G., Mange, D., Moreno, J.M., Rosenberg, J. and Villa, A. (2003) POEtic Tissue: An Integrated Architecture for Bio-inspired Hardware. In Tyrrell, Haddow, Torresen (eds.) *Proceedings of the International Conference on Evolvable Systems*. Berlin: Springer Verlag, 129-140.
 37. Wilson, E.O. (1971) *The Insect Societies*. Cambridge, MA: Harvard University Press.