Multimedia Contents

76. Evolutionary Robotics

Stefano Nolfi, Josh Bongard, Phil Husbands, Dario Floreano

Evolutionary Robotics is a method for automatically generating artificial brains and morphologies of autonomous robots. This approach is useful both for investigating the design space of robotic applications and for testing scientific hypotheses of biological mechanisms and processes. In this chapter we provide an overview of methods and results of Evolutionary Robotics with robots of different shapes, dimensions, and operation features. We consider both simulated and physical robots with special consideration to the transfer between the two worlds.

Evolutionary robotics is a method for the automatic creation of autonomous robots [76.1]. It is inspired by the Darwinian principle of selective reproduction of the fittest, captured by evolutionary algorithms [76.2]. In evolutionary robotics, robots are considered as autonomous artificial organisms that develop their own control system and body configuration in close interaction with the environment without human intervention. Drawing inspiration from principles of biological self-organization, evolutionary robotics includes elements of evolutionary, neural, developmental, and morphological systems. The idea that an evolutionary process could drive the generation of control systems dates back to at least the 1950s [76.3] with a more explicit form appearing in the mid 1980s with the ingenious thought experiments by neuroscientist *Valentino Braitenberg* on neurally driven vehicles [76.4]. In the early 1990s, the first generation of simulated artificial organisms with a genetic code describing the neural circuitry and morphology of a sensory motor system began evolving on computer screens [76.5–8]. At that time, real robots were still complicated and expensive machines that required specialized programming techniques and skillful manipulation. Towards the end of that period, a new generation of robots started to emerge that shared important characteristics with simple biological systems: robustness, simplicity, small size, flexibility, and modularity [76.9, 10]. Above all, those robots were designed so that they could be programmed and manipulated by people without engineering training. Those technological achievements, together with the growing influence of biological inspiration in artificial intelligence [76.11], coincided with the first evolutionary experiments on real robots [76.12–14] (VIDEO 39 and VIDEO 371), and the term evolutionary robotics was coined [76.15].

76.1 Method

The major methodological steps in evolutionary robotics proceed as follows (Fig. 76.1). An initial population of different artificial chromosomes, each encoding the control system (and possibly the morphology) of a robot, is randomly created. Each of these chromosomes is then decoded into a corresponding controller, for example a neural network (NN), and downloaded into the processor of the robot. The robot is then let free to act (move, look around, manipulate the environment) according to a genetically specified controller while its performance for a given task is automatically evaluated. Performance evaluation is done by a fitness function that measures how fast and straight the robot moves, how frequently it collides with obstacles, etc. This procedure is repeated for all chromosomes of the population. The fittest individuals (those that have received more fitness points) are allowed to reproduce by generating copies of their chromosomes with the addition of random modifications introduced by genetic operators (e.g., mutations and exchange of genetic material). The newly obtained population is then tested again on the same robot. This process is repeated for a number of generations until an individual is born which satisfies the fitness function set by the user. The control system of evolved robots, encoded in an artificial genome, is therefore generated by a repeated process of selective reproduction, random mutation, and genetic recombination, similarly to what happens in natural evolution (VIDEO 119).

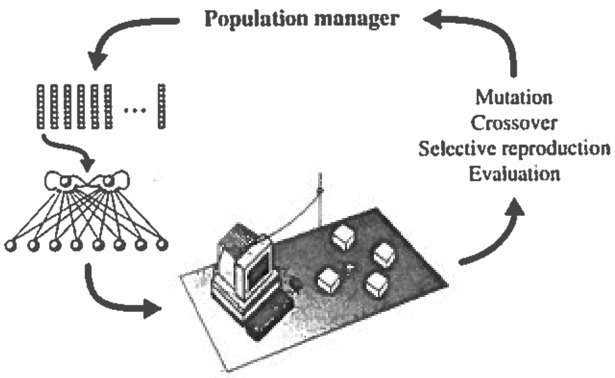

Fig. 76.1 Evolutionary experiments on a single robot. Each individual of the population is decoded into a corresponding neurocontroller which reads sensory information and sends motor commands to the robot every 300 ms while its fitness is automatically evaluated and stored away for reproductive selection

76.2 First Steps

In an early experiment on robot evolution without human intervention, carried out at Ecole Polytechnique Fédérale de Lausanne (EPFL) [76.12], a small wheeled robot was evolved for navigation in a looping maze (Fig. 76.2). The Khepera robot has a diameter of 55 mm and two wheels with controllable velocities in both directions of rotation. It also has eight infrared sensors, six on one side and two on the other side, that can function either in active mode to measure distance from obstacles or in passive mode to measure the amount of (infrared) light in the environment. The robot was connected to a desktop computer through rotating contacts that provided both power supply and data exchange through a serial port (VIDEO 39).

A simple genetic algorithm [76.16] was used to evolve the synaptic strengths of a neural network composed of eight sensory neurons and two motor neurons. Each sensory unit was clamped to one of the eight active infrared sensors whose value was updated every 300 ms. Each motor unit received weighted signals from the sensory units and from the other motor unit, plus a recurrent connection with itself with a 300 ms delay. The net input of the motor units was offset by a modifiable threshold and passed through a logistic squashing function. The resulting outputs, in the range [0, 1], were used to control the two motors so that an output of 1 generated maximum rotation speed in one direction, an output of 0 generated maximum rotation speed in the opposite direction, and an output of 0.5 did not generate any motion in the corresponding wheel. A population of 80 individuals, each coding the synaptic strengths and threshold values of the neural controllers, was initialized with all weights set to small random values centered around zero. Each individual was tested on the physical robot for 80 sensorimotor cycles (approximately 24 s) and evaluated at every cycle according to a fitness function with three components measured onboard the robot

$$\phi = V(1 - \sqrt{\Delta v})(1 - i) , \quad (76.1)$$

where V is the average rotation speed of the two wheels, Δv is the absolute value of the algebraic difference between the signed speed values of the wheels (positive is one direction, negative the other), and i is the normalized activation value of the infrared sensor with the highest activity. The first component is maximized by speed, the second by straight motion, and the third by distance from objects.

During the first 100 generations, both average and best fitness values grew steadily, as shown in Fig. 76.3. A fitness value of 1.0 would correspond to a robot moving straight at maximum speed in an open space and therefore was not attainable in the looping maze shown in Fig. 76.2, where some of the sensors were often active and where several turns were necessary to navigate. Fig. 76.4 shows the trajectory of the best individual of the last generation.

Although the fitness function did not specify in what direction the robot should navigate (given that it was perfectly circular and that the wheels could rotate in both directions), after a few generations all the best individuals moved in the direction corresponding to the side with the highest number of sensors. Individuals moving in the other direction had a higher probability of colliding into corners without detecting them and thus disappeared from the population. Furthermore, the cruising speed of the best evolved robots was approximately half of the maximum speed that could be technically achieved and did not increase even when the evolutionary experiment was continued up to 200 generations. Further analysis revealed that this self-limitation of the navigation speed had an adaptive function because, considering the sensory and motor

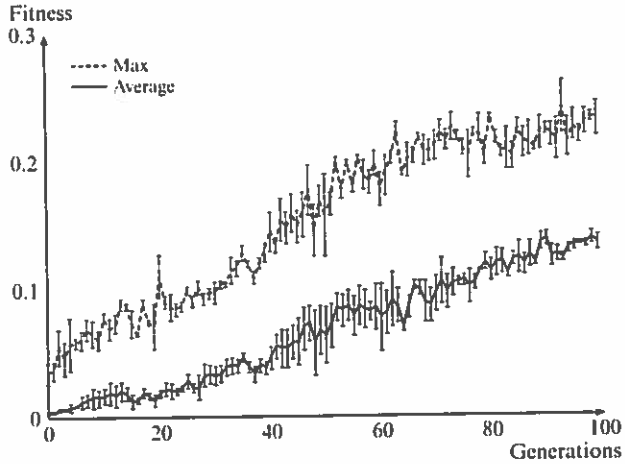

Fig. 76.3 Average fitness of the population and fitness of the best individual at each generation (*error bars* show standard error over three runs from different initial populations)

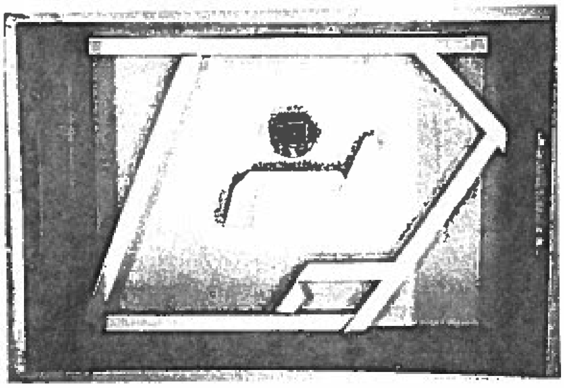

Fig. 76.2 Bird's-eye view of the desktop Khepera robot in the looping maze

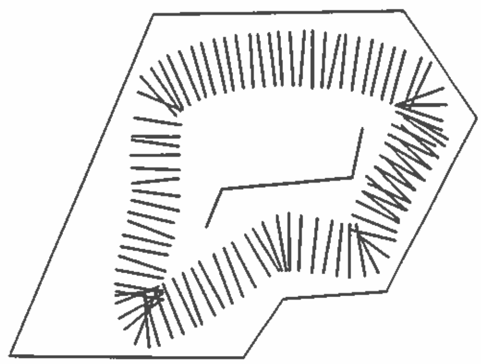

Fig. 76.4 Trajectory of the robot with the best neural controller of the last generation. Segments represent the axis between the two wheels. Data were recorded and plotted every 300 ms using an external laser positioning device

refresh rate together with the response profile of the distance sensors, robots that traveled faster had a higher risk of colliding with walls before detecting them; they gradually disappeared from the population.

Despite its simplicity, this experiments shows that evolution can discover solutions that match not only the computational requirements of the task to be solved, but also the morphological and mechanical properties of the robot in relation to its physical environment.

76.2.1 Evolution of Neural Controllers for Walking

Over the past 20 years or so, there has been a growing body of work on evolving controllers for various kinds of walking robots – a nontrivial sensorimotor coordination task. Early work in this area concentrated on evolving dynamical network controllers for simple (abstract) simulated insects (often inspired by cockroach studies) which were required to walk in simple environments [76.17, 18]. Earlier, *Beer* et al. had introduced a neural architecture for locomotion based on studies of cockroaches [76.19], which is shown in Fig. 76.5. The promise of this work soon led to versions of this methodology being used on real robots. Probably the first success in this direction was by *Lewis* et al. [76.14, 20] who evolved a neural controller for a simple hexapod robot using coupled oscillators built from continuous-time, leaky-integrator, artificial neurons. All evaluations were done on the actual robot with each leg connected to its own pair of coupled neurons, leg swing being driven by one neuron and leg elevation by the other. These pairs of neurons were cross connected, in a manner similar to that used by *Beer* and *Gallagher* [76.18] (Fig. 76.5), to allow coordination between the legs. In order to speed up the process, they employed staged evolution where first an oscillator capable of moving a leg was evolved and then an architecture based on these oscillators was further evolved to develop walking. The robot was able to execute an efficient tripod gait on flat surfaces.

Gallagher et al. [76.21] described experiments where neural networks controlling locomotion in an artificial insect were evolved in simulation and then successfully downloaded onto a real hexapod robot. This machine was more complex than Lewis et al.'s, with a greater number of degrees of freedom per leg. In this approach, each leg was controlled by a fully connected network of five continuous-time, leaky-integrator neurons, each receiving a weighted sensory input from that leg's angle sensor. Initially the architecture shown in Fig. 76.5 was used, with the connection weights and neuron time constants and biases under genetic control. This produced efficient tripod gaits for walking on flat surfaces. In order to produce a wider range of gaits operating at a number of speeds such that rougher terrain could be successfully negotiated, a different distributed architecture, more inspired by stick insect studies, was found to be more effective [76.22].

Galt et al. [76.23] used a genetic algorithm to derive the optimal gait parameters for a Robug III robot, an eight-legged, pneumatically powered walking and climbing robot. The individual genotypes represented parameters defining each leg's support period and the timing relationships between leg movements. These parameters were used as inputs to a mechanistic finite-state machine pattern-generating algorithm that drove the locomotion. Such algorithms, which are often used in conventional walking machines, rely on relatively simple control dynamics and do not have the same potential for the kind of sophisticated multigait coordination that complex dynamical neural network archi-

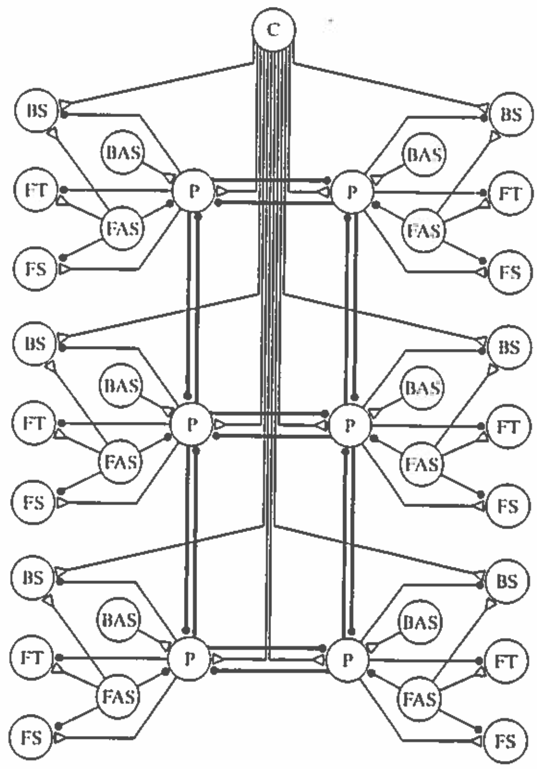

Fig. 76.5 Schematic diagram of a distributed neural network for the control of locomotion as used by *Beer* et al. [76.19]. Excitatory connections are denoted by *open triangles* and inhibitory connections are denoted by *filled circles*. C, command neuron; P, pacemaker neuron; FT, foot motor neuron; FS and BS, forward swing and backward swing motor neurons; FAS and BAS, forward and backward angle sensors

tectures, such as those described in this section, have been shown to produce. However, controllers were successfully evolved for a wide range of environments and to cope with damage and systems failure (although an individual controller had to be tuned to each environment; they were not able to self-adapt across a wide range of conditions). *Gomi* and *Ide* [76.24] evolved the gaits of an eight-legged robot (Fig. 76.6) using genotypes made of eight similarly organized sets of genes, each gene coding for leg motion characteristics such as the amount of delay after which the leg begins to move, the direction of the leg's motion, the end positions of both vertical and horizontal swings of the leg, and the vertical and horizontal angular speed of the leg. After a few dozen generations, where evaluation was on the robot, a mixture of tetrapod and wave gaits was obtained. Using the cellular encoding [76.25] developmental approach – which genetically encodes a grammar-tree program that controls the division of cells growing into a dynamical recurrent neural network of the kind used by Beer et al. – *Gruau* and *Quatramaran* [76.26] evolved a single-leg neural controller for the same eight-legged robot used by Gomi and Ide. This generated a smooth and fast quadrupod locomotion gait (VIDEO 372 and VIDEO 373). *Kodjabachian* and *Meyer* [76.27] extended this work to develop more sophisticated locomotion behaviors. *Jakobi* [76.28] successfully used his minimal simulation techniques (described in Sect. 76.3) to evolve controllers for the same eight-legged robot as Gruau. Evolution in simulation took less than 2 h on what would today be regarded as a very slow computer, and was then successfully transferred to the real robot. Jakobi evolved modular controllers based on Beer's continuous recurrent networks to control the robot as it engaged in walking about its environment, avoiding obstacles and seeking out goals depending on the sensory input. The robot could smoothly change gait, move backward and forward, and even turn on the spot. More recent work has used similar architectures to those explored by the researchers mentioned above, to control more mechanically sophisticated robots such as the Sony Aibo [76.29].

Recently there has been successful work on evolving coupled oscillator style neural controllers for the highly unstable dynamic problem of biped walking. *Reil* and *Husbands* [76.30] showed that accurate physics based simulations employing physics-engine software could be used to develop controllers able to generate successful bipedal gaits (VIDEO 374). *Reil* et al. have now significantly developed this technology to exploits its commercial possibilities, in the animation and games industries, for the real-time control of physically simulated three-dimensional (3-D) humanoid characters engaged in a variety of motor behaviors (refer to [76.31] for further details). Coupled neural oscillators have been evolved also to control the swimming pattern of articulated, snake-like, underwater robots using physics-based simulations [76.32].

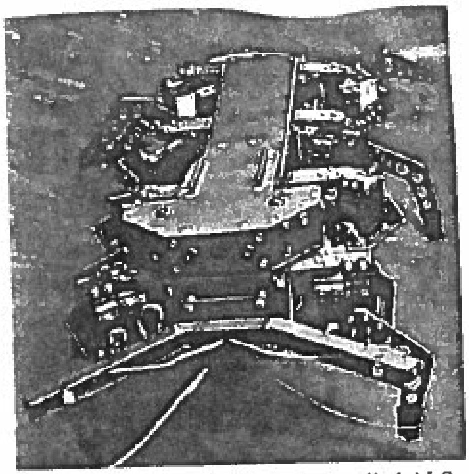

Fig. 76.6 The octopod robot built by Applied AI Systems Inc.

Vaughan has taken related work in another direction. He has successfully applied evolutionary robotics techniques to evolve a simulation of a 3-D ten-degree-of-freedom bipedal robot. This machine demonstrates many of the properties of human locomotion. By using passive dynamics and compliant tendons, it conserves energy while walking on a flat surface. Its speed and gait can be dynamically adjusted and it is capable of adapting to discrepancies in both its environment and its body's construction [76.33]. Parameters describing the body shape (leg segment lengths, hip width, etc.) and properties of a continuous dynamical neural network controller were under genetic control. The machine started out as a passive dynamic walker [76.34] on a slope, and then throughout the evolutionary process the slope was gradually lowered to a flat surface. The machine demonstrated resistance to disturbance while retaining passive dynamic features such as a passive swing leg. *Wischmann* and *Passemann* independently took a very similar approach [76.35]. Vaughan's original machine did not have a torso, but he has also successfully applied the method to a simplified two-dimensional (2-D) machine with a torso above the hips. When pushed, this dynamically stable bipedal machine walks either forward or backwards just enough to release the pressure placed on it. It is also able to adapt to external and internal perturbations as well as variations in body size and mass [76.36]. These biped examples make use of the co-evolution of body morphology and neural controller, an idea also used in earlier, more abstract, work on evolved bipedal locomotion by *Endo* et al. [76.37]. Although possible changes to the body morphology are quite tightly constrained, nonetheless this aspect was important. This theme is covered in more detail in the later section on Evolving Bodies, which also describes recent examples of evolved walking behaviors in the context of body-brain co-evolution.

McHale and *Husbands* [76.38, 39] have compared many forms of evolved neural controllers for bipedal and quadrupedal walking machines. Recurrent dynamical continuous time networks and GasNets (described in Sect. 76.7.3) were shown to have advantages in most circumstances. The vast majority of the studies mentioned above were conducted for relatively benign environments. Notwithstanding this observation, we can conclude that the more complex dynamical neural network architectures, with their intricate dynamics, generally produce a wider range of gaits and generate smoother, more adaptive locomotion than the more standard use of systems based on finite-state machines employing parameterized rules governing the timing and coordination of individual leg movements [76.40].

76.3 Simulation and Reality

Few of the experiments in the previous section were carried out entirely on physical robots because:

1. Evolution may take a long time, especially if it is carried out on a single robot that incarnates the bodies of all the individuals of the evolving population.
2. The physical robot can be damaged because populations always contain a certain number of poorly performing individuals (for example, colliding against walls) by effect of random mutations.
3. Restoring the environment to initial conditions between trails of different individuals or populations (for example, replenishing the arena with objects) may not always be feasible without human intervention.
4. Evolution of morphologies and evolution of robots that can grow during their lifetime is almost impossible with today's technology without some level of human intervention.

For those reasons, researchers often resort to evolution in simulation and transfer the evolved controllers to the physical robot. However, it is well known that programs that work well in simulations may not function properly in the real world because of differences in sensing, actuation, and in the dynamic interactions between robot and environment [76.41]. This *reality gap* is even more evident in adaptive approaches, such as evolutionary robotics, where the control system and morphology are gradually crafted through the repeated interactions between the robot and the environment. Therefore, robots will evolve to match the specificities of the simulation, which differ from the real world. Although these issues clearly rule out any simulation based on grid worlds or pure kinematics, over the last 10 years simulation techniques have dramatically improved and resulted in software libraries that model reasonably well dynamical properties such as friction, collision, mass, gravity, and inertia [76.42]. These software tools allow one to simulate articulated robots of variable morphology and their environment as fast as, or faster than, real time in a desktop computer.

Nonetheless, even physics-based simulations include small discrepancies that can accumulate over time and result in very different behavior from reality (for example, a robot may get stuck against a wall in simulation whereas it can get free in reality, or vice versa). Also, physics-based simulations cannot account for diversity of response profiles of the individual sensors, motors, and gears of a physical robot. Several methods can be used to cope with these problems and improve the quality of the transfer from simulation to reality.

A widely used method consists of adding independent noise to the values of the sensors provided by the model and to the end position of the robot computed by the simulator [76.43]. Some software libraries allow the introduction of noise at several levels of the simulation. This solution prevents evolution from finding solutions that rely on the specificities of the simulation model. Another method consists in sampling the actual sensor values of the real robot positioned at several angles and distances from objects of different texture. Those values are then stored in a look-up table and retrieved with the addition of noise according to the position of the robot in the environment [76.44]. This method proved to be very effective for generating controllers that transfer smoothly from simulation to reality. A drawback of this sampling method is that it does not scale up well to high-dimensional sensors (e.g., vision) or geometrically complicated objects.

Another method, also known as minimal simulations, consists of modeling only those characteristics of the robot and environment that are relevant for the emergence of desired behaviors [76.45]. These characteristics, which are referred to as base-set features, should be accurately modeled in simulation. Instead, all the other characteristics, which are referred to as implementation aspects, should be randomly varied across several trials of the same individual in order to ensure that evolving individuals do not rely on implementation aspects, but rely on base-set features only. Base-set features must also be varied to some extent across trials in order to ensure some degree of robustness of the individual with respect to base-set features, but this variation should not be so large that reliably fit controllers fail to evolve at all. This method allows very fast evolution of complex robot–environment situations, as in the example of the hexapod walk described in Sect. 76.2.1. A drawback of minimal simulations is that it is not always easy to tell in advance which are the base-set features that are relevant for the desired behavior.

Yet another method consists of the coevolution of the robot (control and/or morphology) and of the simulator parameters that are most likely to differ from the real world and that may affect the quality of the transfer [76.46]. This method consists of coevolving two populations, one encoding the properties of the robot and one encoding the parameters of the simulator. Coevolution happens in several passes through a two-stage process. In stage one, a randomly generated population of robots are evolved in the default simulator and the best individual is tested on the real robot while the time series of sensory values are recorded. In stage two, the population of simulators is evolved to reduce the difference between the time series recorded on the real robot and the time series obtained by testing evolved robots within the simulator. The best evolved simulator is then used for stage one where a new randomly generated population is evolved and the best individual is tested on the real robot to generate the time series for stage two of simulator evolution. This two-stage coevolution is repeated several times until the error between simulated and real robot behavior is the smallest possible. It has been shown that approximately 20 passes of the two-stage process are sufficient to evolve a good control system that could be transferred to an articulated robot. In that case, the real robot was used to test only 20 individuals.

A recent approach tackles the simulation to reality transfer problem by using a multi-objective formulation of ER in which two main objectives are optimized via a Pareto-based multi-objective evolutionary algorithm: (1) the fitness and (2) the transferability [76.47]. To evaluate the transferability a simulation-to-reality disparity measure was defined in terms of the difference in behavior between simulation and reality for any given controller. This measure is approximated for each member of the population and the method has successfully been demonstrated for walking behaviors [76.47].

Finally, another method consists of genetically encoding and evolving the learning rules of the control system, rather than its parameters (e.g., connection strengths). The parameters of the decoded control system are always initialized to small random values at the beginning of an individual lifetime and must self-organize using the learning rules [76.48]. This method prevents evolution from finding a set of control parameters that fit the specificities of the simulation model, and encourages emergence of control systems that remain adaptive to partially unknown environments. When such an evolved individual is transferred to the real robot, it will develop online its control parameters according to the genetically evolved learning rules and taking into account the specificities of the physical world. This method is described in more detail in Sect. 76.7.2 on evolution of learning.

76.4 Behavior as a Complex Adaptive System

Behavior is a dynamical process resulting from nonlinear interactions (occurring at a fast time rate) between the agent's control system, its body, and the environment [76.49, 50]. At any time step, the environment and the agent–environment relation influence the body and the motor reaction of the agent, which in turn influences the environment and/or the agent–environmental relation (Fig. 76.7). Sequences of these interactions lead to a dynamical process where the contributions of the different aspects (i.e., the robot's control system, the robot's body, and the environment) cannot be separated. This implies that even complete knowledge of the elements governing the interactions provides little insight into the behavior emerging from these interactions [76.51, 52].

An interesting property of evolutionary robotics is that it can enable to synthesize robots displaying a certain behavioral capacity without specifying the manner in which such capacity should be realized and/or the combination of elementary behaviors that should be produced and combined to achieve the desired overall capacity. This allows the evolving robots to discover and exploit behaviors that emerge from the interactions between the robot control system, the robot body and the environment and/or from the interaction between previously developed behavioral capacities [76.52]. The synthesis and exploitation of emergent properties, in turn, often allows evolving robots to discover solutions that rely on relatively parsimonious control policy and/or body structures.

For an example of how evolving robots can solve an adaptive task on the basis of a simple control policy, thanks to the possibility to exploit properties emerging from the agent/environmental interactions, let us consider the case of a Khepera robot placed in an arena surrounded by walls (Fig. 76.8) that should evolve an ability to forage by finding and remaining close to a food object (i. e., a cylindrical object) [76.53]. The robot is provided with eight infrared sensors and two motors controlling the desired speed of the two corresponding wheels. From the point of view of an external observer, solving this problem requires robots able to:

1. Explore the environment until an obstacle is detected.
2. Discriminate whether the obstacle detected is a wall or a cylindrical object.
3. Approach or avoid the object depending on the object type.

A detailed analysis of the sensory patterns experienced by the robot indicated that the task of discriminating the two objects is far from trivial since the two classes of sensory patterns experienced by robots close to a wall and close to cylindrical objects overlap significantly. However, robots evolved for the ability to solve this task resorted to a strategy that does not require to explicitly discriminate of the two types of objects [76.53]. This solution (VIDEO 116) consists in reacting to sensory states so that the robot/environmental dynamics converge into a limit cycle near the cylindrical object, in which the robot keep moving forth and back and left and right, and not near a wall (Fig. 76.9).

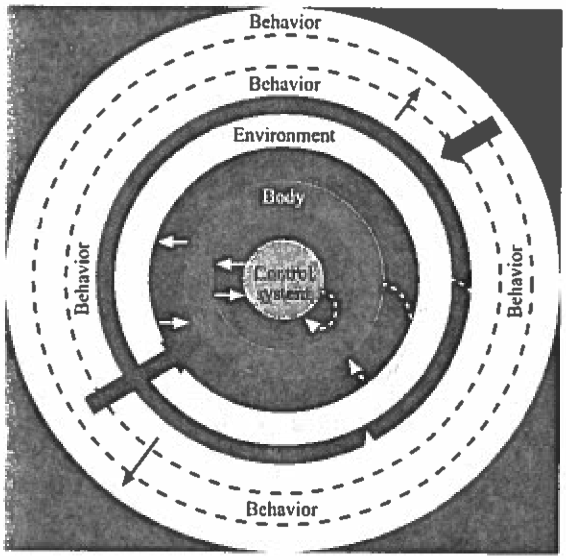

Fig. 76.7 Schematization of how: (1) behavior can emerge from several nonlinear interactions, occurring at fast time rates, between the agent's control system, its body, and the environment, and (2) behavior can display a multi-level organization in which the robot/environmental interactions and the interaction between lower-level behaviors give rise to higher level behaviors that later affect the interactions from which they originate

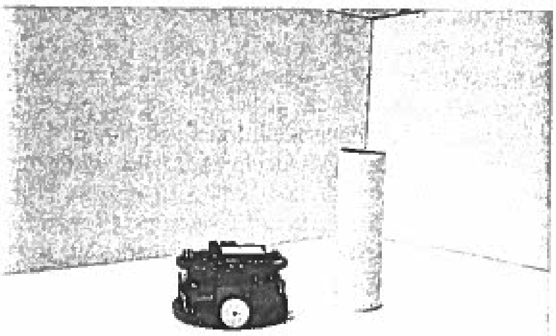

Fig. 76.8 The environment and the robot. The environment consists of an arena of 60×35 cm and contains a cylindrical objects placed at a randomly selected location

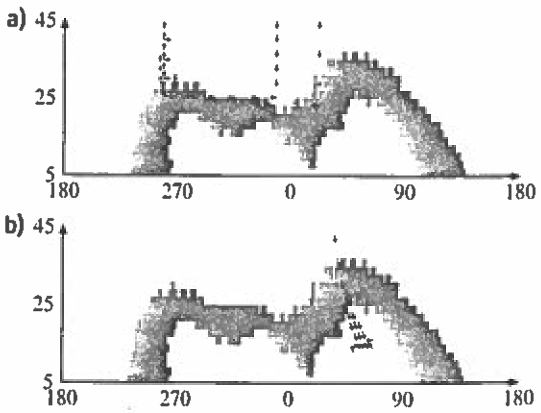

Fig. 76.9a,b Angular trajectories of an evolved robot close to a wall (**a**) and to a cylinder (**b**). The picture was obtained by placing the robot at a random position in the environment, leaving it free to move for 500 cycles, and recording its relative movements with respect to the two types of objects for distances smaller than 45 mm. For sake of clarity, *arrows* are used to indicate the relative direction, but not the amplitude of movements

The possibility to discover and rely on these forms of emergent behavior allows evolving robots to find computationally simple solutions to apparently complex problems. Indeed, the foraging task described above can be solved by a Khepera robot provided with a simple reactive controller (i.e., a feedforward neural network with eight sensory neurons that encode the state of the corresponding infrared sensors directly connected to two motor neurons that set the desired speed of the two wheels).

76.4.1 Behavior Recombination and Re-Use

Evolving robots can recombine and re-use acquired elementary behavioral capacities to produce higher level behaviors. This has been demonstrated in a series of simulation experiments in which a population of humanoid robots provided with an articulated arm, a camera, and a touch sensor on the palm of the hand, have been evolved for the ability to execute two-words imperative sentences constituted by the combination of three action and three object words (reach, touch, move, red-object, green-object, and blue-object) encoded by six corresponding binary sensors [76.54].

During the evolutionary process the robots were evaluated for the ability to *comprehend* seven out of the nine possible sentences (that can be generated by combining the three action and the three object words) by executing the seven corresponding behaviors. The robots were then post-evaluated also on the other two sentences not experienced during the evolutionary process.

Some of the evolved robots were able to develop the required skills and to generalize their capacities to the two new sentences by executing the appropriate corresponding behaviors (VIDEO 411). Differently from the other individuals, the robots able to generalize where characterized by a hierarchical organization in which the nine behaviors were produced by combining over time a set of elementary behaviors and in which the same elementary behaviors was re-used to produce different high-level behaviors. More specifically the robots able to generalize displayed a reach-X behavior (that consisted in moving the arm toward a red, or green, or blue object), a touch behavior (that consisted in moving the hand until the object is touched irrespectively from the color), and a move behavior (that consists in keep moving the hand also after the object has been touched irrespectively from the color) and combined these lower-level behaviors in a compositional manner to produce the nine required higher-level actions. This means that, for example, the same *reach red-object* behavior was used incombination with the *touch* or the *push* behavior to produce a *touch the red-object* and a *push the red-object* behavior. For other works discussing the emergence and the role of multi-level behavioral organizations see [76.52].

76.4.2 Sensory-Motor Coordination

By acting robots inevitably modify the robot-environmental relation and/or the environment and consequently the stimuli that they will experience next. By exploiting the possibility to actively influence the perceived stimuli through actions, robots can find adaptive solutions based on parsimonious control policies. Artificial evolution constitutes an effective method for discovering such type of solutions that are often hard to imagine from the point of view of a human observer. Indeed, examples of clever use of sensory-motor coordination abounds in the evolutionary robotics literature.

Let us consider, for instance, the case of a Khepera robot endowed with infrared and wheels speed sensors, that can forage by ramaining close to large cylindrical objects (food) while avoiding small cylindrical objects (dangers) [76.55]. From a passive perspective, that does not take into account the fat that the robot can self-select useful stimuli through action, the ability to discriminate between sensory stimuli experienced near small and large cylindrical objects requires a relatively complex control policy since the two classes of stimuli strongly overlap in the robot's perceptual space. On the other hand, the exploitation of sensory-motor coordination can can allow the robots to simplify the discrimination problem.

Indeed, evolving robots tend to converge on a rather simple solution that consists in circling around the cylindrical objects, as soon as an object is perceived, and in using the differential speed of the left and right wheels sensed during the execution of the object-circling behavior to decide to keep circling around the object (in the case of small differential speeds) or to abandon the object (in the case of large differential speeds). Indeed, the execution of the object-circling behavior allows the robots to experience sensory stimuli on the wheel sensors that are well differentiated for small and large objects. This, in turn, allows them to solve the object discrimination problem with a rather simple but reliable control policy. For other examples see [76.53, 56, 57].

76.5 Evolving Bodies

Most evolutionary robotics experiments – and most robotics experiments in general – assume that the body plan of the robot has already been designed; an optimization method is used to improve the control policy only (The term *body plan* is here used to denote all aspects of a robot's design other than its control policy. Such design considerations include the robot's mechanical layout and material properties as well as its sensor and motor distributions.). This emphasis belies an assumption within the field of robotics, which is that control policy design is non-intuitive and thus should be automated, while choosing an appropriate robot body plan is intuitive and thus can be manually designed.

However, it has been shown that the careful design of the robot's body can have large and desirable impacts on its resulting behavior. For example, proper curvature on the underside of a biped robot's feet (along with other settings) can allow it to walk down a declined plane with no control policy at all [76.58]. Or, that modifications to an anthropomorphic robot arm and hand can facilitate the evolution of active categorical perception [76.59] (Active categorical perception occurs when a robot or animal actively interacts with objects of interest, and the sensory stimulation resulting from this physical interaction allows for categorization of those objects.). These results fit with the view of embodied behavior, as outlined in Fig. 76.7: because behavior arises from the interaction between a robot's body and its environment, alterations to the robot's body will alter the resulting behavior.

This suggests that it is useful to automatically improve not just a robot's control policy but also its body plan. Evolutionary algorithms are a uniquely well suited tool for this task because, unlike many learning methods, they do not make assumptions about the structure of the system being optimized: the length of a robot's leg – or the number of legs – can be evolved just as easily as can the strength of a synaptic connection in an artificial neural network.

76.5.1 Co-Evolving Body and Brains

The Sussex group was the first to demonstrate the evolution of robot morphology: they evolved sensor placements on a physical robot [76.15], although the other aspects of the robot's body plan remain fixed. A year later *Sims* [76.61] demonstrated an evolutionary algorithm that improved the structure and parameters of the robots' body plans and control policies. Although Sim's *creatures* were virtual and operated in a simulated environment, the robots exhibited a wide range of intuitive and non-intuitive body plans that allowed them to swim, walk or compete over a limited resource [76.62].

Funes and *Pollack* [76.63] demonstrated that it was possible to evolve three-dimensional forms in simulation, build them in reality, and have the physical structure act similarly to the originally-evolved simulated structure. This was followed by work from the same group in which robots evolved in simulation were manufactured as physical robots using 3-D printing technology [76.60] (Fig. 76.10). Although only the plastic frame of the robot was printed and the electronics and battery had to be manually added, this served as a demonstration that, in principle, robot design could be automated using evolutionary algorithms and robot manufacture could be automated using rapid prototyping [76.64].

Since then, a number of research groups have evolved robot body plans and control policies simultaneously for various purposes. Some researchers have adapted this approach for studying biological questions. For example *Long* et al. [76.65] have evolved the stiffness of artificial tails attached to physical swimming robots: robots with tails of differing stiffness have differing abilities to swim fast or turn well. This provides a unique experimental tool for investigating how backbones originally evolved in early vertebrates. *Clark* et al. [76.66] also evolved the material properties of part of a robot fish, but focussed in this case on evolving the stiffness and shape of its fins in simulation. This project had an engineering aim: the evolved fins were

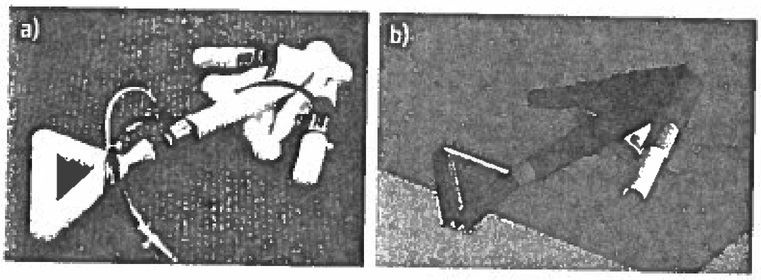

Fig.76.10a,b An example evolved robot from the GOLEM project [76.60]. **(a)** The virtual robot, as originally evolved in the simulated environment. **(b)** The physical robot comprised of a 3-D-printed plastic frame and manually-added electronics

manufactured and tested on a physical robot fish and were found to aid desirable swimming behavior.

Evolutionary algorithms can be used to explore the space of possible robot body plans, but they can also be used to explore metamorphosis, or how a robot's body plan might change over its lifetime. In recent work *Bongard* [76.67] compared two approaches to evolving walking behaviors for upright legged robots (VIDEO 771). In the first approach, upright legged robots were evolved until successful walking was discovered. In the second approach, locomotion was first evolved for legless, anguilliform robots that gradually grew legs while moving. As evolution proceeded in this second approach, later generations of robots gradually lost this infant legless body plan and instead were born with the upright, legged body plan. It was found that walking evolved for the upright legged robots more rapidly in the second approach, and that the evolved controllers were more robust. The explanation for this result is that the legless robot provides a form of scaffolding that accelerates search: it is easier for evolution to generate locomotion for the legless robot because with the anguilliform body plan the robot cannot fall over (Scaffolding is the phenomenon in which a teacher introduces some aspect into the learner's environment that helps the learner to grasp a concept and then later refine the concept when the scaffold is removed [76.68]. The canonical example of scaffolding is training wheels for bicycles.). This locomotion strategy is then refined subsequently by evolution to successfully control the upright and legged (and thus unstable) robot.

The simultaneous evolution of robot body plans and control policies offers other avenues for investigating the relationship between body, brain and environment. In [76.69] it was found that more complexly-shaped robots were produced when evolved to walk over rough terrain than when evolved to walk over flat terrain (VIDEO 772). This was due to the fact that robots evolved in rough terrain evolved appendages and hooks to gain purchase between outcroppings and then pull or push themselves forward.

Most recently *Hiller* and *Lipson* [76.70] have demonstrated the evolution of soft robots: this requires evolving not just the control policy of the robot and its physical shape, but also the material properties of each voxel comprising the machine. This allows for complex three-dimensional patterning of soft and rigid material throughout the robot, which can be exploited by evolution to produce locomotion [76.71]. Soft robots are an ideal vehicle for demonstrating the power of evolutionary algorithms: the design and control of such machines is highly nonintuitive, making manual design extremely difficult.

76.5.2 Self-Modeling

The evolution of robot body plans can be useful not just for robot design but also for increasing the adaptivity of a physical robot once it has been designed and deployed. For example in [76.72] it was shown that a physical robot could be equipped with an onboard simulator that the robot could use to continuously evolve models of itself. These self-models reflect the mechanical construction of the robot. This method was found useful for robots that might sustain unanticipated damage such as the mechanical separation of a leg: the robot diagnoses the damage; it then evolves a simulated damaged robot that accurately reflects the physical damage; it evolves a compensatory control policy internally using the simulated, damaged robot; and finally the physical robot uses the internally evolved compensatory control policy to continue moving despite its injury. Figure 76.11 outlines this method.

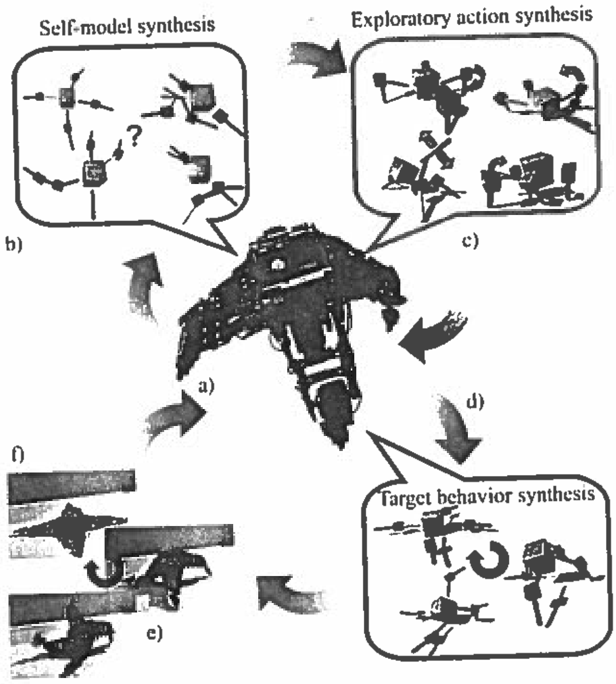

Fig.76.11a–f The machine begins by performing a random action and collects the resulting sensor-motor data (**a**). An evolutionary algorithm evolves a population of simulated robots that, when they move, produce similar sensor-motor data as the physical machine (**b**). Another algorithm searches for a new action for the physical robot to perform (**c**). The physical robot performs the new action, and re-evolves self-models to explain the result of the first action and the new action ((**a**)). After several cycles of this self-modeling, the best self-model is used to evolve new behaviors (**d**). Finally, the physical robot executes these newly-evolved behaviors (**e**,**f**)

A method for generalizing this to robot swarms was demonstrated in [76.73]. Each robot in the swarm maintained its own self-modeling engine, but would periodically export its best self-model and control policy to others in the swarm. The result of this was that if one robot was damaged and recovered, a second robot that suffered similar damage recovered more rapidly.

Finally, instead of modeling the self, a robot could create a model of another robot in its vicinity. Even better, the robot could evolve a model of the other robot's intentions and use this information to aid or thwart the other robot's actions, as demonstrated in [76.74]. Although the robots did not model each other's body plans, this ability to model others in general is known as Theory of Mind. One could imagine how increasing levels of recursion of such embedded mind reading could provide continued evolutionary pressure toward increasingly intelligent machines.

76.6 Seeing the Light

Pioneering experiments on evolving visually guided behaviors were performed at Sussex University [76.75] on a specially designed gantry robot (Fig. 76.12, see also VIDEO 371). Discrete-time dynamical recurrent neural networks and visual sampling morphologies were concurrently evolved: the brain was developed in tandem with the visual sensor [76.13, 76.77]. The robot was designed to allow real-world evolution by having *off-board* power and processing so that the robot could be run for long periods while being monitored by automatic fitness evaluation functions. A charge-coupled device (CCD) camera points down towards a mirror angled at 45° as shown in Fig. 76.12. The mirror can rotate around an axis perpendicular to the camera's image plane. The camera is suspended from the gantry, allowing motion in the X, Y, and Z dimensions. This effectively provides an equivalent to a wheeled robot with a forward-facing camera when only the X and Y dimensions of translation are used. The additional dimension allows flying behaviors to be studied.

The apparatus was initially used in a manner similar to the real-world experiments on navigation in the looping maze with the miniature mobile robot described in Sect. 76.2. A number of visually guided navigation behaviors were successfully achieved, including navigating around obstacles, tracking moving targets, and discriminating between different objects [76.76]. The evolutionary process was incremental. The ability to distinguish between two different targets was evolved on top of the single target-finding behavior. The chromosome was of dynamic length so the neurocontroller was structurally further developed by evolution to achieve the new task (neurons and connections added). In the experiment illustrated in Figs. 76.12 and 76.13,

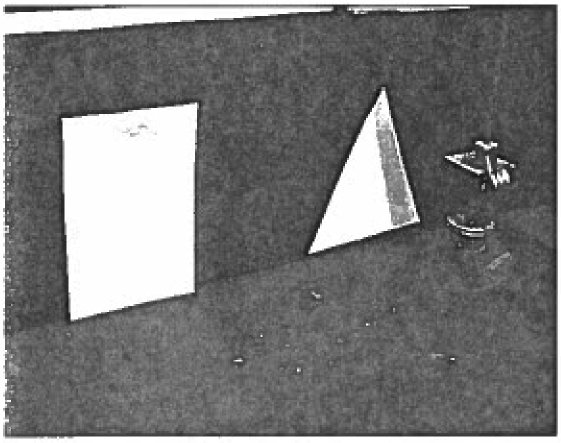

Fig. 76.12 The gantry robot used in the visual discrimination task. The camera inside the top box points down at the inclined mirror, which can be turned by the stepper motor beneath. The lower plastic disk is suspended from a joystick to detect collisions with obstacles

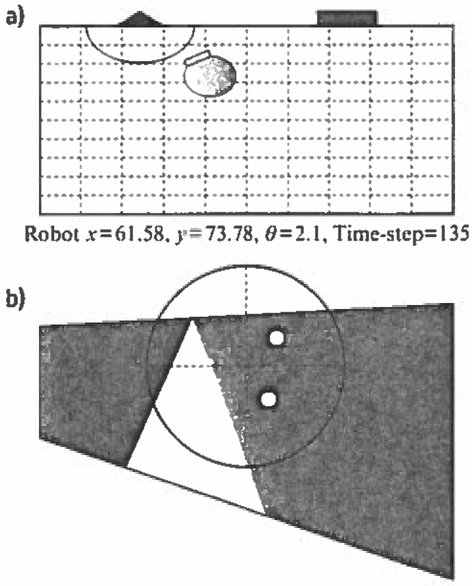

Fig. 76.13a,b The shape discrimination task. **(a)** The position of the robot in the arena, showing the target area in front of the triangle. **(b)** The robot camera's field of view showing the visual patches selected by evolution for sensory input

starting from a random position and orientation, the robot had to move to the triangle rather than the rectangle. This had to be achieved irrespective of the relative positions of the shapes and under very noisy lighting conditions. Recurrent neural network controllers were evolved in conjunction with visual sampling morphologies. Only genetically specified patches from the camera image were used (by being connected to input neurons according to the genetic specification). The rest of the image was thrown away. This resulted in extremely minimal systems using only two or three pixels of visual information, yet still able to perform the task reliably under highly variable lighting conditions [76.13, 76].

This was another example of staged, or incremental, evolution to obtain control systems capable of solving problems that are either too complex or may profit from an evolutionary methodology that discovers, preserves, and builds upon subcomponents of the solution. For an evolutionary method that incorporate strategies to explicitly address this issue, interested readers may refer to [76.78]. However, staged evolution remains a poorly explored area of evolutionary robotics that deserves further study and a more principled approach [76.79] in order to achieve increasingly complex robotic systems.

76.6.1 Coevolution of Active Vision and Feature Selection

Machine vision today can hardly compete with biological vision despite the enormous power of computers. 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 take time to explore the image over time, searching for features and dynamically integrating information over time.

Active vision is the sequential and interactive process of selecting and analyzing parts of a visual scene [76.80–82]. *Feature selection* instead is the development of sensitivity to relevant features in the visual scene to which the system selectively responds, e.g., [76.83]. Each of these processes has been investigated and adopted in machine vision. However, the combination of active vision and feature selection is still largely unexplored. An intriguing hypothesis is that coevolution of active vision and feature selection could greatly simplify the computational complexity of vision-based behavior by facilitating each other's task.

This hypothesis was investigated in a series of experiments [76.84] on coevolution of active vision and feature selection for behavioral systems equipped with a primitive moving retina and a deliberately simple neural architecture (Fig. 76.14). The neural architecture was composed of an artificial retina and two sets of output units. One set of output units determined the movement and zooming factor of the retina, and the other set of units determined the behavior of the system, such as the response of a pattern-recognition system, the control parameters of a robot, or the actions of a car driver. The neural network was embedded in a behavioral system and its input/output values were updated every 300 ms while its fitness was computed. Therefore, the synaptic weights of this network were responsible for both the visual features on which the system based its behavior and for the motor actions necessary to search for those features.

In a first set of experiments, the neural network was embedded in a simulated pan–tilt camera and asked to discriminate between triangles and squares of different size that could appear at any location of a screen (Fig. 76.15a), a perceptual task similar to that explored with the gantry robot described in Sect. 76.5. The visual system was free to explore the image for 60 s while continuously reporting whether the current screen showed a triangle or a square. The fitness was proportional to

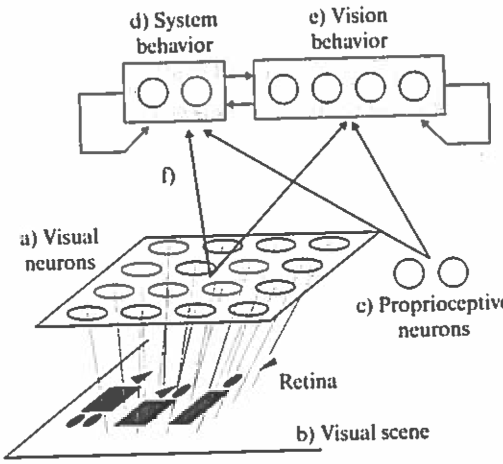

Fig. 76.14a–f The neural architecture of the active vision system is composed of: (a) a grid of visual neurons with nonoverlapping receptive fields whose activation is given by (b) the grey level of the corresponding pixels in the image; (c) a set of proprioceptive neurons that provide information about the movement of the vision system; (d) a set of output neurons that determine the behavior of the system (pattern recognition, car driving, robot navigation); (e) a set of output neurons that determine the behavior of the vision system; and (f) a set of evolvable synaptic connections. The number of neurons in each subsystem can vary according to the experimental settings

the amount of correct responses accumulated over the 60 s for several screenshots containing various instances of the two shapes. Evolved systems were capable of correctly identifying the type of shape with 100% accuracy after a few seconds despite the fact that this recognition problem is not linearly separable and that the neural network does not have hidden units, which in theory are necessary to solve nonlinearly separable tasks. Indeed, the same neural network presented with the same set of images and trained with supervised learning, but without the possibility to actively explore the scene, was not capable of solving the task. The evolved active vision system developed sensitivity to vertical edges, oriented edges and corners, and used its movement to search for these features in order to tell whether the shape was a triangle or a square. These features, which are also found in the early visual system of almost all animals, are invariant to size and location.

In a second set of experiments, the neural network was embedded in a simulated car and was asked to drive over several mountain circuits (Fig. 76.15b). The simulator was a modified version of a car race video game. The neural network could move the retina across the scene seen through the windscreen at the driver's seat and control the steering, acceleration, and braking of the car. The fitness was inversely proportional to the time taken to complete the circuits without exiting the road. Evolved networks completed all circuits with time laps competitive to those of well-trained students controlling the car with a joystick. The evolved network started by searching for the edge of the road and tracked its relative position with respect to the edge of the windscreen in order to control steering and acceleration. This behavior was supported by the development of sensitivity to oriented edges.

In a third set of experiments, the neural network was embedded in a real mobile robot with a pan–tilt camera that was asked to navigate in a square arena with low walls located in an office (Fig. 76.16). The fitness was proportional to the amount of straight motion measured over two minutes. Robots that hit the walls because they watched people or other irrelevant features of the office had lower fitness than robots that could perform long straight paths and avoid walls of the arena. Evolved robots tended to fixate the edge between the floor and the walls of the arena, and turned away from the wall when the size of its retinal projection became larger than a threshold (VIDEO 36). This combination of sensitivity to oriented edges and looming is also found in the visual circuits of several insects and birds.

In a further set of experiments [76.85], the visual pathway of the neural network was augmented by an intermediate set of neurons whose synaptic weights could be modified by Hebbian learning [76.86] while the robot moved in the environment. All the other synaptic weights were genetically encoded and evolved. The results showed that lifelong development of the receptive fields improved the performance of evolved robots and allowed robust transfer of evolved neural controllers from simulated to real robots, because the receptive fields developed sensitivity to features encountered in the environment where they happen to be born (see also the section above on simulation and reality). Furthermore, the results showed that the development of visual receptive fields was significantly and consistently affected by active vision as compared to the development of receptive fields passively exposed to the same set of sample images. In other words, robots evolved with active vision developed sensitivity to a smaller subset of features in the environment and actively tracked those features to maintain a stable behavior.

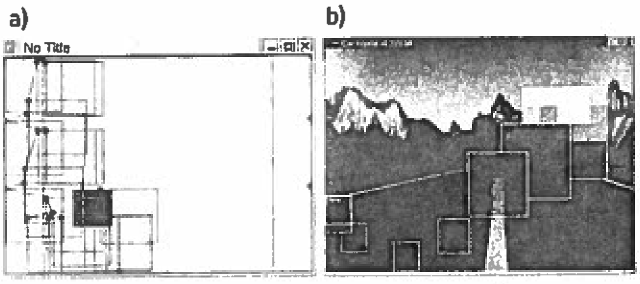

Fig. 76.15 **(a)** An evolved individual explores the screen searching for the shape and recognizes it by the presence of a vertical edge. **(b)** Search for the edge of the road at the beginning of a drive over a mountain road

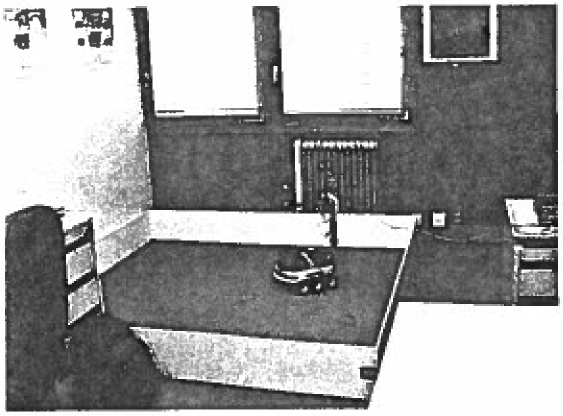

Fig. 76.16 A mobile robot with a pan–tilt camera is asked to move within the walled arena in the office environment

76.7 Computational Neuroethology

Evolutionary robotics is also used to investigate open questions in neuroscience and cognitive science [76.87–90] because it offers the vantage point of a behavioral system that interacts with its environment [76.91]. Although the results should be carefully considered when drawing analogies with biological organisms, evolutionary robotics can generate and test hypotheses that could be further investigated with mainstream neuroscience methods.

For example, the active vision system with Hebbian plasticity described in the previous section was used to answer a question raised by *Held* and *Hein* [76.92] in the 1960s. The authors devised the apparatus shown in Fig. 76.17 where the free movements of a kitten (*active kitten*) were transmitted to a second kitten that was carried in a gondola (*passive kitten*). The second kitten could move its head, but its feet did not touch the ground. Consequently, the two kitten received almost identical visual stimulation, but only one of them received that stimulation as a result of body self-movement. After a few days in that environment, only the active kitten displayed normal behavior in several visually guided tasks. The authors suggested the hypothesis that proprioceptive motor information resulting from generation of actions was necessary for the development of normal, visually guided behavior.

The kitten experiments were replicated by cloning an evolved robot controller and randomly initializing the synaptic values of the adaptive visual pathways in both clones. One cloned robot was then left free to move in a square environment while the other cloned robot was forced to move along imposed trajectories, but was free to control its camera position, just like the passive kitten [76.94]. The results indicated that the visual receptive fields and behaviors of passive robots differ significantly from those of active robots. Furthermore, passive robots that were later left free to move were no longer capable of properly avoiding walls. A thorough analysis of neural activation correlated with behavior of the robot and even transplantation of neurons across active and passive robots revealed that the poor performance was due to the fact that passive robots could not completely select the visual features they were exposed to. Consequently, passive robots developed sensitivity to features that were not functional to their normal behavior and interfered with other dominant features in the visual field. Whether this explanation also hold for living animals remains to be further investigated, but at least these experiments indicated that motor feedback is not necessary to explain the pattern of pathological behavior observed in animals and robots.

76.7.1 Emergence of Place Cells

Let us now consider the case of an animal exploring an environment and periodically returning to its nest to feed. It has been speculated that this type of situation requires the formation of spatial representations of the environment that allow the animal to find its way home [76.95]. Different neural models with various degrees of complexity and biological detail that could provide such functionality have been proposed [76.96, 97].

Would a robot evolved under similar survival conditions develop a spatial representation of the environment and, if so, what type of representation would that be? These questions were explored using the same Khepera robot and evolutionary methodology described in Sect. 76.2 for reactive navigation in the looping maze. The environment was a square arena with a small patch on the floor in a corner where the robot could instantaneously recharge its (simulated) battery (Fig. 76.18). The environment was located in a dark room with a small light tower over the *recharging station*.

The sensory system of the robot was composed of eight distance sensors, two ambient-light sensors (one

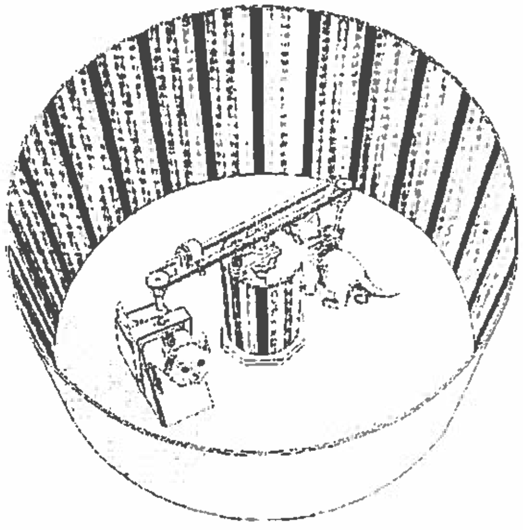

Fig. 76.17 The original apparatus in [76.92], where the gross movements of a kitten moving almost freely were transmitted to a second kitten that was carried in a gondola. Both kittens were allowed to move their head. They received essentially the same visual stimulation because of the unvarying pattern on the walls and the center post of the apparatus (after [76.93], with permission)

on each side), one floor-color sensor, and a sensor for battery charge level. The battery lasted only 20 s and had a linear discharge. The evolutionary neural network included five fully connected internal neurons between sensory and motor neurons. The same fitness function described in Sect. 76.2 for navigation in the looping maze was used, except for the middle term which had been used to encourage straight navigation in the looping maze. The fitness value was computed every 300 ms and accumulated over the life span of the individual. Therefore, individuals who discovered where the charger was could live longer and accumulate more fitness by exploring the environment (individuals were killed if they survived longer than 60 s to limit the experimentation time).

The same physical robot evolved for 10 days and nights as both the fitness and life span of individuals continued to increase (Fig. 76.19). After approximately 200 generations, the robot was capable of navigating around the environment, covering long trajectories while avoiding both walls and the recharging area (VIDEO 418). When the battery was almost discharged it initiated a straight navigation towards the recharging area and exited immediately after battery recharge to resume navigation. Best evolved individuals always entered the recharging area one or two seconds before full discharge of the battery. That implies that robots must somehow calibrate the timing and trajectory of their homing behavior depending on where they happened to be in the environment.

In order to understand how that behavior could possibly be generated, a set of neuroethological measures were performed using a laser positioning device that provided exact position and orientation of the robot every 300 ms. By correlating the robot position and behavior with the activation of the internal neurons in real time while the evolved individual freely moved in the environment, it was possible to see that some neurons specialized for reactive behaviors, such as obstacle avoidance, forward motion, and battery monitoring. Other neurons instead displayed more complex activation patterns. One of them revealed a pattern of activation levels that depended on whether the robot was oriented facing the light tower or facing the opposite direction (Fig. 76.20). In the former case, the activation pattern reflected zones of the environment and paths typically followed by the robot during exploration and homing. For example, the robot trajectory towards the recharging area never crossed the two *gate walls* visible in the activation maps around the recharging station. When the robot faced the opposite direction, the same neuron displayed a gradient field orthogonally aligned with the recharging area. This gradient provides an indication of the distance from the recharging area. Interestingly, this pattern of activity is not significantly affected by the charge level of the battery.

The functioning of this neuron reminds of the classic findings on the hippocampus of the rat brain where some neurons (also known as *place cells*) selectively fire when the rat is in specific areas of the environment [76.98]. Also, the orientation-specific pattern of neural activation measured on the evolved robot is rem-

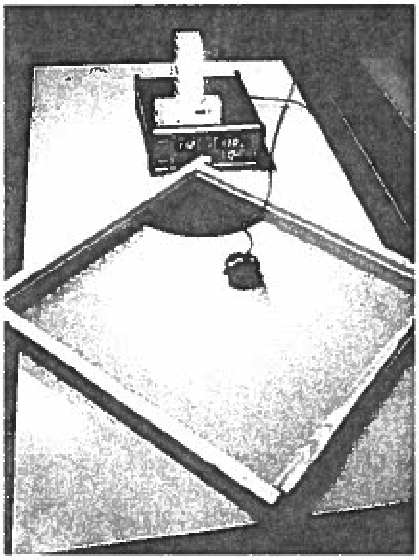

Fig. 76.18 Bird's eye view of the arena with the light tower over the recharging station and the Khepera robot

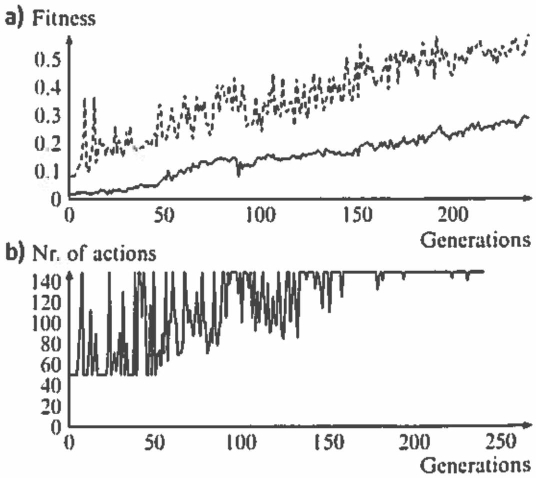

Fig. 76.19 (a) Average population fitness (*continuous line*) and fitness of the best individual (*dotted line*). (b) Life span of the best individuals measured as number of sensorimotor cycles, or actions. Individuals start with a full battery which lasts 50 actions (20 s), if not recharged. The maximum life span is 150 actions

iniscent of the so-called *head-direction neurons* in the rat hippocampus, which are positioned nearby place cells, whose firing patterns depend on the rat heading direction with respect to an environmental landmark [76.99]. Although the analogy between brains of evolved robots and of biological organisms should not be taken too literally, these results indicate that the two organisms converge towards a functionally similar neural strategy, which may be more efficient to address this type of situation than a strategy that does not rely on representations (but only on reactive strategies such as random motion, light following, or dead reckoning).

76.7.2 Spiking Neurons

The great majority of biological neurons communicate using self-propagating electrical pulses called *spikes*, but from an information-theoretic perspective it is not yet clear how information is encoded in the spike train. Connectionist models [76.100], by far the most widespread, assume that what matters is the *firing rate* of a neuron, that is, the average quantity of spikes emitted by the neuron within a relatively long time window (for example, over 100 ms). Alternatively, what matters is the average number of spikes of a small population of neurons at a give point. In these models the real-valued output of an artificial neuron represents the firing rate, possibly normalized relatively to the maximum attainable value. Pulsed models [76.101], instead, are based on the assumption that the *firing time*, that is, the precise time of emission of a single spike, may convey important information [76.102]. Spiking neuron models have slightly more complicated dynamics of synaptic and membrane integration. Depending on one's theory of what really matters, connectionist or spiking models are used.

However, designing circuits of spiking neurons that display a desired functionality is still a challenging task. The most successful results in the field of robotics obtained so far focused on the first stages of sensory processing and on relatively simple motor control [76.103, 104]. Despite these implementations, there are not yet methods for developing complex spiking circuits that could display minimally cognitive functions or learn behavioral abilities through autonomous interaction with a physical environment.

Artificial evolution represents a promising methodology to generate networks of spiking circuits with desired functionalities expressed as behavioral criteria (fitness function). Evolved networks could then be examined to detect what communication modality is used and how that correlates with observed behavior of the robot.

Floreano and *Mattiussi* [76.105] evolved a fully connected network of spiking neurons for driving a vision-based robot in an arena painted with black stripes of variable size against a white background (Fig. 76.21). The Khepera robot used in these experiments was equipped with a vision turret composed of one linear array of grayscale photoreceptors spanning a visual field of 36°. The output values of a bank

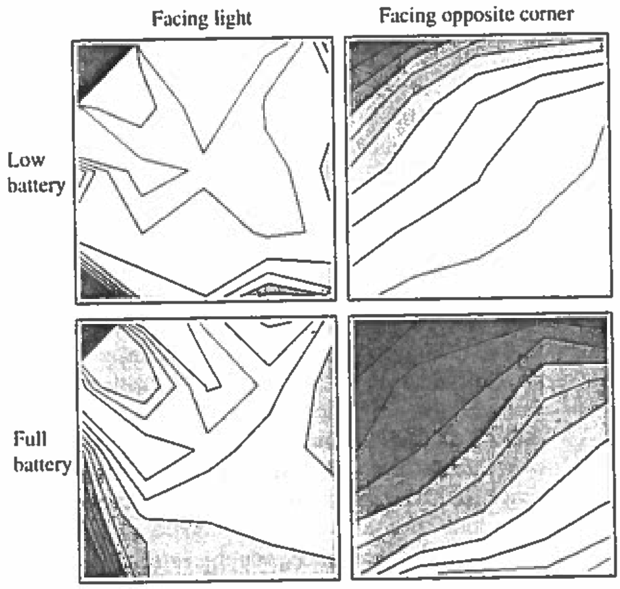

Fig. 76.20 Activation levels (brightness proportional to activation) of an internal neuron plotted over the environment while the robot was positioned at various locations in each of the four conditions (facing recharging area or not, discharged battery or not). The recharging area is located at the *top left corner* of each map

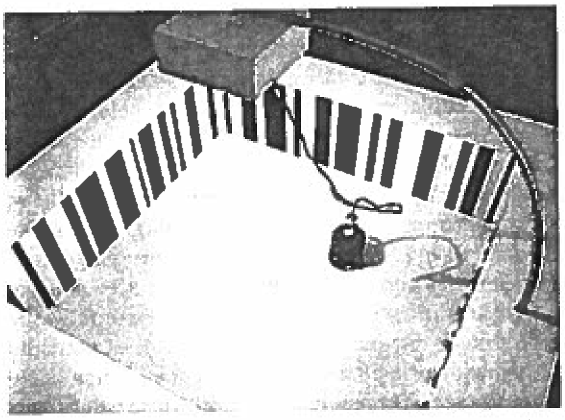

Fig. 76.21 A network of spiking neurons is evolved to drive the vision-based robot in the arena. The light below the rotating contacts allows continuous evolution also overnight

of local contrast detection filters were converted in spikes (the stronger the contrast, the larger the number of spikes per second) sent to ten fully connected spiking neurons implemented according to the spike response model [76.106]. The spike series of a subset of these neurons was translated into motor commands (more spikes per second corresponded to faster rotation of the wheel). The fitness function was the amount of forward translation of the robot measured over 2 min. Consequently robots that turned in place or hit the walls had comparatively lower fitness than robots that could move straight and turn only to avoid walls. The genome of these robots was a bit string that encoded only the sign of the neurons and the presence of synaptic connections. Existing connections were set to 1 and could not change during the lifetime of the robot.

Evolution reliably discovered very robust spiking controllers in approximately 20 generations, approximately 30 h of evolution on the real robot (VIDEO 37). Evolved robots could avoid not only the walls, but any object positioned in front of them. Detailed analysis of the best evolved controllers revealed that neurons did not exploit time differences between spikes, which one would have expected if optic flow was used to detect distance from walls. Instead, they simply used the number of incoming spikes (firing rate) as an indication of when to turn. When the robot perceived a lot of contrast it would go straight, but when the contrast decreased below a certain threshold (indicating that it approached an object), it started to turn away. This extremely efficient and simple result seems to be in contrast with theories of optic flow detection in insects and may be worth considering as an alternative hypothesis for vision-based behavior.

Spiking neural networks turned out to be more evolvable than connectionist models (at least for this task). One possible explanation is that spiking neurons have subthreshold dynamics that, to some extent, can be affected by mutations without immediately affecting the output of the network.

The robust results and compact genetic encoding encouraged the authors to use an even simpler model of spiking neuron so that the entire neural network could be mapped in less than 50 bytes of memory. The evolutionary algorithm was also reduced to a few lines of code and the entire system was implemented within a programmable intelligent computer (PIC) microcontroller without the need for any external computer for data storage. The system was used for a sugar-cube robot (Fig. 76.22) that autonomously and reliably developed the ability to navigate around a maze in less than an hour [76.107]. Interestingly, evolved spiking controllers developed a pattern of connections where spiking neurons received connections from a small patch of neighboring sensors, but not from other sensors, and were connected only to neighboring spiking neurons. This pattern of connectivity is also observed in biological systems and encourages specialization of neurons to sensory features.

76.7.3 GasNets

This section describes another style of artificial neural network strongly inspired by those parts of contemporary neuroscience that emphasize the complex electrochemical nature of real nervous systems. In particular, they make use of an analogue of *volume signaling*, whereby neurotransmitters freely diffuse into a relatively large volume around a nerve cell, potentially affecting many other neurons [76.108, 109]. This exotic form of neural signaling does not sit easily with classical pictures of brain mechanisms and is forcing a radical rethink of existing theory [76.110–113]. The class of artificial neural networks developed to explore artificial volume signaling are known as GasNets [76.114]. These are essentially standard neural networks augmented by a chemical signaling system comprising a diffusing *virtual* gas which can modulate the response of other neurons. A number of GasNet variants, inspired by different aspects of real nervous systems, have been explored in an evolutionary robotics context as artificial nervous systems for mobile autonomous robots. They have been shown to be significantly more evolvable, in terms of speed to a good solution, than other forms of neural networks for a variety of robot tasks and behaviors [76.38, 114–116]. They are being investigated as potentially useful engineering tools and as a way of gaining helpful insights into biological systems [76.112, 117–119].

By analogy with biological neuronal networks, GasNets incorporate two distinct signaling mechanisms, one *electrical* and one *chemical*. The underlying *electrical* network is a discrete-time-step recurrent neural network with a variable number of nodes. These nodes

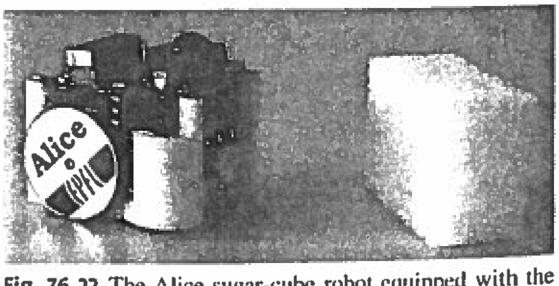

Fig. 76.22 The Alice sugar-cube robot equipped with the evolutionary spiking neural network implemented within its PIC microcontroller

are connected by either excitatory or inhibitory links (Fig. 76.23).

In addition to this underlying network in which positive and negative *signals* flow between units, an abstract process loosely analogous to the diffusion of gaseous modulators is at play. Some units can emit virtual *gases* which diffuse and are capable of modulating the behavior of other units by changing the profile of their output functions. The networks occupy a 2-D space; the diffusion processes mean that the relative positioning of nodes is crucial to the functioning of the network. Spatially, the gas concentration varies as an inverse exponential of the distance from the emitting node with a spread governed by a parameter r with the concentration set to zero for all distances greater than r. The total concentration of gas at a node is determined by summing the contributions from all other emitting nodes.

For mathematical convenience, in the original GasNet there are two *gases*, one whose modulatory effect is to increase the transfer function gain parameter and one whose effect is to decrease it. Thus the gas does not alter the electrical activity in the network directly but rather acts by continuously changing the mapping between input and output for individual nodes, either directly of by stimulating the production of further virtual gas. The general form of the diffusion is based on the properties of a (real) single-source neuron as modeled in detail by *Philippides* et al. [76.112, 117]. The modulation chosen is motivated by what is known of NO modulatory effects at synapses [76.120]. For full details see [76.114].

Various extensions of the basic GasNet have been produced. Two in particular are strongly inspired by contemporary neuroscience. The *plexus model* is directly inspired by a type of signaling seen in the mammalian cerebral cortex in which the NO signal is generated by the combined action of many fine NO-producing fibers, giving a targeted *cloud* which is distant from the neurons from which the fiber plexus emanates [76.118]. In the plexus GasNet, which models this form of signaling at an abstract level, the spatial distribution of gas concentration has been modified to be uniform over the area of affect. The center of this gas diffusion cloud is under genetic control and can be distant from the controlling node (which, by analogy, is the source of the *plexus*) [76.116]. All other details of the models are identical to the original GasNet model, as described earlier. The *receptor* GasNet incorporates an aspect of biological neuronal networks that has no analog in the vast majority of artificial neural networks (ANNs): the role of receptor molecules. Although neuroscience is a long way from a full understanding of receptor mechanisms, a number of powerful systems level ideas can be abstracted.

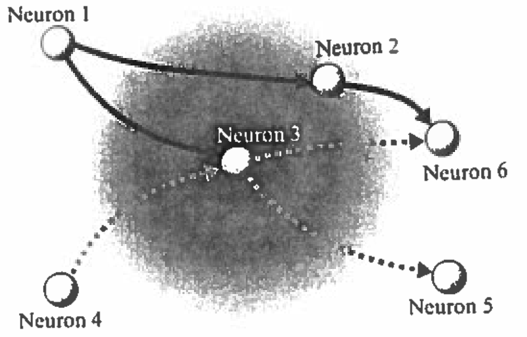

A GasNet. Neuron 3 is emitting gas, and modulating neuron 2 despite there being no synaptic connection.

Fig. 76.23 A basic GasNet showing positive (*solid*) and negative (*dashed*) *electrical* connections and a diffusing virtual gas creating a *chemical* gradient

Details of the receptor variant are similar to the basic GasNet except there is now *only one* virtual gas and each node in the network can have one of three discrete quantities (zero, medium, maximum) of a number of possible receptors. The modulation the diffusing neurotransmitter affects at a neuron depends on which receptors are present. The strength of a modulation at a node is proportional to the product of the gas concentration at the node and the relevant receptor quantity. In the original GasNet, any node that was in the path of a diffusing transmitter would be modulated in a fixed way. The receptor model allows site-specific modulations, including no modulation (no receptors) and multiple modulations at a single site ([76.116] for further details).

Although most of the GasNet variants described in this section have been successfully used in a number of robotic tasks, their evolvability and other properties were thoroughly compared on a version of the (gantry) robot visual discrimination task described in Sect. 76.5, VIDEO 375. All aspects of the networks were under genetic control: the number of nodes, the connectivity and, in the case of the GasNets, all parameters governing volume signaling (including the position of the nodes and whether or not they were *virtual gas* emitters). The visual sampling morphology was also under evolutionary control. The original basic GasNet was found to be significantly more evolvable than a variety of other styles of connectionist neural networks as well as a GasNet with the volume signaling disabled. Successful GasNet controllers for this task tended to be rather minimal, in terms of numbers of nodes and connections, while possessing complex dynamics [76.114].

Later experiments comparing the basic GasNet with the plexus and receptor variants showed the latter two to be considerably more evolvable than the former, with the receptor GasNet being particularly successful [76.116]. These GasNet experiments demonstrated that the intricate network dynamics made possible by the artificial volume signaling mechanisms can be readily harnessed to generate adaptive behaviors in autonomous agents. They also throw up such questions as why GasNets are more evolvable than many other forms of ANN and why there is a difference in evolvability between GasNet variants. In order to gain insight into what factors are most important in GasNet evolvability, several other varieties were studied, including non-spatial GasNets where the diffusion process is replaced by explicit gas connections with complex dynamics and version with other forms of modulation and diffusion [76.119]. Detailed comparative studies of these variants with each other, and with other forms of ANN, were performed using the visual discrimination task described above [76.116, 119].

The comparative studies revealed that the rich dynamics and additional timescales introduced by the gas played an important part in enhanced evolvability, but were not the whole story [76.116, 119]. The particular form of modulation was also important – multiplicative or exponential modulation (in the form of changes to the transfer function) were found to be effective, but additive modulations were not. The former kind of modulations may well confer evolutionary advantages by allowing nodes to be sensitive to different ranges of input (internal and sensory) in different contexts. The spatial embedding of the networks also appears to play a role in producing the most effective coupling between the two distinct signalling processes (*electrical* and *chemical*). By exploiting a loose, flexible coupling between the two processes, it is possible to significantly reduce destructive interference between them, allowing one to be *tuned* against the other while searching for good solutions [76.115, 116, 121]. Similar forces may be at play in spiking neural networks, where sub-threshold and spiking dynamics interact with each other, which although not yet compared to GasNets, were shown to be more evolvable than connectionist networks. Measurements of the degree of coupling in the GasNets variants versus speed of evolution supported this view [76.116]; the receptor GasNet, for which the evolutionary search process has the most direct control over the degree of coupling between the signaling processes, and which has a bias towards a loose coupling, was by far the most evolvable [76.116].

Analysis of GasNet solutions often reveals high levels of degeneracy, with functionally equivalent subnetworks occurring in many different forms, some involving gas and some not [76.121]. Their genotype to phenotype mapping (where the phenotype is robot behavior) is also highly degenerate with many different ways of achieving the same outcome (e.g., moving node positions, changing gas diffusion parameters or adding new connections can all have the same effect). This is especially true when variable length genotypes are used to efficiently sculpt solutions in a search space of variable dimensions. The levels of degeneracy are generally significantly higher than when using connectionist networks. These properties partly explain the robustness and adaptability of GasNets in noisy environments and are another important factor in their evolvability (there are many paths to the same phenotypical outcome with reduced probabilities of lethal mutations) [76.119, 122]. In the most successful varieties of GasNet, multi-scale dynamics, modulation and spatial embedding act in concert to produce highly evolvable degenerate networks.

These and ongoing investigations indicate that explicitly dealing with the electrochemical nature of nervous systems is likely to be an increasingly fruitful area of research, both for evolutionary robotics and for neuroscience, that will likely force us to broaden our notions of what behavior-generating mechanisms might look like.

Because of its ability to explore whole classes of underspecified models, ER is being increasingly used to develop or explore neural models aimed at answering specific questions in neuroscience [76.88–90] or to probe new theories about possible neural mechanisms [76.90]. One intriguing recent hypothesis is that one of the forms of plasticity on which the brain relies is itself a form of evolution via natural selection acting within neural tissue [76.123, 124]. The units of selection in this case are activity and connection patterns which are copied between groups of neurons. Irrespective of whether or not it occurs in nature (and it might), this kind of mechanism could be employed in a whole new kind of evolutionary robotics.

76.8 Evolution and Learning

Evolution and learning (or phylogenetic and ontogenetic adaptation) are two forms of biological adaptation that differ in space and time. Evolution is a process of selective reproduction and substitution based on the existence of a population of individuals displaying variability at the genetic level. Learning, instead, is a set of modifications taking place within each single individual during its own life time. Evolution and learning operate on different time scales. Evolution is a form of adaptation capable of capturing relatively slow environmental changes that might encompass several generations (e.g., the perceptual characteristics of food sources for a given species). Learning, instead, allows an individual to adapt to environmental modifications that are unpredictable at the generational level. Learning might include a variety of mechanisms that produce adaptive changes in an individual during its lifetime, such as physical development, neural maturation, variation of the connectivity between neurons, and synaptic plasticity. Finally, whereas evolution operates on the genotype, learning affects only the phenotype, and phenotypic modifications cannot directly modify the genotype.

Researchers have combined evolutionary techniques and learning techniques (supervised or unsupervised learning algorithm such us reinforcement learning or Hebbian learning; for a review see [76.125]). These studies have been conducted with two different purposes:

1. Identifying the potential advantage of combining these two methods from the point of view of developing robust and effective robots.
2. Understanding the role of the interaction between learning and evolution in nature.

Within an evolutionary perspective, learning has several different adaptive functions. First, it might allows individuals to adapt to changes that occur too quickly to be tracked by evolution [76.126]. Secondly, learning might allows robots to use information extracted during their interaction with environment to develop adaptive characters ontogenetically without necessarily discovering these characters through genetic variations and without encoding these characters in their genome. To understand the importance of this aspect, we should consider that evolutionary adaptation is based on an explicit but concise indication of how well an individual robot coped with its environment – the fitness value of a robot. Ontogenetic adaptation, on the contrary, is based on extremely rich information – the state of the sensors while the robot interacts with its environment. This huge amount of information encodes very indirectly how well an individual is doing in different phases of its lifetime or how it should modify its behavior to increase its fitness. However, evolving robots that have acquired a predisposition to exploit this information to produce adaptive changes during their lifetime might be able to develop adaptive characteristics on the fly, thus leading to the possibility to produce *complex* phenotypes on the basis of parsimonious genotypes. Finally, learning can help and guide evolution. Although physical changes of the phenotype, such as strengthening of synapses during learning, cannot be written back into the genotype, *Baldwin* [76.127] and *Waddington* [76.128] suggested that learning might indeed affect the evolutionary course in subtle but effective ways. Baldwin's argument was that learning accelerates evolution because suboptimal individuals can reproduce by acquiring during life necessary features for survival. However, variation occurring during successive generation might lead to the discovery of genetic traits that lead to the establishment of the same characteristics that were previously acquired thorough lifetime learning. This latter aspect of Baldwin's effect, namely indirect genetic assimilation of learned traits, has been later supported by scientific evidence and defined by *Waddington* [76.128] as a canalization effect.

Learning however, also has costs such as: (1) a delay in the ability to acquire fitness (due to the need to develop fit behavior ontogenetically), and (2) increased unreliability due to the fact that the possibility to develop certain abilities ontogenetically is subjected to partially unpredictable characteristics of the robot–environment interaction [76.129]. In the next two subsections we describe two experiments that show some of the potential advantages of combining evolution and learning.

76.8.1 Learning to Adapt to Fast Environmental Variations

Consider the case of a Khepera robot that should explore an arena surrounded by black or white walls to reach a target placed in a randomly selected location [76.126]. Evolving robots are provided with eight sensory neurons that encode the state of the four corresponding infrared sensors and two motor neurons that control the desired speed of the two wheels. Since the color of the walls change every generation and since the color significantly affects the intensity of the response of the infrared sensors, evolving robots should develop an ability to *infer* whether they are currently located in an environment with white or black walls and learn to modify their behavior during lifetime. That is, robots should avoid walls only when the infrared sensors are almost fully activated in the case of arenas with white walls, while they should avoid walls even when the infrared sensors are slightly activated in the case of arenas with black walls.

Robots were provided with a neural controller (Fig. 76.24) including four sensory neurons that encoded the state of four corresponding infrared sensors;

two motors neurons that encoded the desired speed of the two wheels; and two teaching neurons that encoded the teaching values used to modify the connection weights from the sensory neurons to the motor neurons during the robots' lifetime. This special architecture allows evolving robots to transform the sensory states experienced by the robots during their lifetime into teaching signals that might potentially lead to adaptive variations during lifetime. Analysis of evolved robots revealed that they developed two different behaviors that are adapted to the particular arena where they happen to be *born* (surrounded by white or black walls). Evolving robots did not inherit an ability to behave effectively, but rather a predisposition to learn to behave. This predisposition to learn involves several aspects such as a tendency to experience useful learning experiences, a tendency to acquire useful adaptive characters through learning, and a tendency to channel variations toward different directions in different environmental conditions [76.126].

76.8.2 Evolution of Learning

In the previous example, the evolutionary neural network learned using a standard learning rule that was applied to all synaptic connections. *Floreano* and *Mondada* [76.130] explored the possibility of genetically encoding and evolving the learning rules associated to the different synaptic connections of a neural network embedded in a real robot. The main motivation of this line of work was to evolve robots capable of adapting to a partially unknown environment, rather than robots adapted to the environment(s) seen during evolution. In order to prevent evolutionary tuning of the neural network to the specificities of the evolutionary environment (which would limit transfer to different environments or transfer from simulation to reality), the synaptic weight values were not genetically encoded. Instead, each synaptic connection in the network was described by three genes that defined its sign, its learning rule, and its learning rate (Fig. 76.25). Every time a genome was decoded into a neural network and downloaded onto the robot, the synaptic strengths were initialized to small random values and could change according to the genetically specified rules and rates while the robot interacted with the environment. Variations of this methodology included a more compact genetic encoding where the learning properties were associated to a neuron instead of a synapse. All synapses afferent to a neuron used its genetically specified rules and rates. Genes could encode four types of Hebbian learning that were modeled upon neurophysiological data and were complementary to each other [76.131].

Experimental results in a nontrivial, multitask environment (Fig. 76.26, VIDEO 40) indicated that this methodology has a number of significant advantages with respect to the evolution of synaptic strengths without learning [76.48]. Robots evolved faster and obtained better fitness values. Furthermore, evolved behaviors were qualitatively different, notably in that they did not exploit minimal solutions tuned to the environment (such as turning only on one side, or turning in circles tuned to the dimensions of the evolutionary arena). Most important, these robots displayed remarkable adaptive properties after evolution. Best evolved individuals: (1) transferred perfectly from simulated

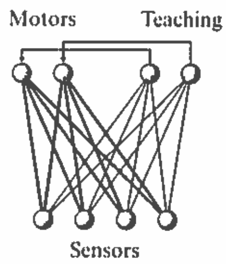

Fig. 76.24 A self-teaching network. The output of the two teaching neurons is used as a teaching value for the two motor neurons. The weights that connect the sensory neurons to the teaching neurons do not vary during the robots' lifetime while the weights that connect the sensory neurons to the motor neurons are modified with an error-correction algorithm

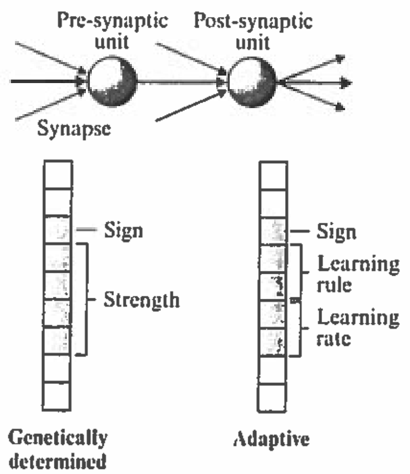

Fig. 76.25 Two methods for genetically encoding a synaptic connection. Genetically determined synapses cannot change during the lifetime of the robot. Adaptive synapses instead are randomly initialized and can change during lifetime of the robot according to the learning rules and rates specified in the genome

to physical robots, (2) accomplished the task when the light and reflection properties of the environment were modified, (3) accomplished the task when key landmarks and target areas of the environment were displaced, and (4) transferred well across morphologically different robotic platforms. In other words, these robots were selected for their ability to solve a partially unknown problem by adapting on the fly, rather than for being a solution to the problem seen during evolution.

In further experiments where the genetic code for each synapse of the network included one gene whose value caused its remaining genes to be interpreted as connection strengths or learning rules and rates, 80% of the synapses *made the choice* of using learning, reinforcing the fact that this genetic strategy has a comparatively stronger adaptive power [76.131]. This methodology could also be used to evolve the morphology of neural controllers were synapses are created at runtime and therefore their strengths cannot be genetically specified [76.132]. Recently, the adaptive properties of this type of adaptive genetic encoding were confirmed also in the context of evolutionary spiking neurons for robot control [76.133].

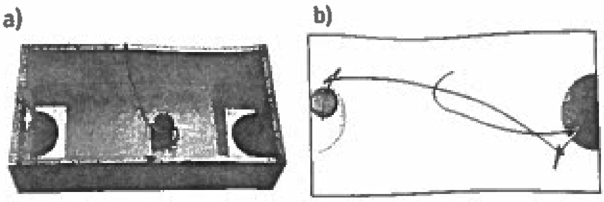

Fig. 76.26 (a) A mobile robot Khepera equipped with a vision module can gain fitness points by staying on the *grey area* only when the light is on. The light is normally off, but it can be switched on if the robot passes over the *black area* positioned on the other side of the arena. The robot can detect ambient light and wall color, but not the color of the floor. (b) Behavior of an individual evolved in simulation with genetic encoding of learning rules

76.9 Evolution of Social Behavior

In the previous sections, we limited our analysis to individual behaviors, i. e., to the evolution of robots placed in an environment that does not include other robots. The evolutionary method, however, can also be applied to evolve social behaviors in which multiple robots situated in the same environment interact between themselves in cooperative or competitive manners.

As we will see, competitive co-evolution is particularly interesting from the point of view of synthesizing progressively more complex capacities and from the point of view of developing solutions that are robust with respect to environmental variations. Cooperative evolution instead is particularly interesting for the possibility to solve problems that cannot be handled by a single robot, because of physical constraints or limited behavioral capabilities [76.134] and to develop solutions that are robust.

76.9.1 Coevolving Predator and Prey Robots

Competitive coevolution, for example the coevolution of two populations of predator and prey robots that are evolved for the ability to catch prey and to escape predators, respectively, has two characteristics that are particularly interesting from an evolutionary robotics perspective. The first aspect is that the competition between populations with different interests might spontaneously lead to a sort of incremental evolutionary process where evolving individuals are faced with progressively more complex challenges (although this is not necessarily the case). Indeed, in initial generations the task of the two populations is relatively simple because opponents have simple and poorly developed abilities on average. After a few generations, however, the abilities of the two populations increase and, consequently, the challenges for each population become more difficult. The second aspect consists of the fact that the environment varies across generations because it includes other coevolving individuals. This implies that coevolving individuals should be able to adapt to ever-changing environments and to develop behaviors that are robust with respect to environmental variations [76.135].

The potential advantages of competitive coevolution for evolutionary robotics have been demonstrated by a set of experiments conducted by *Floreano* and *Nolfi* [76.136, 137] where two populations of robots were evolved for the ability to catch prey and escape predators, respectively (Fig. 76.27).

The results indicated that both predator and prey robots tended to vary their behavior throughout generations without converging on a stable strategy. The behavior displayed by individuals at each generation tended to be tightly adapted to the counter-strategy exhibited by the opponent of the same generation (VIDEO 3[illegible]). This evolutionary dynamic however does not really lead to long-lasting progress because, after an initial evolutionary phase, the coevolutionary

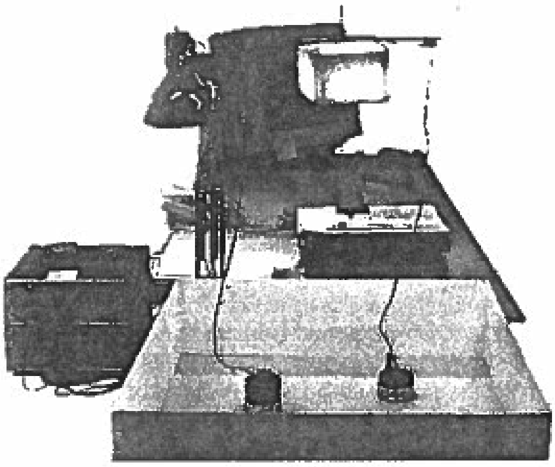

Fig. 76.27 Experimental setup. The predator and prey robot (*from left to right*) are placed in an arena surrounded by walls and are allowed to interact for several trials starting at different randomly generated orientations. Predators are selected on the basis of the percentage of trials in which they are able to catch (i.e., to touch) the prey, and prey on the basis of the percentage of trials in which they were able to escape (i.e., to not be touched by) predators. Predators have a vision system, whereas the prey have only short-range distance sensors, but can go twice as fast as the predator. Collision between the robots is detected by a conductive belt at the base of the robots

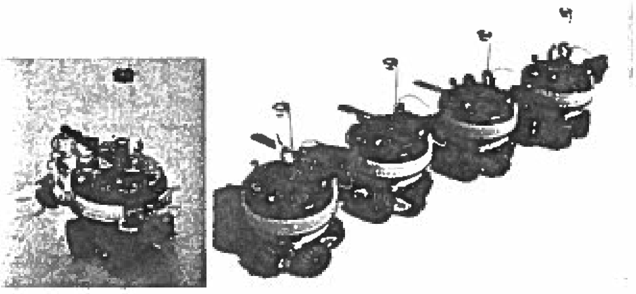

Fig. 76.28 An s-bot and a simulated swarm-bot consisting of four s-bots assembled in chain formation

process led to a limit cycle dynamic where the same small set of behavioral strategies recycled over and over again along generations [76.137]. This limit cycle dynamic can be explained by considering that prey robots tended to vary their behavior in order to disorient predators as soon as predators become effective against the current behavioral strategies exhibited by prey robots.

However, experiments [76.138] where robots were allowed to change their behavior on the fly on the basis of unsupervised Hebbian learning rules showed that the evolutionary phase where coevolving robots were able to produce real progress was significantly longer, and evolved predators displayed an ability to effectively cope with prey exhibiting different behavioral strategies by adapting their behavior on the fly to the prey's behavior. Prey instead tended to display behavior that changed in unpredictable ways.

Further experiments showed that competitive coevolution can solve problem that the evolution of a single population cannot. *Nolfi* and *Floreano* [76.137] demonstrated that the attempt to evolve predators robot for the ability to catch a fixed pre-evolved prey produced lower performance with respect to control experiments where predators and prey were coevolved at the same time.

76.9.2 Evolving Cooperative Behavior

As testified by social insects, colonies of simple cooperating individuals can display remarkable capacities and exhibit self-organising behaviors in which the spatio-temporal pattern observed at the system level emerge from numerous interactions among the individual robots. On the other hand, designing collective robotic systems of this sort constitutes a difficult problem due to the indirect relationship between the desired group behavior and the characteristics of the individual robots. By evaluating the robotic system as a whole (i.e., by selecting the robots on the basis of the global behavior that emerge from a large number of robot/environmental and robot/robot interactions), Evolutionary Robotics provides a means for discovering effective behavioral solutions and simple and robust control policies [76.139].

Recent research showed that teams of evolved robots can:

1. Develop robust and effective coordinated behavior [76.140, 141]
2. Collaborate by assuming complementary role [76.141, 142] (VIDEO 376)
3. Display self-organizing properties [76.143]
4. Develop and use communicative capabilities [76.144–146].

Moreover, some of the research carried in this area demonstrated how evolutionary robotics experiments can contribute to model biological phenomena, e.g., to identify the evolutionary conditions that enable the emergence of cooperative communicative behaviors [76.146] or the mechanisms enabling the evolution of effective division of labour strategies [76.147]

Here we briefly review a series of experiments where swarm-bots [76.148], i. e., groups of autonomous robots capable of assembling by physically connecting together, were evolved for the ability to display coordinated motion (VIDEO 115). Each individual robot consisted of a main platform (chassis) and turret that could actively rotate with respect to each other (Fig. 76.28). The chassis included tracks with teethed wheels for navigation on both rough and flat terrain, and infrared sensors pointing to the ground. The turret included a gripper, sixteen light-sensors distributed around the body, a loudspeaker, three microphones, and a traction sensor placed between the turret and the chassis to detect the direction and the intensity of the traction force that the turret exerts on the chassis. Swarm-bots were formed by several robots provided with identical neural controllers and assembled together so to form a single physical entity.

By evolving the neural controllers of these Swarm-bots, *Baldassare* et al. [76.140] demonstrated how the robots can display a robust and effective coordinated capacities that allow the individuals to negotiate and converge on a coherent direction and to keep moving along that direction by compensating the disalignments originating during motion. Such behavioral capacity was robust enough to allow a smooth transfer from simulation to reality and to allow the robots to to generalize their capacity to rugged terrains. In an extended experiments in which the s-bots were also equipped with infrared sensors, speakers and microphones, the evolved swarm-bots also showed a capacity to avoid dangers (e.g., holes) by coordinately changing direction as soon as one s-bot detected a hole [76.149].

Evolved swarm-bots generalized their coordinated motion capabilities also when they were tested in different conditions (e.g., when they were assembled in much more numerous groups and/or in different topologies, or when they had to also carry heavy objects by pushing and pulleding them in a coordinated manner). Finally, when placed in new environmental conditions (e.g., in environment with obstacle and walls), the swarm-bots spontaneously displayed new behavioral skills (related to acquired skills), such as the ability to cooperatively avoid obstacles, without any further adaptation. This ability to display new related behaviors, in new behavioral conditions, emerged as a result of the dynamical process originating from the interaction of the same robots with the new environmental conditions [76.142, 150].

76.9.3 Evolution of Communication

Communication represents a key aspect in collective behaviors. Recent research in evolutionary robotics has demonstrated how sophisticated communication capabilities can emerge and evolve in population of robots selected for the ability to perform tasks requiring coordination and/or cooperation.

The analysis of these experiments indicate that communicative interactions often originate as the result of cues, that provide useful information to other robots, produced inadvertently during the execution of specific behaviors [76.146, 151]. The presence of these cues create the basis for the development of an ability to react to them in an adaptive way thus leading to the establishment of adaptive communicative interactions in which robots produce signals and react to detected signals adaptively. The establishment of these forms of communicative interactions then create the adaptive conditions for the co-evolution of signalling and response strategies [76.152] (VIDEO 117).

The reliability and stability of the resulting communication system depend on the level of relatedness (i. e., genetic similarity) between robots and the level at which they were selected [76.146]. Robots that are genetically higly related or that are selected on the basis of the behavior exhibited by the group evolve reliable signals and stable communicative conventions. In contrast, when relatedness between robots is low and selection is acting at the level of the individuals, the evolutionary process might lead to the emergence of instable, ineffective and in some case deceptive communication forms [76.153, 154].

The evolution of communication is strongly interlinked with the evolution of other behavioral capacities [76.155]. Indeed, after all, robots need to develop appropriate behaviors to access and/or generate the information to be communicated and/or to react to detected signals appropriately. The co-adaptation of behavioral and communicative skills might lead to prolonged innovation phases in which the development of behavioral capacities create the adaptive conditions for the development of communication capacities and vice versa [76.151, 152]. Moreover, the co-adaptation of behavioral and communication capacities tend to lead to highly contingent evolutionary processes in which the capacities possessed by the population at a certain evolutionary phase strongly influence the outcome of the successive phases [76.152, 156] (VIDEO 117).

76.10 Evolutionary Hardware

In recent years, technology advancements have allowed researchers to explore evolution of electronic circuits. In this section, we briefly summarize some foundational work in this direction.

76.10.1 Evolvable Hardware Robot Controllers

In most of the work discussed so far some form of genetically specified neural network, implemented in software, has been at the center of the robot control system. Work on a related approach of evolving control systems directly onto hardware dates back to *Thompson*'s work in the mid 1990s [76.157]. In contrast to hardware controllers that are designed or programmed to follow a well-defined sequence of instructions, evolved hardware controllers are directly configured by evolution and then allowed to behave in real time according to semiconductor physics. By removing standard electronics design constraints, the physics can be exploited to produce highly nonstandard and often very efficient and minimal systems [76.158].

Thompson [76.157] used artificial evolution to design an onboard hardware controller for a two-wheeled autonomous mobile robot engaged in simple wall-avoidance behavior in an empty arena. Starting from a random orientation, and position near the wall, the robot had to move to the center of the arena and stay there using limited sensory input (Fig. 76.29). The direct current (DC) motors driving the wheels were not allowed to run in reverse and the robot's only sensors were a pair of time-of-flight sonars rigidly mounted on the robot, pointing left and right.

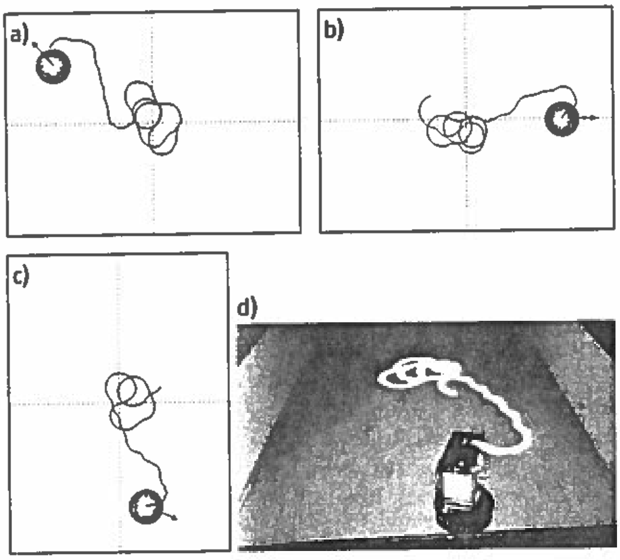

Fig.76.29a–d Wall-avoidance behavior of a robot with an evolved hardware controller in virtual reality **(a–c)** and the real world **(d)**

Thompson's approach made use of a so-called dynamic state machine (DSM) – a kind of generalized read-only memory (ROM) implementation of a finite-state machine where the usual constraint of strict synchronization of input signals and state transitions are relaxed (in fact put under evolutionary control). The system had access to a global clock whose frequency was also under genetic control. Thus evolution determined whether each signal was synchronized to the clock or allowed to flow asynchronously. This allowed the evolving DSM to be tightly coupled to the dynamics of interaction between the robot and environment and for evolution to explore a wide range of systems dynamics. The process took place within the robot in a kind of *virtual reality* in the sense that the real evolving hardware controlled the real motors, but the wheels were just spinning in the air. The movements that the robot would have actually performed if the wheels had been supporting it were then simulated and the sonar echo signals that the robot was expected to receive were supplied in real time to the hardware DSM. Excellent performance was attained after 35 generations, with good transfer from the virtual environment to the real world (Fig. 76.29).

Shortly after this research was performed, particular types of field programmable gate arrays (FPGAs) which were appropriate for evolutionary applications became available. FPGAs are reconfigurable systems allowing the construction of circuits built from basic logic elements. Thompson exploited their properties to demonstrate evolution directly in the chip. By again relaxing standard constraints, such as synchronizing all elements with a central clock, he was able to develop very novel forms of functional circuits, including a controller for a Khepera robot using infrared sensors to avoid obstacles [76.158, 159].

Following Thompson's pioneering work, *Keymeulen* et al. evolved a robot control system using a Boolean function approach implemented on gate-level evolvable hardware [76.160]. This system acted as a navigation system for a mobile robot capable of locating and reaching a colored ball while avoiding obstacles. The robot was equipped with infrared sensors and an vision system giving the direction and distance to the target. A programmable logic device (PLD) was used to implement a Boolean function in its disjunctive form. This work demonstrated that such gate-level evolvable hardware was able to take advantage of the correlations in the input states and to exhibit useful generalization abilities, thus allowing the evolution of robust behavior in simulation followed by a good transfer into the real world.

In a rather different approach, *Ritter* et al. used an FPGA implementation of an onboard evolutionary algorithm to develop a controller for a hexapod robot [76.161]. *Roggen* et al. devised a multicellular reconfigurable circuit capable of evolution, self-repair, and adaptation [76.162], and used it as a substrate for evolving spiking controllers of a wheeled robot [76.163]. Although evolved hardware controllers are not widely used in evolutionary robotics, they still hold out the promise of some very useful properties, such as robustness to faults, which make them interesting for extreme condition applications such as space robotics.

76.11 Closing Remarks

Evolutionary robotics is a young and integrated approach to robot development without human intervention where machines change and adapt by capitalizing on the interactions with their environment. Despite initial skepticism by mainstream and applied robotics practitioners and even by pioneers of this approach [76.164], over the years the field has been constantly growing with new methods and approaches for evolving more complex, efficient, and sometimes surprising robotic systems. In some areas, such as morphology and self-assembly, evolutionary robotics is still the most widely used and powerful approach.

Evolutionary robotics is not only a method for automatic robot development inspired by biology, but also a tool for investigating open questions in biology concerning evolutionary, developmental, and brain dynamics. Its richness and fecundity make us believe that this approach will continue to grow and progress towards the creation of a new species of machines capable of self-evolution.

To gain a pratical knowledge, interested readers might use software libraries such as framework for autonomous robotics simulation and analysis (FARSA) [76.165], an open-software tool that permit to carry on evolutionary robotics experiments based on a variety of robotic platforms and to replicate and vary some of the experiments described in this chapter [76.166].

Video-References

VIDEO 36 Visual navigation of mobile robot with pan-tilt camera
available from http://handbookofrobotics.org/view-chapter/76/videodetails/36

VIDEO 37 Visual navigation with collision avoidance
available from http://handbookofrobotics.org/view-chapter/76/videodetails/37

VIDEO 38 Coevolved predator and prey robots
available from http://handbookofrobotics.org/view-chapter/76/videodetails/38

VIDEO 39 Evolution of collision-free navigation
available from http://handbookofrobotics.org/view-chapter/76/videodetails/39

VIDEO 40 Online learning to adapt to fast environmental variations
available from http://handbookofrobotics.org/view-chapter/76/videodetails/40

VIDEO 41 iCub language comprehension
available from http://handbookofrobotics.org/view-chapter/76/videodetails/41

VIDEO 114 Resilent machines through continuous self-modeling
available from http://handbookofrobotics.org/view-chapter/76/videodetails/114

VIDEO 115 A swarm-bot of eight robots displaying coordinated motion
available from http://handbookofrobotics.org/view-chapter/76/videodetails/115

VIDEO 116 Discrimination of objects through sensory-motor coordination
available from http://handbookofrobotics.org/view-chapter/76/videodetails/116

VIDEO 117 Evolution of cooperative and communicative behaviors
available from http://handbookofrobotics.org/view-chapter/76/videodetails/117

VIDEO 118 Exploration and homing for battery recharge
available from http://handbookofrobotics.org/view-chapter/76/videodetails/118

VIDEO 119 Introduction to evolutionary robotics at EPFL
available from http://handbookofrobotics.org/view-chapter/76/videodetails/119
VIDEO 371 Evolution of visually guided behavior on Sussex gantry robot
available from http://handbookofrobotics.org/view-chapter/76/videodetails/371
VIDEO 372 Evolved walking in an Octpod
available from http://handbookofrobotics.org/view-chapter/76/videodetails/372
VIDEO 373 Evolved homing walk on rough ground
available from http://handbookofrobotics.org/view-chapter/76/videodetails/373
VIDEO 374 Evolved bipedal walking
available from http://handbookofrobotics.org/view-chapter/76/videodetails/374
VIDEO 375 Evolved GasNet visualization
available from http://handbookofrobotics.org/view-chapter/76/videodetails/375
VIDEO 376 Evolved group coordination
available from http://handbookofrobotics.org/view-chapter/76/videodetails/376
VIDEO 771 Morphological change in an autonomous robot
available from http://handbookofrobotics.org/view-chapter/76/videodetails/771
VIDEO 772 More complex robots evolve in more complex environments
available from http://handbookofrobotics.org/view-chapter/76/videodetails/772

References

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

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

76.3 A.M. Turing: Computing machinery and intelligence, Mind **LIX 236**, 433–460 (1950)

76.4 V. Braitenberg: *Vehicles. Experiments in Synthetic Psychology* (MIT, Cambridge 1984)

76.5 R.D. Beer: *Intelligence as Adaptive Behavior: An Experiment in Computational Neuroethology* (Academic, Boston 1990)

76.6 D. Parisi, F. Cecconi, S. Nolfi: Econets: Neural networks that learn in an environment, Network **1**, 149–168 (1990)

76.7 P. Husbands, I. Harvey: Evolution versus design: Controlling autonomous robots, Integrating Percept. Plan. Action, Proc. 3rd IEEE Annu. Conf. Artif. Intell. Simul. Plan. (1992) pp. 139–146

76.8 D. Floreano, O. Miglino, D. Parisi: Emergent complex behaviors in ecosystems of neural networks. In: *Parallel Architectures and Neural Networks*, ed. by E. Caianiello (World Scientific, Singapore 1991)

76.9 R.A. Brooks: Intelligence without representation, Artif. Intell. **47**, 139–159 (1991)

76.10 F. Mondada, E. Franzi, P. Ienne: Mobile robot miniaturization: A tool for investigation in control algorithms, Proc. 3rd Int. Symp. Exp. Robotics, Tokyo, ed. by T. Yoshikawa, F. Miyazaki (1993) pp. 501–513

76.11 L. Steels (Ed.): *The Biology and Technology of Intelligent Autonomous Agents*, NATO ASI (Springer, Berlin, Heidelberg 1995)

76.12 D. Floreano, F. Mondada: Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot, Proc. 3rd Int. Conf. Simul. Adapt. Behav.: Anim. Animat. 3, ed. by D. Cliff, P. Husbands, J.A. Meyer, S.W. Wilsonpages (MIT, Cambridge 1994) pp. 402–410

76.13 I. Harvey, P. Husbands, D.T. Cliff: Seeing the light: Artificial evolution, real vision, Proc. 3rd Int. Conf. Simul. Adapt. Behav.: Anim. Animat. 3, ed. by D.T. Cliff, P. Husbands, J.-A. Meyer, S. Wilson (MIT, Cambridge 1994) pp. 392–401

76.14 M.A. Lewis, A.H. Fagg, A. Solidum: Genetic programming approach to the construction of a neural network for a walking robot, Proc. IEEE Int. Conf. Robotics Autom. (ICRA) (1992) pp. 2618–2623

76.15 D. Cliff, I. Harvey, P. Husbands: Explorations in evolutionary robotics, Adapt. Behav. **2**, 73–110 (1993)

76.16 D.E. Goldberg: *Genetic Algorithms in Search, Optimization and Machine Learning* (Addison-Wesley, Reading City 1989)

76.17 H. de Garis: Genetic programming: Evolution of time dependent neural network modules which teach a pair of stick legs to walk, Proc. 9th Eur. Conf. Artif. Intell. (ECAI), Stock. (1990) pp. 204–206

76.18 R.D. Beer, J.C. Gallagher: Evolving dynamical neural networks for adaptive behavior, Adapt. Behav. **1**, 94–110 (1992)

76.19 R.D. Beer, H.J. Chiel, L.S. Sterling: Heterogeneous neural networks for adaptive behavior in dynamic environments. In: *Neural Information Processing Systems*, Vol. 1, ed. by D. Touretzky (Morgan Kauffman, San Mateo 1989) pp. 577–585

76.20 M.A. Lewis, A.H. Fagg, G. Bekey: Genetic algorithms for gait synthesis in a hexapod robot. In: *Recent Trends in Mobile Robots*, ed. by Y. Zheng (World Scientific, Singapore 1994) pp. 317–331

76.21 J. Gallagher, R. Beer, M. Espenschiel, R. Quinn: Application of evolved locomotion controllers to a hexapod robot, Robotics Auton. Syst. **19**(1), 95–103 (1996)

76.22 R.D. Beer, R.D. Quinn, H.J. Chiel, R.E. Ritzmann: Biologically inspired approaches to robotics, Commun. ACM **40**, 31–38 (1997)

76.23 S. Galt, B.L. Luk, A.A. Collie: Evolution of smooth and efficient walking motions for an 8-legged robot, Proc. 6th Eur. Workshop Learn. Robots, Brighton (1997)

76.24 T. Gomi, K. Ide: Emergence of gaits of a legged robot by collaboration through evolution, IEEE World Congr. Comput. Intell. (IEEE Press, New York 1998)

76.25 F. Gruau: Automatic definition of modular neural networks, Adapt. Behav. 3(2), 151–183 (1995)

76.26 F. Gruau, K. Quatramaran: Cellular encoding for interactive evolutionary robotics, Proc. 4th Eur. Conf. Artif. Life, ed. by P. Husbands, I. Harvey (MIT, Cambridge 1997) pp. 368–377

76.27 J. Kodjabachian, J.A. Meyer: Evolution and development of neural networks controlling locomotion, gradient following and obstacle avoidance in artificial insects, IEEE Trans. Neural Netw. **9**, 796–812 (1998)

76.28 N. Jakobi: Running across the reality gap: Octopod locomotion evolved in a minimal simulation, Lect. Notes Comput. Sci. **1468**, 39–58 (1998)

76.29 R. Téllez, C. Angulo, D. Pardo: Evolving the walking behavior of a 12 DOF quadruped using a distributed neural architecture, Lect. Notes Comput. Sci. **3853**, 5–19 (2006)

76.30 T. Reil, P. Husbands: Evolution of central pattern generators for bipedal walking in real-time physics environments, IEEE Trans. Evol. Comput. 6(2), 10–21 (2002)

76.31 NaturalMotion: http://www.naturalmotion.com

76.32 B. von Haller, A.J. Ijspeert, D. Floreano: Co-evolution of structures and controllers for Neubot underwater modular robots, Lect. Notes Comput. Sci. **3630**, 189–199 (2005)

76.33 E. Vaughan, E.A. Di Paolo, I. Harvey: The evolution of control and adaptation in a 3D powered passive dynamic walker, Proc. 9th Int. Conf. Simul. Synth. Living Syst. Artif. Life IX, ed. by J. Pollack, M. Bedau, P. Husbands, T. Ikegami, R. Watson (MIT, Cambridge 2004) pp. 139–145

76.34 T. McGeer: Passive walking with knees, Proc. IEEE Conf. Robotics Autom. (ICRA) (1990) pp. 1640–1645

76.35 S. Wischmann, F. Passeman: From passive to active dynamic 3D bipedal walking – An evolutionary approach, Proc. 7th Int. Conf. Climbing Walk. Robots (CLAWAR 2004), ed. by M. Armada, P. González de Santos (Springer, Berlin, Heidelberg 2005) pp. 737–744

76.36 E. Vaughan, E.A. Di Paolo, I. Harvey: The tango of a load balancing biped, Proc. 7th Int. Conf. Climbing Walk. Robots (CLAWAR), ed. by M. Armada, P. González de Santos (2005)

76.37 K. Endo, F. Yamasaki, T. Maeno, H. Kitano: A method for co-evolving morphology and walking pattern of biped humanoid robot, Proc. IEEE Int. Conf. Robotics Autom. (ICRA) (2002) pp. 2775–2780

76.38 G. McHale, P. Husbands: Quadrupedal locomotion: Gasnets, CTRNNs and hybrid CTRNN/PNNs compared, Proc. 9th Int. Conf. Simul. Synth. Living Syst. (Artif. Life IX), ed. by J. Pollack, M. Bedau, P. Husbands, T. Ikegami, R. Watson (MIT, Cambridge 2004) pp. 106–112

76.39 G. McHale, P. Husbands: GasNets and other evolvable neural networks applied to bipedal locomotion, Proc. 8th Int. Conf. Simul. Adapt. Behav.: Anim. Animat. 8, ed. by S. Schaal (MIT, Cambridge 2004) pp. 163–172

76.40 J.F. Laszlo, M. van de Panne, E. Fiume: Limit cycle control and its application to the animation of balancing and walking, Proc. 23rd Annu. Conf. Comp. Graph. Interact. Tech., ACM (1996) pp. 155–162

76.41 R.A. Brooks: Artificial life and real robots, Proc. 1st Eur. Conf. Artif. Life., Toward a Pract. Auton. Syst., ed. by F.J. Varela, P. Bourgine (MIT, Cambridge 1992) pp. 3–10

76.42 R. Featherstone, D. Orin: Robot dynamics: Equations and algorithms, Proc. IEEE Int. Conf. Robotics Autom. (ICRA) (2000) pp. 826–834

76.43 N. Jakobi, P. Husbands, I. Harvey: Noise and the reality gap: The use of simulation in evolutionary robotics, Lect. Notes Comput. Sci. **929**, 704–720 (1995)

76.44 O. Miglino, H.H. Lund, S. Nolfi: Evolving mobile robots in simulated and real environments, Artif. Life **2**, 417–434 (1996)

76.45 N. Jakobi: Half-baked, ad-hoc and noisy: Minimal simulations for evolutionary robotics, Proc. 4th Eur. Conf. Art. Life, ed. by P. Husbands, I. Harvey (MIT, Cambridge 1997) pp. 348–357

76.46 J.C. Bongard, H. Lipson: Nonlinear system identification using coevolution of models and tests, IEEE Trans. Evol. Comput. 9(4), 361–384 (2005)

76.47 S. Koos, J. Mouret, S. Doncieux: Crossing the reality gap in evolutionary robotics by promoting transferable controllers, Proc. 12th Annu. Conf. Genetic Evol. Comput. ACM (2010) pp. 119–126

76.48 J. Urzelai, D. Floreano: Evolution of adaptive synapses: Robots with fast adaptive behavior in new environments, Evol. Comput. 9, 495–524 (2001)

76.49 H.R. Maturana, F.J. Varela: *Autopoiesis and Cognition: The Realization of the Living* (Reidel, Dordrecht 1980)

76.50 R.D. Beer: A dynamical systems perspective on agent-environment interaction, Artif. Intell. **72**, 173–215 (1995)

76.51 P. Funes, B. Orme, E. Bonabeau: Evolving emergent group behaviors for simple humans agents, Proc. 7th Eur. Conf. Artif. Life, ed. by J. Dittrich, T. Kim (Springer, Berlin, Heidelberg 2003) pp. 76–89

76.52 S. Nolfi: Behavior and cognition as a complex adaptive system: Insights from robotic experiments. In: *Philosophy of Complex Systems*, ed. by C. Hooker (Elsevier, Amsterdam 2009) pp. 443–466

76.53 S. Nolfi: Power and limits of reactive agents, Neurocomputing **42**, 119–145 (2002)

76.54 E. Tuci, T. Ferrauto, A. Zeschel, G. Massera, S. Nolfi: An Experiment on behaviour generalisation and the emergence of linguistic compositionality in evolving robots, IEEE Trans. Auton. Mental Dev. **3**, 176–189 (2011)

76.55 C. Scheier, R. Pfeifer, Y. Kunyioshi: Embedded neural networks: Exploiting constraints, Neural Netw. **11**, 1551–1596 (1998)

76.56 S. Nolfi, D. Marocco: Active perception: A sensorimotor account of object categorization, Proc. 7th Int. Conf. Simul. Adapt. Behav.: Anim. Animat. 7, ed. by B. Hallam, D. Floreano, J. Hallam, G. Hayes, J.-A. Meyer (MIT, Cambridge, MA 2002) pp. 266–271

76.57 E. Tuci, G. Massera, S. Nolfi: Active categorical perception of object shapes in a simulated anthropomorphic robotic arm, IEEE Trans. Evol. Comput. **14**, 885–899 (2010)

76.58 S. Collins, A. Ruina, R. Tedrake, M. Wisse: Efficient bipedal robots based on passive-dynamic walkers, Science **307**(5712), 1082–1085 (2005)

76.59 J.C. Bongard: Innocent until proven guilty: Reducing robot shaping from polynomial to linear time, IEEE Trans. Evol. Comput. **15**(4), 571–585 (2011)

76.60 H. Lipson, J.B. Pollack: Automatic design and manufacture of artificial lifeforms, Nature **406**, 974–978 (2000)

76.61 K. Sims: Evolving 3D morphology and behaviour by competition, Artif. Life **1**(4), 28–39 (1994)

76.62 Karl Sims: Evolved virtual creatures, evolution simulation, https://www.youtube.com/watch?v=JBgG_VSP7f8 (1994)

76.63 P. Funes, J. Pollack: Evolutionary body building: Adaptive physical designs for robots, Artif. Life **4**(4), 337–357 (1998)

76.64 Golem Evolutionary Robotics: https://www.youtube.com/watch?v=sLtXXFw_q8c&playnext=1&list=PL396A15596535B451&feature=results_video

76.65 J. Long: *Darwin's devices: What evolving robots can teach us about the history of life and the future of technology* (Basic Books, New York 2012)

76.66 A.J. Clark, J.M. Moore, J. Wang, X. Tan, P.K. McKinley: Evolutionary design and experimental validation of a flexible caudal fin for robotic fish, Artif. Life **13**, 325–332 (2012)

76.67 J. Bongard: Morphological change in machines accelerates the evolution of robust behavior, Proc. Natl. Acad. Sci. **108**(4), 1234–1239 (2011)

76.68 M. Dorigo, M. Colombetti: *Robot shaping: An experiment in behavior engineering* (MIT, Cambridge 1997)

76.69 J.E. Auerbach, J.C. Bongard: On the relationship between environmental and morphological complexity in evolved robots, Proc. 14th Int. Conf. Genetic Evol. Comput. Conf., ACM (2012) pp. 521–528

76.70 J. Hiller, H. Lipson: Automatic design and manufacture of soft robots, IEEE Trans. Robotics **28**(2), 457–466 (2012)

76.71 Evolved Soft Robots: https://www.youtube.com/watch?v=RrgZoo1-z_Y

76.72 J. Bongard, V. Zykov, H. Lipson: Resilient machines through continuous self-modeling, Science **314**(5802), 1118–1121 (2006)

76.73 J.C. Bongard: Accelerating self-modeling in cooperative robot teams, IEEE Trans. Evol. Comput. **13**(2), 321–332 (2009)

76.74 K.J. Kim, H. Lipson: Towards a theory of mind in simulated robots, Proc. 11th Annual Conf. Companion Genetic Evol. Comput. Conf. Late Break. Pap. ACM (2009) pp. 2071–2076

76.75 I. Harvey, P. Husbands, D.T. Cliff, A. Thompson, N. Jakobi: Evolutionary robotics: The Sussex approach, Robotics Auton. Syst. **20**, 205–224 (1997)

76.76 P. Husbands, I. Harvey, D. Cliff, G. Miller: Artificial evolution: A new path for AI?, Brain Cogn. **34**, 130–159 (1997)

76.77 N. Jakobi: Evolutionary robotics and the radical envelope of noise hypothesis, Adapt. Behav. **6**, 325–368 (1998)

76.78 K.O. Stanley, R. Miikkulainen: Evolving neural networks through augmenting topologies, Evol. Comput. **10**(2), 99–127 (2002)

76.79 M.A. Arbib: Self-reproducing automata – Some implications for theoretical biology. In: *Towards a Theoretical Biology*, 2nd edn., ed. by C.H. Waddington (Edinburgh Univ. Press, Edinburgh 1969) pp. 204–226

76.80 J. Aloimonos, I. Weiss, A. Bandopadhay: Active vision, Int. J. Comput. Vis. **1**(4), 333–356 (1987)

76.81 R. Bajcsy: Active perception, Proc. IEEE **76**(8), 996–1005 (1988)

76.82 D.H. Ballard: Animate vision, Artif. Intell. **48**(1), 57–86 (1991)

76.83 P.J. Hancock, R.J. Baddeley, L.S. Smith: The principal components of natural images, Network **3**, 61–70 (1992)

76.84 D. Floreano, T. Kato, D. Marocco, E. Sauser: Coevolution of active vision and feature selection, Biol. Cybern. **90**(3), 218–228 (2004)

76.85 D. Floreano, M. Suzuki, C. Mattiussi: Active vision and receptive field development in evolutionary robots, Evol. Comput. **13**(4), 527–544 (2005)

76.86 T.D. Sanger: Optimal unsupervised learning in a single-layer feedforward neural network, Neural Netw. **2**, 459–473 (1989)

76.87 I. Harvey, E.A. Di Paolo, R. Wood, M. Quinn, E. Tuci: Evolutionary robotics: A new scientific tool for studying cognition, Artif. Life **11**(1-2), 79–98 (2005)

76.88 A. Seth: Causal connectivity of evolved neural networks during Behaviour, Netw. Comput. Neural Syst. **16**(1), 35–54 (2005)

76.89 E. Izquierdo, S. Lockery: Evolution and analysis of minimal neural circuits for klinotaxis in Caenorhabditis elegans, J. Neurosci. **30**, 12908–12917 (2010)

76.90 P. Husbands, R.C. Moioli, Y. Shim, A. Philippides, P.A. Vargas, M. O'Shea: Evolutionary robotics and neuroscience. In: *The Horizons of Evolutionary Robotics*, ed. by P.A. Vargas, E.A. Di Paolo, I. Harvey, P. Husbands (MIT, Cambridge 2013) pp. 17–64

76.91 D.T. Cliff: Computational neuroethology: A provisional manifesto, Proc. 1st Int. Conf. Simul. Adapt. Behav.: Anim. Animat., ed. by J.-A. Meyer, S.W. Wilson (MIT, Cambridge 1991) pp. 29–39

76.92 R. Held, A. Hein: Movement-produced stimulation in the development of visually guided behavior, J. Comp. Physiol. Psychol. **56**(5), 872–876 (1963)

76.93 R. Held: Plasticity in sensory-motor systems, Sci. Am. **213**(5), 84–94 (1965)

76.94 M. Suzuki, D. Floreano, E.A. Di Paolo: The contribution of active body movement to visual development in evolutionary robots, Neural Netw. **18**(5/6), 656–665 (2005)

76.95 S. Healy (Ed.): *Spatial Representations in Animals* (Oxford Univ. Press, Oxford 1998)

76.96 N.A. Schmajuk, H.T. Blair: Place learning and the dynamics of spatial navigation: A neural network approach, Adapt. Behav. **1**, 353–385 (1993)

76.97 N. Burgess, J.G. Donnett, K.J. Jeffery, J. O'Keefe: Robotic and neuronal simulation of the hippocampus and rat navigation, Philos. Trans. R. Soc. **352**, 1535–1543 (1997)

76.98 J. O'Keefe, L. Nadel: *The Hippocampus as a Cognitive Map* (Clarendon, Oxford 1978)

76.99 J.S. Taube, R.U. Muller, J.B. Ranck Jr.: Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis, J. Neurosci. **10**, 420–435 (1990)

76.100 D.E. Rumelhart, J. McClelland, P.D.P. Group: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (MIT, Cambridge 1986)

76.101 W. Maas, C.M. Bishop (Eds.): *Pulsed Neural Networks* (MIT, Cambridge 1999)

76.102 F. Rieke, D. Warland, R. van Steveninck, W. Bialek: *Spikes:: Exploring the Neural Code* (MIT, Cambridge 1997)

76.103 G. Indiveri, P. Verschure: Autonomous vehicle guidance using analog VLSI neuromorphic sensors, Lect. Notes Comput. Sci. **1327**, 811–816 (1997)

76.104 M.A. Lewis, R. Etienne-Cummings, A.H. Cohen, M. Hartmann: Toward biomorphic control using custom aVLSI CPG chips, Proc. IEEE Int. Conf. Robotics Autom. (ICRA) (2000) pp. 494–500

76.105 D. Floreano, C. Mattiussi: Evolution of spiking neural controllers for autonomous vision-based robots. In: *Evolutionary Robotics. From Intelligent Robotics to Artificial Life*, ed. by T. Gomi (Springer, Tokyo 2001) pp. 38–61

76.106 W. Gerstner, J.L. van Hemmen, J.D. Cowan: What matters in neuronal locking?, Neural Comput. **8**, 1653–1676 (1996)

76.107 D. Floreano, Y. Epars, J.C. Zufferey, C. Mattiussi: Evolution of spiking neural circuits in autonomous mobile robots, Int. J. Intell. Syst. **21**(9), 1005–1024 (2006)

76.108 J.A. Gally, P.R. Montague, G.N. Reeke, G.M. Edelman: The NO hypothesis: Possible effects of a short-lived, rapidly diffusible signal in the development and function of the nervous system, Proc. Natl. Acad. Sci. **87**(9), 3547–3551 (1990)

76.109 J. Wood, J. Garthwaite: Models of the diffusional spread of nitric oxide: Implications for neural nitric oxide signaling and its pharmacological properties, Neuropharmacology **33**, 1235–1244 (1994)

76.110 T.M. Dawson, S.N. Snyder: Gases as biological messengers: Nitric oxide and carbon monoxide in the brain, J. Neurosci. **14**(9), 5147–5159 (1994)

76.111 J. Garthwaite, C.L. Boulton: Nitric oxide signaling in the central nervous system, Annu. Rev. Physiol. **57**, 683–706 (1995)

76.112 A.O. Philippides, P. Husbands, M. O'Shea: Four-dimensional neuronal signaling by nitric oxide: A computational analysis, J. Neurosci. **20**(3), 1199–1207 (2000)

76.113 C. Hölscher: Nitric oxide, the enigmatic neuronal messenger: Its role in synaptic plasticity, Trends Neurosci. **20**, 298–303 (1997)

76.114 P. Husbands, T. Smith, N. Jakobi, M. O'Shea: Better living through chemistry: Evolving GasNets for robot control, Connect. Sci. **10**(4), 185–210 (1998)

76.115 T.M.C. Smith, P. Husbands, M. O'Shea: Local evolvability, neutrality, and search difficulty in evolutionary robotics, Biosystems **69**, 223–243 (2003)

76.116 A.O. Philippides, P. Husbands, T. Smith, M. O'Shea: Flexible couplings: Diffusing neuromodulators and adaptive robotics, Artif. Life **11**(1-2), 139–160 (2005)

76.117 A.O. Philippides, P. Husbands, T. Smith, M. O'Shea: Structure based models of NO diffusion in the nervous system. In: *Computational Neuroscience: A Comprehensive Approach*, ed. by J. Feng (CRC, Boca Raton 2004) pp. 97–130

76.118 A.O. Philippides, S.R. Ott, P. Husbands, T. Lovick, M. O'Shea: Modeling co-operative volume signaling in a plexus of nitric oxide synthase-expressing neurons, J. Neurosci. **25**(28), 6520–6532 (2005)

76.119 P. Husbands, A. Philippides, P. Vargas, C. Buckley, P. Fine, E.A. Di Paolo, M. O'Shea: Spatial, temporal and modulatory factors affecting GasNet evolvability in a visually guided robotics task, Complexity **16**(2), 35–44 (2010)

76.120 D. Barañano, C. Ferris, S. Snyder: A typical neural messenger, Trends Neurosci. **24**(2), 99–106 (2001)

76.121 T.M.C. Smith, P. Husbands, A. Philippides, M. O'Shea: Neuronal plasticity and temporal adaptivity: Gasnet robot control networks, Adapt. Behav. **10**(3/4), 161–184 (2002)

76.122 G. Edelman, J. Gally: Degeneracy and complexity in biological systems, Proc Natl. Acad. Sci. USA **98**, 13763–13768 (2001)

76.123 C. Fernando, K. Karishma, E. Szathmáry: Copying and evolution of neuronal topology, PLoS ONE **3**(11), e3775 (2008)

76.124 C. Fernando, E. Szathmáry, P. Husbands: Selectionist and evolutionary approaches to brain function: A critical appraisal, Front. Comput. Neurosci. **6**, 24 (2012)

76.125 S. Nolfi, D. Floreano: Learning and evolution, Auton. Robots **7**, 89–113 (1999)

76.126 S. Nolfi, D. Parisi: Learning to adapt to changing environments in evolving neural networks, Adapt. Behav. **1**, 75–98 (1997)

76.127 J.M. Baldwin: A new factor in evolution, Am. Nat. **30**, 441–451 (1896)

76.128 C.H. Waddington: Canalization of development and the inheritance of acquired characters, Nature **150**, 563–565 (1942)
76.129 G. Mayley: Landscapes, learning costs, and genetic assimilation, Evol. Comput. **4**, 213–234 (1997)
76.130 D. Floreano, F. Mondada: Evolution of plastic neurocontrollers for situated agents, Proc. 4th Int. Conf. Simul. Adapt. Behav.: Anim. Animat. 4, ed. by P. Maes, M. Matarić, J.A. Meyer, J. Pollack, H. Roitblat, S. Wilson (MIT, Cambridge 1996) pp. 402–410
76.131 D. Floreano, J. Urzelai: Evolutionary robots with online self-organization and behavioral fitness, Neural Netw. **13**, 431–443 (2000)
76.132 D. Floreano, J. Urzelai: Neural morphogenesis, synaptic plasticity, and evolution, Theory Biosci. **120**(3-4), 225–240 (2001)
76.133 E. Di Paolo: Evolving spike-timing-dependent plasticity for single-trial learning in robots, Philos. Trans. R. Soc. Lond. **361**, 2299–2319 (2003)
76.134 Y.U. Cao, A.S. Fukunaga, A. Kahng: Cooperative mobile robotics: Antecedents and directions, Auton. Robots **4**, 7–27 (1997)
76.135 S. Nolfi: Co-evolving predator and prey robots, Adapt. Behav. **20**, 10–15 (2012)
76.136 D. Floreano, S. Nolfi: God save the red queen! Competition in co-evolutionary robotics, Proc. 2nd Conf. Genetic Program., ed. by J.R. Koza, K. Deb, M. Dorigo, D. Foegel, B. Garzon, H. Iba, R.L. Riolo (Morgan Kaufmann, San Francisco, CA 1997) pp. 398–406
76.137 S. Nolfi, D. Floreano: Co-evolving predator and prey robots: Do *arm races* arise in artificial evolution?, Artif. Life **4**(4), 311–335 (1998)
76.138 D. Floreano, S. Nolfi: Evolution versus design: Controlling autonomous robots, Proc. 4th Eur. Conf. Artif. Life, ed. by P. Husbands, I. Harvey (MIT, Cambridge 1997) pp. 378–387
76.139 V. Trianni, S. Nolfi: Evolving collective control, cooperation and distributed cognition. In: *Handbook of Collective Robotics – Fundamentals and Challenges*, ed. by S. Kernbach (CRC, Boca Raton 2012) pp. 246–276
76.140 G. Baldassarre, V. Trianni, M. Bonani, F. Mondada, M. Dorigo, S. Nolfi: Self-organised coordinated motion in groups of physically connected robots, IEEE Trans. Syst. Man Cybern. **37**, 224–239 (2007)
76.141 M. Quinn, L. Smith, G. Mayley, P. Husbands: Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors, Philos. Trans. R. Soc. Lond. **361**, 2321–2344 (2003)
76.142 G. Baldassarre, D. Parisi, S. Nolfi: Coordination and behavior integration in cooperating simulated robots, Proc. 8th Int. Conf. Simul. Adapt. Behav.: Anim. Animat. 8 (MIT, Cambridge 2003) pp. 385–394
76.143 V. Sperati, V. Trianni, S. Nolfi: Self-organised path formation in a swarm of robots, Swarm Intell. **5**, 97–119 (2011)
76.144 M. Quinn: Evolving communication without dedicated communication channels, Proc. 6th Eur. Conf. Artif. Life, ed. by J. Kelemen, P. Sosik (Springer, Berlin, Heidelberg 2001) pp. 357–366
76.145 D. Marocco, S. Nolfi: Self-organization of communication in evolving robots, Proc. 10th Int. Conf. Artif. Life, ed. by L. Rocha, L. Yeager, M. Bedau, D. Floreano, R. Goldstone, A. Vespignani (MIT, Cambridge 2006) pp. 178–184
76.146 D. Floreano, S. Mitri, S. Magnenat, L. Keller: Evolutionary conditions for the emergence of communication in robots, Curr. Biol. **17**, 514–519 (2007)
76.147 M. Waibel, D. Floreano, S. Magnenat, L. Keller: Division of labour and colony efficiency in social insects: Effects of interactions between genetic architecture, colony kin structure and rate of perturbations, Proc. Royal Soc. B Biol. Sci. **273**, 1815–1823 (2006)
76.148 F. Mondada, G. Pettinaro, A. Guignard, I. Kwee, D. Floreano, J.L. Deneubourg, S. Nolfi, L.M. Gambardella, M. Dorigo: Swarm-bot: A new distributed robotic concept, Auton. Robots **17**, 193–221 (2004)
76.149 V. Trianni, S. Nolfi, M. Dorigo: Cooperative hole-avoidance in a swarm-bot, Robotics Auton. Syst. **54**, 97–103 (2006)
76.150 G. Baldassarre, S. Nolfi, D. Parisi: Evolving mobile robots able to display collective behavior, Artif. Life **9**, 255–267 (2003)
76.151 S. Nolfi: Evolution of communication and language in evolving robots. In: *Current Perspective on the origin of language*, ed. by C. Lefebvre, B. Comrie, H. Cohen (Cambridge Univ. Press, Cambridge 2013)
76.152 J. De Greef, S. Nolfi: Evolution of implicit and explicit communication in a group of mobile robots. In: *Evolution of Communication and Language in Embodied Agents*, ed. by S. Nolfi, M. Mirolli (Springer, Berlin, Heidelberg 2010) pp. 179–214
76.153 S. Mitri, D. Floreano, L. Keller: The evolution of information suppression in communicating robots with conflicting interests, Proc. Natl. Acad. Sci. **106**, 15786–15790 (2009)
76.154 S. Mitri, D. Floreano, L. Keller: Relatedness influences signal reliability in evolving robots, Proc. Royal Soc. B Biol. Sci. **278**, 378–383 (2011)
76.155 S. Nolfi: Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours, Connect. Sci. **3-4**, 231–248 (2005)
76.156 S. Wischmanna, D. Floreano, L. Keller: Historical contingency affects signaling strategies and competitive abilities in evolving populations of simulated robots, Proc. Natl. Acad. Sci. **109**, 864–868 (2011)
76.157 A. Thompson: Evolving electronic robot controllers that exploit hardware resources, Lect. Notes Artif. Intell. **929**, 640–656 (1995)
76.158 A. Thompson: *Hardware Evolution: Automatic Design of Electronic Circuits in Reconfigurable Hardware by Artificial Evolution*, Distinguished Dissertation Series (Springer, Berlin, Heidelberg 1998)

76.159 A. Thompson: Artificial evolution in the physical world. In: *Evolutionary Robotics. From Intelligent Robots to Artificial Life (ER'97)*, ed. by T. Gomi (AAI Books, Ottawa 1997) pp. 101–125

76.160 D. Keymeulen, M. Durantez, M. Konaka, Y. Kuniyoshi, T. Higuchi: An evolutionary robot navigation system using a gate-level evolvable hardware, Lect. Notes Comput. Sci. **1259**, 193–209 (1996)

76.161 G. Ritter, J.-M. Puiatti, E. Sanchez: Leonardo and discipulus simplex: An autonomous, evolvable six-legged walking robot, Lect. Notes Comput. Sci. **1586**, 688–696 (1999)

76.162 D. Roggen, D. Floreano, C. Mattiussi: A morphogenetic evolutionary system: Phylogenesis of the POETIC circuit, Lect. Notes Comput. Sci. **2606**, 153–164 (2003)

76.163 D. Roggen, S. Hofmann, Y. Thoma, D. Floreano: Hardware spiking neural network with run-time reconfigurable connectivity in an autonomous robot, NASA/DoD Conf. Evolv. Hardw., ed. by J. Lohn, R. Zebulum, J. Steincamp, D. Keymeulen, A. Stoica, M.I. Fergusonpages (2003) pp. 189–198

76.164 M. Matarić, D. Cliff: Challenges in evolving controllers for physical robots, Robotics Auton. Syst. **19**(1), 67–83 (1996)

76.165 G. Massera, T. Ferrauto, O. Gigliotta, S. Nolfi: FARSA: An open software tool for embodied cognitive science, Proc. 12th Eur. Conf. Artif. Life, ed. by P. Lio, O. Miglino, G. Nicosia, S. Nolfi, M. Pavone (MIT, Cambridge 2013) pp. 454–538

76.166 Framework for Autonomous Robotics Simulation and Analysis: http://laral.istc.cnr.it/farsa