

4 Evolutionary Robotics

Stefano Nolfi

4.1 Introduction

Evolutionary robotics (Nolfi and Floreano 2000; Nolfi et al. 2016; Nolfi, 2021) is a method that allows the creation of robots capable of developing the ability to perform one or more functions as a result of an adaptation process analogous to natural evolution.

Robots are considered to be autonomous artificial organisms that adapt in close interaction with the environment without human intervention. The role of the experimenter is limited to the specification of the fitness function—that is, the criteria used to evaluate the performance level of the robots—and to the specification of the characteristics of the robots that are not subjected to the adaptive process. The remaining characteristics are encoded in a vector of parameters (genotype) and evolved through an evolutionary algorithm (Rechenberg 1973; Goldberg and Holland 1988). In the majority of cases, the evolving robots are provided with neural network controllers. The connection weights of the network, which determine the behavior of the robot, are encoded in the genotype and evolved. Eventually, the architecture of the neural network (Stanley and Miikkulainen 2002; Durr, Mattiussi, and Floreano 2006) and/or the morphology of the robot can be encoded in the genotype and evolved (Sims 1994; Lipson and Pollack 2000; Auerbach and Bongard 2012; Hiller and Lipson 2012).

The evolutionary process is realized by creating an initial population of genotypes generated randomly and then repeating the following steps for a certain number of generations: 1) create a population of robots with the characteristics specified in the corresponding genotypes, 2) allow the robots to interact with their environment for a finite amount of time and calculate a scalar value (fitness) that rates the performance of each robot with respect to a given problem, and 3) create a new population of genotypes composed of copies with random variations of the genotypes of the fittest robots.

An important aspect to consider is that the utilization of a fitness function that rewards the robot for performing a given function—for example, foraging—can drive the development of several behavioral and cognitive capacities that are instrumental to the achievement of that function, such as avoiding obstacles and dangers, orienting and navigating in the environment, discriminating relevant objects, integrating sensory information over time and later using it to appropriately regulate the robot's behavior, and so on. The analysis of the

way in which these capacities are realized and integrated in evolving robots can provide valuable information from the perspective of modeling the organization and the development of similar capacities in natural systems.

Evolutionary robotics has been applied to the study of a wide range of phenomena, including embodied cognition, sensorimotor coordination, integration of behavioral and cognitive skills, social and collective behaviors, internal models, and interaction between evolution and learning. In the following sections, I will describe a few representative examples of the work conducted in these areas.

4.2 Evolving Bodies and Brains: Morphological Computation

The behavioral and cognitive skills of robots or animals are dynamical properties that unfold in time and arise from a large number of interactions between the agent's nervous system, body, and environment (Chiel and Beer 1997; see also chapter 11). The dynamical process originating from the interactions depends on the characteristics of the agent's body and brain. This implies that varying the characteristics of the body and/or of the brain can shape the dynamical process.

An example of behavior that can be realized by shaping the characteristics of the body or of the brain is walking on a declining plane. Indeed, it can be produced either by brainless robots with passive joints and carefully designed body morphologies (McGeer 1990; Collins et al. 2005) or by highly controlled robots lacking the morphological features of the former robots (Chestnutt et al. 2005). The term “morphological computation” (Pfeifer et al. 2006; Paul 2006; see also chapter 1) has been introduced to indicate processes performed by the body that otherwise would have to be performed by the brain. Solutions exploiting morphological computation are often advantageous in terms of energy efficiency and robustness with respect to alternative solutions (Pfeifer and Bongard 2006).

The possibility of adapting both the body plan and the control policy of robots permits the selection of solutions that are simpler and more effective within the spectrum of those available—that is, among solutions relying primarily on morphological computation or on control. Moreover, it permits the generation of solutions in which the morphological and control features are coadapted. Evolutionary robotics constitutes an ideal approach for adapting both the policy and the morphology of robots since it is a model-free method that does not make any assumption about the structure of the adaptive system. Moreover, unlike alternative model-free training methods, it permits the adaptation of any type of parameter, including a combination of qualitatively different parameters. The number of body parts forming the body of the robot, the relative position of these parts, the physical properties of each body part, and the characteristics of the joints among body parts can be encoded in the genotype and evolved together with the characteristics of the neural network of the robot. This is typically realized by using genotypes that encode growing rules, which determine how the initial “embryo” grows and differentiates, rather than using genotypes that directly encode the property of a fully formed robot.

In a pioneering work in this area, Sims (1994) demonstrated how artificial evolution can be used to evolve the morphology and the control policy of simulated creatures capable of swimming, walking, and grabbing objects while competing with other creatures. Lipson and

Pollack (2000) later used a similar approach to evolve simulated walking robots that are then manufactured using a three-dimensional printer and spare electronic components.

Since that time, this approach has been used for various purposes. For example, Long (2012) evolved the stiffness of artificial tails of swimming robots to investigate how backbones evolved in early vertebrates. By evolving robots in environments of varying complexity, Auerbach and Bongard (2012) showed how the complexity of the evolved morphology correlates with the complexity of the environment. For example, robots evolved to walk on irregular terrain develop morphologies that include appendages missing in robots evolved over flat terrain. Hiller and Lipson (2012) demonstrated how evolving robots made of cells with different material properties arranged in evolved topologies can produce a variety of locomotion behaviors. These behaviors originate from simple periodic expansion/contraction actions produced by some of the cells and from the physical interactions among the cells composing the robot body and among the cells and the environment. These simulated robots composed of multiple cells can then be transformed into artificial living creatures by assembling ectoderm and cardiac stem cells in the same three-dimensional spatial configuration (Kriegman et al. 2020). Remarkably, these artificial living creatures are able to locomote and to explore their aqueous environment autonomously for days.

4.3 Sensorimotor Coordination

In agents that are embodied and situated, the role of perception cannot be separated by that of action and vice versa. What an agent perceives is determined by what it does, and what an agent does can be determined by what the agent needs to perceive.

The existence of a close link between perception and action draws on a number of distinct traditions in philosophy, in psychology, and in the cognitive sciences. It is at the core of the ecological theory of perception developed by Gibson (1979) and of several other fundamental contributions (Arbib 1989; Varela, Thomson, and Rosh 1991; Maturana and Varela 1987; Thelen and Smith 1994; Berthoz 2000; O'Regan and Noë 2001; Noë 2004; Clark 1998, 1999). The coupling of the sensory and the motor process can be indicated with the term “sensorimotor coordination” (Dewey 1981 [1986]).

Evolutionary robotics constitutes an ideal framework for studying the role of sensorimotor coordination in the development of behavioral and cognitive skills. The first reason is that the evolutionary process leaves the evolving robots free to determine the way in which they achieve their adaptive goals. Consequently, the robots are free to coordinate their perceptual and action processes in ways that are functional to the achievement of their objectives. The second reason is that the evolutionary process is driven by a fitness measure that rates the overall performance of the robot—that is, the sum of rewards obtained over an extended evaluation period. This permits variations that enhance the coordination between the sensory and action process to be identified and retained regardless of whether the time interval between actions and associated rewards is immediate or delayed.

Indeed, sensorimotor coordination plays a crucial role in practically all experiments carried out by evolving robots. The first demonstration was reported in an experiment in which a wheeled robot provided with infrared sensors and situated in an arena surrounded by walls was evolved for the ability to find and remain near a cylindrical object (Nolfi 1996, 2005). Interestingly, the evolved robots did not solve the problem by internally processing

the experienced sensory states in order to discriminate the stimuli corresponding to walls and cylinders, a strategy that was actually challenging since the stimuli experienced near cylinders and walls strongly overlap in sensory space. They instead solved the task by reacting to the stimuli to produce behavioral attractors—that is, oscillatory behavior generated by alternating move-forward/move-backward and turn-left/turn-right actions, near cylinders but not near walls. In other words, they exploited the fact that the execution of the same actions has different perceptual consequences near walls or cylinders that can lead to the production of the two required differentiated behaviors. This experiment can be replicated with the Evorobotpy software tool available from <https://github.com/snolfi/evorobotpy> (see the instruction for running the ErDiscrim experiment in Nolfi 2021, chapter 13).

In an extended version of this experiment, in which the robot was provided with proprio-sensors that encoded the speed of the robot's wheels, Scheier, Pfeifer, and Kunyioshi (1998) observed the evolution of a qualitatively different sensorimotor strategy that exploits actions to self-select easy-to-interpret stimuli. In this case the evolved robots displayed a wall-following behavior near walls and cylinders of moving straight along the wall and turning around the cylinder, respectively. They then used the perceived offset between the speed of the left and right wheel to keep producing the wall-following behavior near cylinders and to move away from walls. In other words, the robots acted to later experience favorable sensory states. They displayed an initial behavior that enabled them to later experience two well-differentiated states on their propriosensors near walls and cylinders.

Qualitatively similar solutions have been observed in more complex robots evolved for the ability to solve more challenging problems. This is the case, for example, of an experiment in which a simulated iCub robot (Sandini, Metta, and Vernon 2004) was evolved for the ability to discriminate spherical and ellipsoid objects on the basis of rough tactile information (Tuci, Massera, and Nolfi 2010). The robot was provided with fourteen motor neurons that encoded the torque produced by seven sets of antagonistic muscles controlling the seven degrees of freedom (DOFs) of the arm and of the wrist, two motor neurons that encoded the desired extension/flexion of the thumb and of the four fingers, and two motor neurons that indicated the category of the object (i.e., spherical or ellipsoid). The sensors of the robot included eight neurons that encoded the current angular position of the DOFs of the arm and of the wrist, five neurons that encoded the extension/flexion of the five corresponding fingers, and ten neurons that encoded the ten touch sensors located on the fingertips and on the palm. Touch sensors binarily encoded whether the corresponding part of the robot body collided with another body. The robots were rewarded for discriminating the shape of the objects experienced during multiple evaluation episodes. They were not rewarded for the production of any specific behaviors and consequently were left free to select behaviors that enabled and/or facilitated the discrimination problem.

The analysis of the evolved robots demonstrates that they did indeed develop manipulation behaviors that enabled them to experience stimuli allowing them to reliably discriminate the two types of objects despite the similarity of the objects' shapes and the limited resolution of the touch sensors. The categorization process involves three phases. In the first part, the robot manipulates the object by wrapping it with its fingers and by moving the object until a suitable hand/object posture is reached. The information contained in the tactile stimuli experienced during this phase increases and finally reaches a high value when a hand/object achieves a suitable posture, which remains almost stable in the remain-

ing part of the episode. During the second phase, the robot starts to produce a categorization answer, keeps producing fine manipulation actions, and keeps integrating the sensory information experienced by eventually reversing its categorization decision. This continues during the third phase, in which the categorization decision is no longer reversible.

The solutions discovered by the evolved robots thus fit the dynamical view of cognition elaborated by Spivey (2007). The extension of the categorization process over time enables the robot to experience useful stimuli and to integrate the conflicting evidence experienced over time in order to maximize the accuracy of the categorization decision.

4.4 On the Relation between Reactive and Cognitive Capabilities

Evolutionary robotics can also be used to study the relation and the integration between behavioral and cognitive capabilities.

As discussed above, morphological computation and sensorimotor coordination can be used to perform processes that the brain would otherwise have to perform. The exploitation of the interaction between the agent and the environment thus permits reliance on solutions that are simpler, from an internal-processing perspective, than solutions that do not rely on these properties. This opens up questions about the relationship between reactive and cognitive capabilities. Do they tend to interact in a synergetic or conflictual manner? And “is cognition truly seamless—implying a gentle, incremental trajectory linking fully embodied responsiveness to abstract thought and off-line reason? Or is it a patchwork quilt, with jumps and discontinuities and with very different kinds of processing and representations serving different needs?” (Clark 1999, 350).

Interesting evidence supporting a synergetic relation and a smooth incremental integration of reactive and cognitive capabilities has been reported in evolutionary experiments addressing the evolution of a robot selected for the ability to navigate in a double T-maze environment (figure 4.1; Carvalho and Nolfi 2016). The robot, which is initially located in an area at the bottom of the central corridor with a randomly varying position and orientation, should

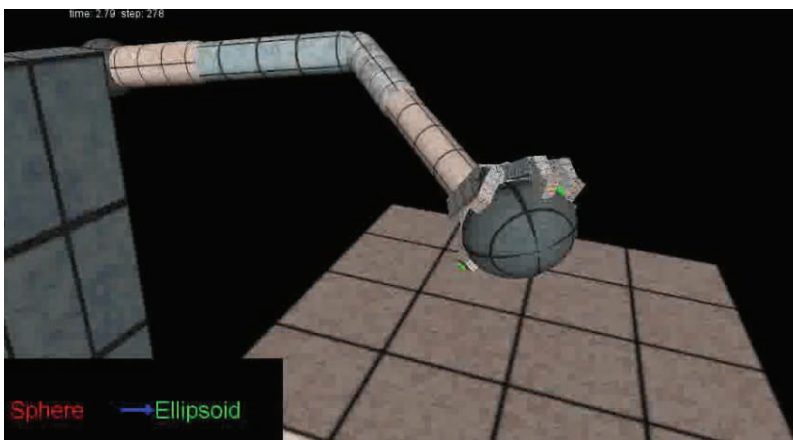


Figure 4.1
The object-discriminating robot.

travel toward a target destination located at one of the four ends of the maze. The correct destination is marked by two green objects located in the central corridor. The robot should thus solve a time-delay problem in which the information experienced while it travels down the central corridor should later influence the direction in which the robot turns when it reaches the first and the second junction.

The analysis of evolving robots indicates that they solve the problem with a strategy that does not require them to store the information extracted from the green object in internal states, recognize the arrival at the first and at the second junction, or turn left or right on the basis of the internal states and of the junction. As shown in figure 4.2, the trajectories produced during different evaluation episodes first converge in the bottom portion of the central corridor and then diverge while the robot perceives the position of the green objects. The initial convergence enables the robot to reduce the differences caused by the varying initial positions and orientations. The divergent process allows the robot to enter into one of four separate basins of attraction of robot/environmental dynamics that bring the robot to the right destination—the destination that matches the relative position of the two green beacons.

The strategy displayed by evolved robots thus exploits a form of cognitive off-loading—that is, the possibility of off-loading an agent's future intention into the external environment (Gilbert 2015a, 2015b). More specifically, the robot off-loads the information experienced in the central corridor by assuming different positions and orientations with respect to the corridor and by then maintaining such positions/orientations. The relative position of the robot in the corridor is then used to turn appropriately left or right at the first and then at the second junction. The trajectories displayed in figure 4.2 are produced by a robot that has no memory. However, similar strategies are produced by robots with memory—that is, by robots provided with recurrent connections in their internal neurons. The possibility of off-loading information

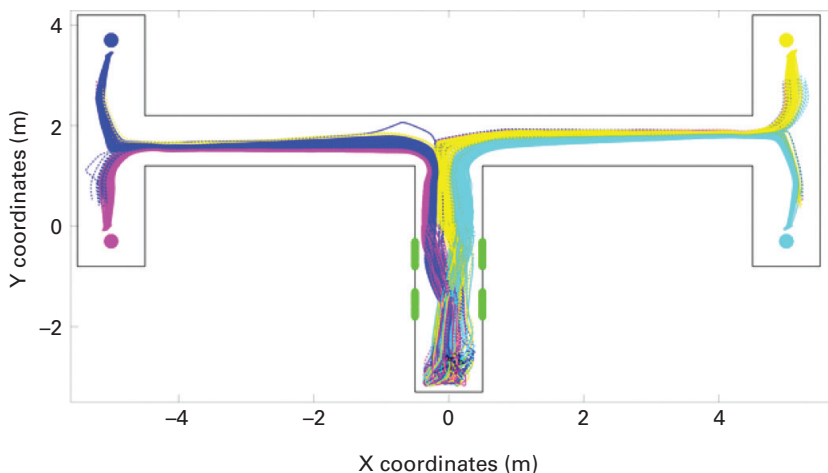


Figure 4.2

Trajectories of a typical evolved robot postevaluated for three hundred episodes. The trajectories (shown in magenta, blue, yellow, and cyan) indicate those produced by the robot during episodes in which it should have navigated toward the destination with the corresponding color. The target destination is marked by the relative position of the two green objects located to the left or right of the central corridor.

in the environment is thus preferred to alternative solutions relying on internal processing independently from the availability of memory.

Interestingly, evolving robots subjected to position perturbations, such as being randomly moved left or right as a result of “gusts of wind” occurring from time to time, solve the problem by developing composite strategies that rely on cognitive off-loading to determine the motor trajectory and on memory to reenter the appropriate basin of attraction after a position perturbation. This and additional control experiments reported in Carvalho and Nolfi (2016) demonstrate how, at least in this domain, reactive strategies do not prevent but rather promote the development of cognitive capabilities. Moreover, they illustrate how the development of cognitive capacities does not lead to the elimination of preexisting reactive capacities but rather to their extension.

4.5 Social and Collective Behavior

In the previous section, we limited our analysis to individual behaviors—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. This can be done simply by situating the evolving robots in environments containing other robots.

This scenario has been used to study the conditions that support the evolution of cooperative behavior. As expected, cooperative behavior readily emerges when a group of interacting robots is formed by genetically related individuals (e.g., individuals possessing identical genotypes) or when selection operates at the level of the colony or swarm (Floreano et al. 2007). When instead the individuals forming the colony are not genetically related and selection operates at the level of individuals, the evolutionary process leads to a dynamic in which cooperation periodically emerges and extinguishes (Mitri, Floreano, and Keller 2009).

The evolution of genetically related robots readily produces self-organizing properties—that is, the spontaneous formation of spatial, temporal, or spatiotemporal structures or functions that emerge from local interactions among individual robots and that are robust with respect to environmental variations (Camazine et al. 2001; see also chapter 5). For example, Sperati, Trianni, and Nolfi (2011) conducted experiments in which a population of wheeled robots was evolved for the ability to forage. The evolving robots developed an ability to arrange themselves in dynamic chains that enabled the colony to efficiently navigate between a nest and a foraging area. These dynamic chains, which self-sustain in the presence of perturbations, allow robots with limited individual sensory capacities to efficiently navigate to the right destination by discovering and storing information on the location of the relevant environmental areas at the level of the colony. Another example of self-organized behavior has been observed in a population of robots capable of self-assembly—in this case, by physically attaching together—to master problems that cannot be solved by individual robots. Robots evolved for the ability to move while attached developed an ability to negotiate a common direction of motion and to keep moving along that direction by compensating for misalignments originating during motion (Baldassarre et al. 2007). Also in this case, the ability to coordinate and to cooperate was robust with respect to variations in the environmental conditions. Indeed, evolved robots

were capable of coordinating independently from the configuration in which they were assembled. Moreover, robots evolved in specific environmental conditions demonstrated the ability to generalize their skills to new environmental conditions. Such generalization capacity included the ability to display new behaviors adapted to the new experienced conditions. For example, robot swarms evolved in an environment with no obstacles demonstrated an ability to avoid obstacles and to rearrange their shape to pass through narrow passages when situated in a mazelike environment with obstacles (Nolfi 2009).

The evolution of collective behavior in robots can also lead to the emergence of task specialization—that is, to individuals capable of assuming different complementary roles that increase the efficacy of the group (Ferrante et al. 2015; Pagliuca and Nolfi 2018).

The evolution of robots selected for the ability to solve a problem that benefits from cooperation has also been used to study the evolution of communication and language (Cangelosi and Parisi 2002; Nolfi and Mirolli 2010; see also chapter 20). In a series of experiments reported in De Greef and Nolfi (2010), the authors analyzed the origin and complexification of the communication system displayed by evolving robots across generations and the origin and transformation of the meaning associated with communication signals. These analyses indicate that the development of communication capabilities is strongly interlinked with the evolution of other capabilities. Robots need to develop appropriate behaviors to access and/or generate the information to be communicated and to react appropriately to detected signals. Interestingly, the development of communication skills scaffolds the development of behavioral skills and vice versa. This leads to the development of integrated capabilities and to a progressive complexification of robots' skills (Nolfi 2013).

Finally, evolutionary robotics experiments have been used to explain why reciprocity, the reciprocal exchange of episodes of help between two partners, is rare in nature (André and Nolfi 2016). This fact contrasts with the predictions generated by game theoretic models that reciprocity should evolve easily (Axelrod and Hamilton 1981). As shown by André and Nolfi (2016), these game theoretic models' predictions are in error because these methods do not model the mechanisms underlying the generation of behavior, a limitation that does not affect evolutionary robotics models. Indeed, the experiments carried out by evolving robots predict correctly that reciprocity is unlikely to evolve, due to the numerous neutral mutations required to generate a reciprocator behavior from individuals that do not reciprocate.

Another line of research has investigated the evolution of social behaviors in competing scenarios—for example, the evolution of a population of robots with conflicting interests. The coevolution of competing species such as predator and prey might favor the synthesis of evolutionary innovations. Indeed, “an adaptation in one lineage (e.g., predators) may change the selection pressure on another lineage (e.g., prey), giving rise to a counteradaptation. If this occurs reciprocally, an unstable runaway escalation of ‘arm races’ may result” (Dawkins and Krebs 1979, 489; Rosin and Belew 1997). In other words, adaptations on one side call for counteradaptations on the other side, and the counteradaptations call for more counteradaptations, and so on, thus producing an escalation process. Moreover, the concurrent evolution of the agents and of the learning environment can lead to a spontaneous, progressive complexification of the adaptive problem. That is to say, a pedagogically sound training process can be produced in which progress in one population

is accompanied by a gradual complexification of the adaptive task caused by parallel progress in the competing population (Rosin and Belew 1997).

Evolutionary experiments performed by evolving predator and prey robots (Cliff and Miller 1995; Nolfi and Floreano 1998) showed that co-evolution does indeed lead to “arms races” that produce a progressive complexification during the initial generations. The evolutionary dynamics, however, later converge in a limit-cycle dynamic in which progress against current competitors (local progress) is accompanied by retrogression with respect to ancient or future competitors. Cycling dynamics of this type were found in natural evolution in a population of side-blotched lizards (*Uta stansburiana*) by Sinervo and Lively (1996) and in *Daphnia* and associated parasites conserved in lake sediment (Decaestecker et al. 2007). More recently, Simione and Nolfi (2017, 2019) showed how long-term global progress can be produced in controlled ecological conditions—that is, in experiments in which the evolving populations are divided into subgroups that normally interact with specific subgroups of the competing population and only occasionally with the remaining competitors.

4.6 Evolution, Development, and Learning

The basic evolutionary method illustrated in the introduction can be extended to incorporate development and learning. In the basic method, the process that maps a genotype into a robot is completed before the robot starts to interact with its environment. In other words, robots are born as fully formed individuals. In extended evolutionary methods, by contrast, the developmental process continues during the period in which the robot interacts with its environment.

A model described in Bongard (2011), in which the evolving robots developed from an anguilliform morphology to a legged morphology while they interacted with the external environment, provides an example. The comparison with control experiments, in which the robots did not transition through the anguilliform body plan, indicates that morphological change accelerates the evolution of robust walking behaviors. A second example is given by a series of experiments reported in Kriegman, Cheney, and Bongard (2018) in which soft robots with developmental morphology were evolved for the ability to move over a surface. The analysis of the interaction between the evolutionary and developmental processes in these experiments enabled the authors to highlight an unknown aspect of genetic assimilation—namely, that the traits that render the agents robust to changes in other traits have a greater probability of becoming genetically assimilated in successive generations than traits that are less robust to genetic variations.

A model in which the brains of the robots keep developing while the robots operate in their environment was studied in Nolfi, Miglino, and Parisi (1994). In this model, the evolving robots were provided with neuron axons that grew and branched by establishing connections with other neurons while the robots operated in the environment. As with real nervous systems, the growth process of axons is influenced both by the activity patterns of the single neurons and by genetic factors (Purves 1994; Quartz and Sejnowski 1997). This leads to the evolution of robots capable of developing brains adapted to the environment in which they are situated—for example, to robots that might or might not develop

a brain area dedicated to processing light and in which development of the area is triggered by the exposure to light (Nolfi, Miglino, and Parisi 1994).

Other works have investigated the combination of evolution and learning (Nolfi and Floreano 1999). In these models the topology of the neural network was fixed, but the connection weights varied while the robots interacted with the environment on the basis of an unsupervised (Floreano, Durr, and Mattiussi 2008), self-supervised (Nolfi and Parisi 1993), or reinforcement-learning algorithm (Schembri, Mirolli, and Baldassarre 2007). The combination of evolution and learning enables evolving robots to adapt to environmental variations that occur within generations. For example, it enables predator robots to modify their behavior on the fly while interacting with a prey robot to display the strategy that is effective against the current encountered prey (Floreano and Nolfi 1997).

4.7 Internal Models

Evolutionary robotics is a model-free approach, a method that permits the robots to develop behavioral and cognitive skills from scratch without the need to rely on a model of the external environment and/or the robot's own self. However, the abilities that the robots develop during their adaptation can include the ability to build and use a model of their own body, a model of the external environment, and/or a forward model that allows the consequences of the robots' actions to be predicted.

Bongard, Zykov, and Lipson (2006) give an example of a robot capable of acquiring a model of its own body. In this work, a physical robot was equipped with an onboard simulator that it used to continually evolve a model of itself. The model consisted of a three-dimensional description of the robot's own body that enabled it to predict the perceptual effects of the actions it could execute without actually performing them. The robot then used the model to cope with damages, such as the mechanical separation of a leg. This was realized by 1) using the offset between the actual and predicted consequences of actions to diagnose the damage, 2) updating the model of the robot's own body to reduce the offset between the predicted and actual consequences of the robot's action, 3) evolving a new control policy capable of operating effectively with the damaged body by using a mental simulation, and 4) using the new control policy to keep operating effectively despite the damage. The availability of the world model thus permits the evolution of a compensatory policy by using the imagined effect of variations of the current policy (mental simulation) as a proxy for the actual effect of variations.

Cully et al. (2015) showed how the ability to recover from damages or faults can be speeded up by learning a behavior-performance map that encodes the correlation between the value of the connection weights and the value of fitness. The map can then be used to introduce mutations that have a higher chance of producing improvements with respect to random mutations.

Gigliotta, Pezzulo, and Nolfi (2011) demonstrated how a robot subjected to sensory deprivation can evolve the ability to react appropriately to sensory stimuli and to self-generate states functionally equivalent to sensory stimuli during sensory deprivation phases in which stimuli are not available. The behavior consists of moving the robot's eye to foveate consecutive portions of the image located over a circular trajectory. In normal phases, the robot can

determine the movement of the eye on the basis of the current perceived color. During blind phases, the robot should use self-generated internal states as proxies for missing sensory states. The analysis of the evolved robots indicates that the problem is not solved by generating states that match the missing sensory states. Rather, it is realized by generating internal states that elicit the appropriate movements but are not necessarily similar to the states that would be experienced in normal conditions.

Finally, Ha and Schmidhuber (2018) demonstrated how agents that determine their actions on the basis of features extracted from the sensory states, by a neural network trained with a self-supervised learning algorithm, outperform agents that determine their actions directly on the basis of the features encoded in sensory states. The problem considered consists of learning to drive in a car-racing environment called CarRacing-v0 (Brockman et al. 2016). The learning agent receives an image containing a top-down view of the car and the environment as input. The features are extracted by 1) a variational autoencoder network (Kingma and Welling 2013; Rezende et al. 2014) trained with the ability to encode perceived images in compact representations that can be used to reconstruct the original image and 2) a long short-term memory (LSTM) network (Hochreiter and Schmidhuber 1997) trained to predict the compressed state of the next perceived image on the basis of the compressed state of the current image and of the action the agent is going to perform. These two networks are pretrained using the images collected by the agent during several evaluation episodes in which the agent moves by performing random actions. The neural network controller of the evolving agents, which receives as input the internal state extracted by the sensors from the two pretrained networks described above, is evolved by using a standard evolutionary method for the ability to drive the car. In a second experiment performed by using the *VizDoom* game problem (Kempka et al. 2016), the authors showed that the autoencoder and LSTM prediction network described above can be used to evolve the agents in virtual worlds imagined by the agents themselves. The solutions evolved in these imagined worlds can then be successfully used to control the agent of a real *VizDoom* game.

4.8 Evolution as a Form of Learning

The evolutionary method can also be used to model ontogenetic learning (Schlesinger 2004). This is because the evolutionary algorithm constitutes one of the simplest yet most effective ways to evolve an embodied neural network through a trial-and-error process based on distal rewards. An example is illustrated in experiments in which an iCub humanoid robot (Sandini, Metta, and Vernon 2004) trained through an evolutionary method develops reaching and grasping skills analogous to those displayed by human infants from two to eighteen months of age (Savastano and Nolfi 2013). During this period, infants display a first transition from sweeping and unsuccessful arm movements to primitive, imprecise reaching and grasping behaviors and then a second transition leading to integrated and effective reaching and grasping behaviors (Konczak et al. 1995; Konczak, Borutta, and Dichgans 1997; Konczak and Dichgans 1997; von Hofsten and Rönqvist 1993; Spencer and Thelen 2000).

As illustrated in figure 4.3 (*left*), the robot is set in an upright position in front of a suspended object. This setting is similar to that used by Hofsten (1982) to analyze the development of



Figure 4.3

The simulated setting (*left*) is derived from experiments carried out on infants (*center and right*) by von Hofsten (1982).

reaching and grasping behavior in infants (figure 4.3, *center and right*). The training of the robot is realized in three phases: 1) a prereaching phase in which the robot has simple prewired reflex behaviors, low visual acuity, and an immature nervous system; 2) a gross-reaching phase in which the robot has improved visual acuity and matured cortical areas; and 3) a fine-reaching phase in which the robot has access to perceptual information that encodes the relative position of the object with respect to the hand.

The analysis of the experiments shows that the lack of internal neural resources during the prereaching phase has an adaptive role (i.e., channels the developmental process toward better solutions during the gross-reaching phase) and a bias role (i.e., represents a necessary condition for the emergence of the exploratory motor-babbling behavior). This suggests that the later involvement of cortical areas (Martin 2005) can play an adaptive role in humans and might have evolved to accomplish this function. Moreover, analysis of the behavior displayed by the robots during the course of the training process shows that the following phenomena observed in infants originate spontaneously: 1) a reduced use (freeze) of the distal DOFs of the arm of the robot during the prereaching phase, 2) an exploratory (motor-babbling) behavior during the prereaching phase, and 3) a temporal regression of the reaching capabilities at the onset of the fine-reaching phase. The fact that these qualitative variations emerge spontaneously indicates that they do not necessarily reflect the presence of additional specific maturational constraints. They can be the manifestation of a general self-structuring process that operates by temporarily reducing the complexity of the motor space, of the sensory space, and of the relevant task space, respectively.

In contrast to reinforcement-learning algorithms (Sutton and Barto 2018) that represent the most common choice to model trial-and-error learning, evolutionary algorithms present advantages and drawbacks. The advantages include the possibility of adapting all the characteristics of the robot, including the robot's morphology and the architecture of the robot's neural network and the ability to operate well in the presence of sparse reward. Reinforcement-learning algorithms, on the other hand, are generally more sample efficient.

The development of new evolutionary algorithms that operate by estimating the local gradient (Hansen and Ostermeier 2001) and eventually rely on stochastic gradient optimizers to vary the adaptive parameters (Salimans et al. 2017) makes the usage of evolutionary methods even more attractive. Indeed, although these gradient-ascent methods can also operate on populations that include multiple parents, they are typically used with popula-

tions composed of a single parent producing several offspring. The evaluation of the offspring is used to estimate the local gradient, which in turn is used to vary the parameters of the parent. This implies that, as in ontogenetic learning, the adaptation process is realized by varying the parameters of a single individual.

As demonstrated by Salimans et al. (2017), modern evolutionary methods represent a scalable alternative to the state-of-the-art reinforcement-learning algorithm (Schulman et al. 2015, 2017). Indeed, they can be used to adapt neural network controllers with millions of parameters by achieving results that are competitive with reinforcement-learning methods. The results have been collected on state-of-the-art benchmarking problems: the Mujoco control problems that require controlling articulated robots (Todorov, Erez, and Tassa 2012) and the Atari games that require controlling game players that receive as input the images of the console (Bellemare et al. 2013).

4.9 Conclusion

Evolutionary robotics is not only a method for automatic robot development inspired by biology but also a tool for investigating open questions concerning natural systems such as, for example, the role of embodiment in cognition, the origins of symbolic communication, the relation between behavioral and cognitive capacities, and the mechanisms supporting the development of cooperative behaviors.

Despite initial skepticism demonstrated by representatives of mainstream disciplines and even by pioneers of the approach (Matarić and Cliff 1996), over the years an increasing number of researchers from a wide range of disciplines have adopted the method. The richness and fecundity of the approach combined with the novel opportunities granted by recent methodological progress suggest that it will continue to play an important role in the future.

Readers interested in acquiring hands-on knowledge on evolutionary robotics can access freely available tools that permit the replication of standard experiments and the design of new experiments (see Auerbach et al. 2014; Massera et al. 2014; Nolfi 2021; see also <https://github.com/snolfi/evorobotpy>).

Additional Reading and Resources

- A recent review of the field: Nolfi, S., J. Bongard, P. Husbands, and D. Floreano. 2016. “Evolutionary Robotics.” In *Springer Handbook of Robotics*, edited by Bruno Siciliano and Oussama Khatib, 1423–1541. 2nd ed. Berlin: Springer Verlag.
- An article that illustrates in more detail the complex adaptive system nature of behavior and cognition in embodied agents: Nolfi, S. 2009. “Behavior and Cognition as a Complex Adaptive System: Insights from Robotic Experiments.” In *Handbook of the Philosophy of Science. Volume 10: Philosophy of Complex Systems*, edited by C. Hooker. General editors: Dov M. Gabbay, Paul Thagard, and John Woods. San Diego: Elsevier.
- A more detailed review of the field: Nolfi, S., and D. Florean, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press, 2000.

- Evorobotpy (Nolfi 2021; <https://github.com/snolfi/evorobotpy2>) is a simple and well-documented tool that can be used to perform evolutionary robotics experiments. The associated documentation (Nolfi 2021, chap. 13) includes tutorials and exercises.
- Farsa (Massera et al. 2014; <https://sourceforge.net/projects/farsa/>) is another software tool that can be used to conduct evolutionary robotics experiments.

References

- André, Jean-Baptiste, and Stefano Nolfi. 2016. "Evolutionary Robotics Simulations Help Explain Why Reciprocity Is Rare in Nature." *Scientific Reports* 6:32785.
- Arbib, Michael A. 1989. *The Metaphorical Brain: Vol. 2, Neural Networks and Beyond*. New York: Wiley.
- Auerbach, Joshua, Deniz Aydin, Andrea Maesani, Przemyslaw Kornatowski, Titus Cieslewski, Grégoire Heitz, Pradeep Fernando, Ilya Loshchilov, Ludovic Daler, and Dario Floreano. 2014. "RoboGen: Robot Generation through Artificial Evolution." In *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, edited by H. Sayama, J. Reiffel, S. Risi, R. Doursat, and H. Lipson. New York: MIT Press.
- Auerbach, Joshua E., and Joshua C. Bongard. 2012. "On the Relationship between Environmental and Morphological Complexity in Evolved Robots." In *Proceedings of the 14th International Conference on Genetic and Evolutionary Computation*, 521–528. ACM Press.
- Axelrod, Robert, and William Donald Hamilton. 1981. "The Evolution of Cooperation." *Science* 211:1390–1396.
- Baldassarre, Gianluca, Vito Trianni, Michael Bonani, Francesco Mondada, Marco Dorigo, and Stefano Nolfi. 2007. "Self-Organised Coordinated Motion in Groups of Physically Connected Robots." *IEEE Transactions on Systems, Man, and Cybernetics* 37 (1): 224–239.
- Bellemare, Marc G., Yavar Naddaf, Joel Veness, and Michael Bowling. 2013. "The Arcade Learning Environment: An Evaluation Platform for General Agents." *Journal of Artificial Intelligence Research* 47:253–279.
- Bernstein, Nikolai. 1967. *The Co-ordination and Regulation of Movements*. Oxford: Pergamon Press.
- Berthoz, Alain. 2000. *The Brain's Sense of Movement*. Cambridge, MA: Harvard University Press.
- Bongard, Joshua C. 2011. "Morphological Change in Machines Accelerates the Evolution of Robust Behavior." *Proceedings of the National Academy of Science* 108:1234–1239.
- Bongard, Josh, Victor Zykov, and Hod Lipson. 2006. "Resilient Machines through Continuous Self-Modelling." *Science* 5802 (314): 1118–1121.
- Brockman, Greg, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. "OpenAI Gym." ArXiv: 1606.01540.
- Brunyé, Tad T., Aaron L. Gardony, Amanda Holmes, and Holly A. Taylor. 2018. "Spatial Decision Dynamics during Wayfinding: Intersection Prompt the Decision-Making Process." *Cognitive Research: Principles and Implications* 8:13.
- Camazine, Scott, Jean-Louis Deneubourg, Nigel R. Franks, James Sneyd, Eric Bonabeau, and Guy Theraula. 2001. *Self-Organization in Biological Systems*. Princeton, NJ: Princeton University Press.
- Cangelosi, Angelo, and Domenico Parisi. 2002. *Simulating the Evolution of Language*. Berlin: Springer Verlag.
- Carvalho, Jônata Tyska, and Stefano Nolfi. 2016. "Cognitive Offloading Does Not Prevent but Rather Promotes Cognitive Development." *PLoS One* 11 (8): e0160679.
- Chestnutt, Joel, Manfred Lau, German Cheung, James Kuffner, Jessica Hodgins, and Takeo Kanade. 2005. "Footstep Planning for the Honda ASIMO Humanoid." In *Proceedings of the IEEE International Conference on Robotics and Automation*. New York: IEEE.
- Chiel, Hillel J., and Randall D. Beer. 1997. "The Brain Has a Body: Adaptive Behavior Emerges from Interactions of Nervous System, Body and Environment." *Trends in Neuroscience* 20:553–557.
- Clark, Andy. 1998. *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press.
- Clark, Andy. 1999. "An Embodied Cognitive Science?" *Trends in Cognitive Sciences* 3:345–351.
- Cliff, Dave, and Geoffrey F. Miller. 1995. "Tracking the Red Queen: Measurement of Adaptive Progress in Coevolutionary Simulations." In *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life*, edited by F. Moran, A. Moreno, J. J. Merelo, and P. Chacon. Berlin: Springer Verlag.
- Collins, Steve, Andy Ruina, Russ Tedrake, and Martijn Wisse. 2005. "Efficient Bipedal Robots Based on Passive-Dynamic Walkers." *Science* 307 (5712): 1082–1085.

- Cully, Antoine, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. "Robots That Can Adapt Like Animals." *Nature* 521:503–507.
- Dawkins, Richard, and John Richard Krebs. 1979. "Arms Races between and within Species." *Proceedings of the Royal Society of London B* 205:489–511.
- Decaestecker, Ellen, Sabrina Gaba, Joost A. M. Raeymaekers, Robby Stoks, Liesbeth Van Kerckhoven, Dieter Ebert, and Luc De Meester. 2007. "Host-Parasite 'Red Queen' Dynamics Archived in Pond Sediment." *Nature* 450 (6): 870–874.
- De Greef, Joachim, and Stefano Nolfi. 2010. "Evolution of Implicit and Explicit Communication in a Group of Mobile Robots." In *Evolution of Communication and Language in Embodied Agents*, edited by S. Nolfi and M. Mirolli. Berlin: Springer Verlag.
- Dewey, John. 1981 [1986]. "The Reflex Arc in Psychology." *Psychological Review* 3:357–370. Reprinted in *The Philosophy of John Dewey*, edited by J. J. McDermott, 136–148. Chicago: University of Chicago Press, 1986.
- Dürr, Peter, Claudio Mattiussi, and Dario Floreano. 2006. "Neuroevolution with Analog Genetic Encoding." In *Proceedings of the Ninth Conference on Parallel Problem Solving from Nature*. Vol. 9. Berlin: Springer.
- Ferrante, Eliseo, Ali Emre Turgut, Edgar Duéñez-Guzmán, Marco Dorigo, and Tom Wenseleers. 2015. "Evolution of Self-Organized Task Specialization in Robot Swarms." *PLoS Computational Biology* 11 (8): e1004273.
- Floreano, Dario, Peter Dürr, and Claudio Mattiussi. 2008. "Neuroevolution: From Architectures to Learning." *Evolutionary intelligence* 1 (1): 47–62.
- Floreano, Dario, Sara Mitri, Stéphane Magnenat, and Laurent Keller. 2007. "Evolutionary Conditions for the Emergence of Communication in Robots." *Current Biology* 17:514–519.
- Floreano, Dario, and Stefano Nolfi. 1997. "Adaptive Behavior in Competing Co-evolving Species." In *Proceedings of the Fourth Conference on Artificial Life*, edited by P. Husband and I. Harvey, 378–387. Cambridge, MA: MIT Press.
- Gigliotta, Onofrio, Giovanni Pezzulo, and Stefano Nolfi. 2011. "Evolution of a Predictive Internal Model in an Embodied and Situated Agent." *Theory in Biosciences* 130 (4): 259–276.
- Gilbert, Sam J. 2015a. "Strategic Offloading of Delayed Intentions into the External Environment." *Quarterly Journal of Experimental Psychology* 68 (5): 971–992.
- Gilbert, Sam J. 2015b. "Strategic Use of Reminders: Influence of Both Domain-General and Task-Specific Metacognitive Confidence, Independent of Objective Memory Ability." *Consciousness and Cognition* 33:245–260.
- Gipson, James J. 1979. *The Ecological Approach to Visual Perception*. Hillsdale, NJ: Lawrence Erlbaum.
- Goldberg, David E., and John Henry Holland. 1988. "Genetic Algorithms and Machine Learning." *Machine Learning* 3 (2–3): 95–99.
- Ha, David, and Jürgen Schmidhuber. 2018. "World Models." arXiv: 1803.10122.
- Hansen, Nikolaus, and Andreas Ostermeier. 2001. "Completely Derandomized Self-Adaptation in Evolution Strategies." *Evolutionary Computation* 9 (2): 159–195.
- Hiller, Jonathan, and Hod Lipson. 2012. "Automatic Design and Manufacture of Soft Robots." *IEEE Transactions on Robotics* 28 (2): 457–466.
- Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. "Long Short-Term Memory." *Neural Computation* 9 (8): 1735–1780.
- Iida, F., and R. Pfeifer. 2005. "Morphological Computation: Connecting Body, Brain and Environment." *Japanese Scientific Monthly* 58 (2): 48–54.
- Kempka, Michał, Marek Wydmuch, Grzegorz Runc, Jakub Toczek, and Wojciech Jaśkowski. 2016. "Vizdoom: A Doom-Based AI Research Platform for Visual Reinforcement Learning." In *IEEE Conference on Computational Intelligence and Games*, 341–348. New York: IEEE.
- Kingma, Diederik P., and Max Welling. 2013. "Auto-Encoding Variational Bayes." arXiv: 1312.6114v10.
- Konczak, Jürgen, Maike Borutta, and Johannes Dichgans. 1997. "The Development of Goal-Directed Reaching in Infants. II. Learning to Produce Task-Adequate Patterns of Joint Torque." *Experimental Brain Research* 113 (3): 465–474.
- Konczak, Jürgen, Maike Borutta, Helge Topka, and Johannes Dichgans. 1995. "The Development of Goal-Directed Reaching in Infants: Hand Trajectory Formation and Joint Torque Control." *Experimental Brain Research* 106 (1): 156–168.
- Konczak, Jürgen, and Johannes Dichgans. 1997. "The Development toward Stereotypic Arm Kinematics during Reaching in the First 3 Years of Life." *Experimental Brain Research* 117 (2): 346–354.
- Kriegman, Sam, Douglas Blackiston, Michael Levin, and Josh Bongard. 2020. "A Scalable Pipeline for Designing Reconfigurable Organisms." *Proceedings of the National Academy of Sciences* 117 (4): 1853–1859.

- Kriegman, Sam, Nick Cheney, and Josh Bongard. 2018. "How Morphological Development Can Guide Evolution." *Scientific Reports* 8:13934.
- Lipson, Hod, and Jordan B. Pollack. 2000. "Automatic Design and Manufacture of Artificial Lifeforms." *Nature* 406: 974–978.
- Long, John. 2012. *Darwin's Devices: What Evolving Robots Can Teach Us about the History of Life and the Future of Technology*. New York: Basic Books.
- Massera, Gianluca, Tomassino Ferrauto, Onofrio Gigliotta, and Stefano Nolfi. 2014. "Designing Adaptive Humanoid Robots through the FARSA Open-Source Framework." *Adaptive Behavior* 22 (3): 255–265.
- Martin, John H. 2005. "The Corticospinal System: From Development to Motor Control." *Neuroscientist* 11 (2): 161–173.
- Matarić, Maja, and Dave Cliff. 1996. "Challenges in Evolving Controllers for Physical Robots." *Robotics and Autonomous Systems* 19 (1): 67–83.
- Maturana, Humberto R., and Francisco J. Varela. 1987. *The Tree of Knowledge: The Biological Roots of Human Understanding*. Boston: Shambhala.
- McGeer, Tad. 1990. "Passive Dynamic Walking." *International Journal of Robotics Research* 9 (2): 62–82.
- Mitri, Sara, Dario Floreano, and Laurent Keller. 2009. "The Evolution of Information Suppression in Communicating Robots with Conflicting Interests." *PNAS* 106 (37): 15786–15790.
- Noë, Alva. 2004. *Action in Perception*. Cambridge, MA: MIT Press.
- Nolfi, Stefano. 1996. "Adaptation as a More Powerful Tool than Decomposition and Integration." In *Proceedings of the Workshop on Evolutionary Computing and Machine Learning, 13th International Conference on Machine Learning*, edited by T. Fogarty and G. Venturini. Bari, Italy: University of Bari Aldo Moro.
- Nolfi, Stefano. 2005. "Categories Formation in Self-Organizing Embodied Agents." In *Handbook of Categorization in Cognitive Science*, edited by H. Cohen and C. Lefebvre. Oxford: Elsevier.
- Nolfi, Stefano. 2009. "Behavior and Cognition as a Complex Adaptive System: Insights from Robotic Experiments." In Vol. 10, *Handbook of the Philosophy of Science: Philosophy of Complex Systems*, edited by C. Hooker. San Diego: Elsevier.
- Nolfi, Stefano. 2013. "Emergence of Communication and Language in Evolving Robots." In *New Perspectives on the Origins of Language*, edited by C. Lefebvre, B. Comrie, and H. Cohen, 533–554. Amsterdam: John Benjamins.
- Nolfi, Stefano. 2021. *Behavioral and Cognitive Robotics: An Adaptive Perspective*. Roma, Italy: Institute of Cognitive Sciences and Technologies, National Research Council (CNR-ISTC). <https://bacrobotics.com/>.
- Nolfi, Stefano, Josh Bongard, Phil Husbands, and Dario Floreano. 2016. "Evolutionary Robotics." In *Handbook of Robotics*, 2nd ed., edited by B. Siciliano and O. Khatib. Berlin: Springer Verlag.
- Nolfi, Stefano, and Dario Floreano. 1998. "Co-evolving Predator and Prey Robots: Do 'Arm Races' Arise in Artificial Life?" *Artificial Life* 4 (4): 311–335.
- Nolfi, Stefano, and Dario Floreano. 1999. "Learning and Evolution." *Autonomous Robots* 7 (1): 89–113.
- Nolfi, Stefano, and Dario Floreano. 2000. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press.
- Nolfi, Stefano, Orazio Miglino, and Domenico Parisi. 1994. "Phenotypic Plasticity in Evolving Neural Networks." In *Proceedings of the International Conference from Perception to Action*, edited by D. P. Gaussier and J-D. Nicoud, 146–157. Los Alamitos, CA: IEEE Computer Society Press.
- Nolfi, Stefano, and Marco Mirolli. 2010. *Evolution of Communication and Language in Embodied Agents*. Berlin: Springer Verlag.
- Nolfi, Stefano, and Domenico Parisi. 1993. "Auto-teaching: Networks That Develop Their Own Teaching Input." In *Proceedings of the Second European Conference on Artificial Life*, edited by J. L. Deneubourg, H. Bersini, S. Goss, G. Nicolis, and R. Dagonnier, 845–862. Brussels: Université Libre de Bruxelles.
- O'Regan, J. K. and A. Noë. 2001. "A Sensorimotor Account of Vision and Visual Consciousness." *Behavioral and Brain Sciences* 24 (5): 939–973.
- Pagliuca, Paolo, and Stefano Nolfi. 2018. "Robust Optimization through Neuroevolution." *PLoS One* 14 (3): e0213193.
- Paul, Chandana. 2006. "Morphological Computation: A Basis for the Analysis of Morphology and Control Requirements." *Robotic Autonomous Systems* 54 (8): 619–630.
- Pfeifer, Rolf, and Joshua C. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, Rolf, Fumiya Iida, and Gabriel Gómez. 2006. "Morphological Computation for Adaptive Behavior and Cognition." In *International Congress Series*, Vol. 1291, 22–29. San Diego: Elsevier Press.

- Purves, Dale. 1994. *Neural Activity in the Growth of the Brain*. Cambridge: Cambridge University Press.
- Quartz, Steven R., and Terrence J. Sejnowski. 1997. "The Neural Basis of Cognitive Development: A Constructivist Manifesto." *Behavioral and Brain Science* 4:537–555.
- Rechenberg, Ingo. 1973. *Evolutionsstrategie—Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*. Stuttgart: Frommann-Holzboog.
- Rezende, Danilo Jimenez, Shakir Mohamed, and Daan Wierstra. 2014. "Stochastic Backpropagation and Approximate Inference in Deep Generative Models." In Vol. 32, *Proceedings of the 31st International Conference on Machine Learning*. Red Hook, NY: JMLR.
- Rosin, Christopher D., and Richard K. Belew. 1997. "New Methods for Competitive Coevolution." *Evolutionary Computation* 5 (1): 1–29.
- Salimans, Tim, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. "Evolution Strategies as a Scalable Alternative to Reinforcement Learning." ArXiv:1703.03864v2.
- Sandini, Giulio, Giorgio Metta, and David Vernon. 2004. "Robotcub: An Open Framework for Research in Embodied Cognition." *International Journal of Humanoid Robotics* 8 (2): 18–31.
- Savastano, Piero, and Stefano Nolfi. 2013. "A Robotic Model of Reaching and Grasping Development." *IEEE Transactions on Autonomous Mental Development* 4 (5): 326–336.
- Scheier, Christian, Rolf Pfeifer, and Yasuo Kuniyoshi. 1998. "Embedded Neural Networks: Exploiting Constraints." *Neural Networks* 11:1551–1596.
- Schembri, Massimiliano, Marco Mirolli, and Gianluca Baldassarre. 2007. "Evolution and Learning in an Intrinsically Motivated Reinforcement Learning Robot." In *Proceedings of the European Conference on Artificial Life*. Berlin: Springer.
- Schlesinger, Matthew. 2004. "Evolving Agents as a Metaphor for the Developing Child." *Developmental Science* 7:154–168.
- Schulman, John, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. 2015. "Trust Region Policy Optimization." In *Proceedings of the 32nd International Conference on Machine Learning*, 1889–1897. PMLR.
- Schulman, John, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. "Proximal Policy Optimization Algorithms." ArXiv: 1707.06347.
- Schwefel, Hans-Paul. 1997. *Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie*. Vol. 26. Basel/Stuttgart: Birkhaeuser.
- Simione, Luca, and Stefano Nolfi. 2017. "Achieving Long-Term Progress in Competitive Co-evolution." In *Proceedings of the IEEE Symposium on Computational Intelligence*. New York: IEEE.
- Simione, Luca, and Stefano Nolfi. 2019. "Long-Term Progress and Behavior Complexification in Competitive Co-evolution." ArXiv: 1909.08303.
- Sims, Karl. 1994. "Evolving 3D Morphology and Behavior by Competition." *Artificial Life* 4:28–39.
- Sinervo, Barry, and Curt M. Lively. 1996. "The Rock-Paper-Scissors Game and the Evolution of Alternative Male Strategies." *Nature* 380:240–243.
- Spencer, John P., and Esther Thelen. 2000. "Spatially Specific Changes in Infants' Muscle Coactivity as They Learn to Reach." *Infancy* 1 (3): 275–302.
- Sperati, Valerio, Vito Trianni, and Stefano Nolfi. 2011. "Self-Organised Path Formation in a Swarm of Robots." *Swarm Intelligence* 5:97–119.
- Spivey, Michael. 2007. *The Continuity of Mind*. New York: Oxford University Press.
- Stanley, Kenneth O., and Risto Miikkulainen. 2002. "Evolving Neural Networks through Augmenting Topologies." *Evolutionary Computation* 10 (2): 99–127.
- Sutton, Richard S., and Andrew G. Barto. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press.
- Thelen, Esther, and Linda B. Smith. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press.
- Todorov, Emanuel, Tom Erez, and Yuval Tassa. 2012. "Mujoco: A Physics Engine for Model-Based Control." In *Proceedings of the 2012 IEEE/RSJ Intelligent Robots and Systems Conference*. New York: IEEE.
- Tuci, Elio, Gianluca Massera, and Stefano Nolfi. 2010. "Active Categorical Perception of Object Shapes in a Simulated Anthropomorphic Robotic Arm." *Transaction on Evolutionary Computation Journal* 14 (6): 885–899.
- Varela, Francisco J., Evan Thompson, and Eleanor Rosch. 1991. *The Embodied Mind*. Cambridge, MA: MIT Press.
- von Hofsten, Claes. 1982. "Eye-Hand Coordination in the Newborn." *Developmental Psychology* 18 (3): 450–461.
- von Hofsten, Claes, and Louise Rönnqvist. 1993. "The Structuring of Neonatal Arm Movements." *Child Development* 64 (4): 1046–1057.

