37. Embodied Intelligence

Angelo Cangelosi, Josh Bongard, Martin H. Fischer, Stefano Nolfi

Embodied intelligence is the computational approach to the design and understanding of intelligent behavior in embodied and situated agents through the consideration of the strict coupling between the agent and its environment (situatedness), mediated by the constraints of the agent's own body, perceptual and motor system, and brain (embodiment). The emergence of the field of embodied intelligence is closely linked to parallel developments in computational intelligence and robotics, where the focus is on morphological computation and sensory–motor coordination in evolutionary robotics models, and in neuroscience and cognitive sciences where the focus is on embodied cognition and developmental robotics models of embodied symbol learning. This chapter provides a theoretical and technical overview of some principles of embodied intelligence, namely morphological computation, sensory–motor coordination, and developmental embodied cognition. It will also discuss some tutorial examples on the modeling of body/brain/environment adaptation for the evolution of morphological computational agents, evolutionary robotics model of navigation and object discrimination, and developmental robotics models of language and numerical cognition in humanoid robots.

37.1 Introduction to Embodied Intelligence

Organisms are not isolated entities which develop their sensory–motor and cognitive skills in isolation from their social and physical environment, and independently from their motor and sensory systems. On the contrary, behavioral and cognitive skills are dynamical properties that unfold in time and arise from a large number of interactions between the agents' nervous system, body, and environment [37.1–7]. Embodied intelligence is the computational approach to the design and understanding of intelligent behavior in embodied and situated agents through the consideration of the strict coupling between the agent and its environment (situatedness), mediated by the constraints of the agent's own body, perceptual and motor system, and brain (embodiment).

Historically, the field of embodied intelligence has its origin from the development and use of bio-inspired computational intelligence methodologies in computer science and robotics, and the overcoming of the limitations of symbolic approaches typical of classical artificial intelligence methods. As argued in *Brooks*' [37.2] seminal paper on *Elephants don't play chess*, the study of apparently simple behaviors, such as locomotion and motor control, permits an understanding of the embodied nature of intelligence, without the requirement to start from higher order abstract skills as those involved in chess playing algorithms. Moreover, the emergence of the field of embodied intelligence is closely linked to parallel developments in robotics, with the focus on morphological computation and sensory–motor coordination in evolutionary and developmental robotics models, and in neuroscience and cognitive sciences with the focus on embodied cognition (EC).

The phenomenon of *morphological computation* concerns the observation that a robot's (or animal's) *body plan* may perform computations: A body plan that allows the robot (or animal) to passively exploit interactions with its environment may perform computations that lead to successful behavior; in another body plan less well suited to the task at hand, those computations would have to be performed by the control policy [37.8–10]. If both the body plans and control policies of robots are evolved, evolutionary search may find robots that exhibit more morphological computation than an equally successful robot designed by hand (see more details in Sect. 37.2).

The principle of *sensory–motor coordination*, which concerns the relation between the characteristics of the agents' control policy and the behaviors emerging from agent/environmental interactions, has been demonstrated in numerous evolutionary robotics models [37.6]. Experiments have shown how adaptive agents can acquire an ability to coordinate their sensory and motor activity so as to self-select their forthcoming sensory experiences. This sensory–motor coordination can play several key functions such as enabling the agent to access the information necessary to make the appropriate behavioral decision, elaborating sensory information, and reducing the complexity of the agents' task to a manageable level. These two themes will be exemplified through the illustration of evolutionary robotics experiments in Sect. 37.3 in which the fine-grained characteristics of the agents' neural control system and body are subjected to variations (e.g. gene mutation) and in which variations are retained or discarded on the basis of their effects at the level of the overall behavior exhibited by the agent in interaction with the environment.

In cognitive and neural sciences, the term *embodied cognition* (EC) [37.11, 12] is used to refer to systematic relationships between an organism's cognitive processes and its perceptual and response repertoire. Notwithstanding the many interpretations of this term [37.13], the broadest consensus of the proponents of EC is that our knowledge representations encompass the bodily activations that were present when we initially acquired this knowledge (for differentiations, [37.14]). This view helps us to understand the many findings of modality-specific biases induced by cognitive computations. Examples of EC in psychology and cognitive science can be sensory–motor (e.g., a systematic increase in comparison time with angular disparity between two views of the same object [37.15]), or conceptual (e.g., better recall of events that were experienced in the currently adopted body posture [37.16]), or emotional in nature (e.g., interpersonal warmth induced by a warm handheld object [37.17]). Such findings were hard to accommodate under the more traditional views where knowledge was presumed symbolic, amodal and abstract and thus dissociated from sensory input and motor output processes.

Embodied cognition experiments in psychology have inspired the design of developmental robotics models [37.18] which exploit the ontogenetic interaction between the developing (baby) robot and its social and physical environment to acquire both simple sensory–motor control strategies and higher order capabilities such as language and number learning (Sect. 37.4).

To provide the reader with both a theoretical and technical understanding of the principles of morphological computation, sensory–motor coordination and developmental EC the following three sections will review the progress in these fields, and analyze in detail some key studies as examples. The presentation of studies on the modeling of both sensory–motor tasks (such as locomotion, navigation, and object discrimination) and of higher order cognitive capabilities (such as linguistic and numerical cognition) demonstrates the impact of embodied intelligence in the design of a variety of perceptual, motor, and cognitive skills.

37.2 Morphological Computation for Body–Behavior Coadaptation

Embodied intelligence dictates that there are certain body plans and control policies that, when combined, will produce some desired behavior. For example, imagine that the desired task is active categorical perception (ACP) [37.19, 20]. ACP requires a learner to actively interact with objects in its environment to classify those objects. This stands in contrast to passive categorization whereby an agent observes objects from a distance – perhaps it is fed images of objects or views them through a camera – and labels the objects according to their perceived class. In order for an animal or robot to perform ACP, it must not only possess a control policy that produces as output the correct class for the object being manipulated, but also some manipulator with which to physically affect (and be affected by) the object.

One consequence of embodied intelligence is that certain pairings of *body* and *brain* produce the desired behavior, and others do not. Returning to the example of ACP, if a robot's arm is too short to reach the objects then it obviously will not be able to categorize them. Imagine now a second robot that possesses an arm of the requisite length but can only bring the back of its hand into contact with the objects. Even if this robot's control policy distinguishes between round and edged objects based on the patterned firing of touch sensors embedded in its palm, this robot will also not be able to perform ACP.

A further consequence of embodied intelligence is that some body plans may require a complex control policy to produce successful behavior, while another body plan may require a simpler control policy. This has been referred to as the morphology and control tradeoff in the literature [37.7]. Continuing the ACP example, consider a third robot that can bring its palm and fingers into contact with the objects, but only possesses a single binary touch sensor in its palm. In order to distinguish between round and edged objects, this robot will require a control policy that performs some complex signal processing on the time series data produced by this single sensor during manipulation. A fourth robot however, equipped with multiple tactile sensors embedded in its palm and fingers, may be able to categorize objects immediately after grasping them: Perhaps round objects produce characteristic static patterns of tactile signals that are markedly different from those patterns produced when grasping edged objects.

The morphology and control tradeoff however raises the question as to what is being traded. It has been argued that what is being traded is computation [37.7, 8]. If two robots succeed at a given task, and each robot is equipped with the simplest control policy that will allow that robot to succeed, but one control policy performs fewer computations than the other control policy, then the body plan of the robot equipped with the simpler control policy must perform the *missing* computations required to succeed at the task.

This phenomenon of a robot's (or animal's) body plan performing computation has been termed *morphological computation* [37.8–10]. *Paul* [37.8] outlined a theoretical robot that uses its body to compute the XOR function. In another study [37.9] it was shown how the body of a vacuum cleaning robot could literally replace a portion of its artificial neural network controller, thus subsuming the computation normally performed by that part of the control policy into the robot's body. *Pfeifer* and *Gomez* [37.21] describe a number of other robots that exhibit the phenomenon of morphological computation.

37.2.1 The Counterintuitive Nature of Morphological Computation

All of the robots outlined by *Pfeifer* and *Gomez* [37.21] were designed manually; in some cases the control policies were automatically optimized. If for each task there are a spectrum of robot body plan/control policy pairings that achieve the task, one might ask where along this spectrum the human-designed robots fall. That is, what mixtures of morphological computation and control computation do human designers tend to favor? The bulk of the artificial intelligence literature, since the field's beginnings in the 1950s, seems to indicate that humans exhibit a cognitive chauvinism: we tend to favor control complexity over morphological complexity. Classical artificial intelligence dispensed with the body altogether: it was not until the 1980s that the role of morphology in intelligent behavior was explicitly stated [37.2]. As a more specific example, object manipulation was first addressed by creating rigid, articulated robot arms that required complex control policies to succeed [37.22]. Later, it was realized that soft manipulators could simplify the amount of control required for successful manipulation (e.g., [37.23]). Most recently, a class of robot manipulators known as *jamming grippers*' was introduced [37.24]. In a jamming gripper, a robot arm is tipped with a bag of granular material such that when air is removed from the bag the grains undergo a phase transition into a *jammed*, solid-like state. The control policies for jamming grippers are much simpler than those required for rigid or even soft multifingered dexterous manipulators: at the limit, the controller must switch the manipulator between just two states (*grip* or *release*), regardless of the object.

Despite the fact that the technology for creating jamming grippers has existed for decades, it took a long time for this class of manipulators to be discovered. In other branches of robotics, one can discern a similar historical pattern: new classes of robot body plan were successively proposed that required less and less explicit control. In the field of legged locomotion for example, robots with *whegs* (wheel-leg hybrids) were shown to require less explicit control than robots with legs to enable travel over rough terrain [37.25].

These observations suggest that robots with more morphological computation are less intuitive for humans to formulate and then design than robots with less morphological computation. However, there may be a benefit to creating robots that exhibit significant amounts of morphological computation. For example, hybrid dynamic walkers require very little control and are much more energy efficient compared to fully actuated legged robots [37.26]. It has been argued that tensegrity robots also require relatively little control compared to robots composed of serially linked rigid components, and this class of robot has several desirable properties such as the ability to absorb and recover from external perturbations [37.9].

So, if robots that exhibit morphological computation are desirable, yet it is difficult for humans to navigate in this part of the space of possible robots, can an automated search method be used to discover such robots?

37.2.2 Evolution and Morphological Computation

One of the advantages of using evolutionary algorithms to design robots, compared to machine learning methods, is that both the body plan and the control policy can be placed under evolutionary control [37.27]. Typically, machine learning methods optimize some of the parameters of a control policy with a fixed topology. However, if the body plans and control policies of robots are evolved, and there is sufficient variation within the population of evolving robots, search may discover multiple successful robots that exhibit varying degrees of morphological computation. Or, alternatively, if morphological computation confers a survival advantage within certain contexts, a phylogeny of robots may evolve that exhibit increasing amounts of morphological computation.

A recent pair of experiments illustrates how morphological computation may be explored. An evolutionary algorithm was employed to evolve the body plans

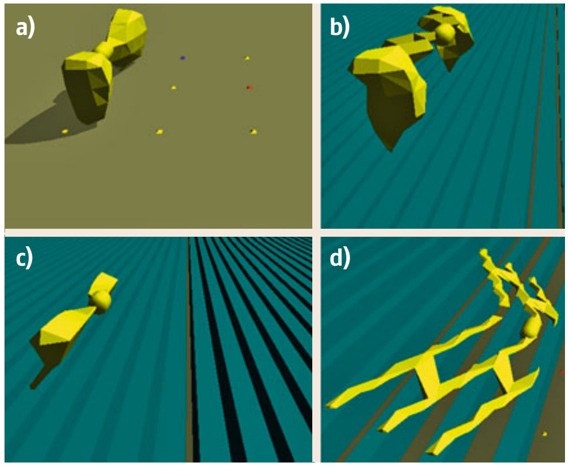

Fig. 37.1a–d A sample of four evolved robots with differing amounts of morphological complexity. **(a)** A simple-shaped robot that evolved to locomote over flat ground. **(b–d)** Three sample robots, more morphologically complex than the robot in **(a)**, that evolved in icy environments (after *Auerbach* and *Bongard* [37.28]). To view videos of these robots see [37.29]

and control policies of robots that must move in one of two environments. The first environment included nothing else other than a flat, high-friction ground plane (Fig. 37.1a). The second environment was composed of a number of low-friction bars that sit atop the high-friction ground plane (Fig. 37.1b–d). These bars can be thought of as ice distributed across a flat landscape. In order for robots to move across the icy terrain, they must evolve appendages that are able to reach down between the icy blocks, come into contact with the high-friction ground, and push or pull themselves forward.

It was found that robots evolved to travel over the ice had more complex shapes than those evolved to travel over flat ground (compare the robot in Fig. 37.1a to those in Fig. 37.1b–d) [37.28]. However, it was also found that the robots that travel over ice had fewer mechanical degrees of freedom (DOFs) than the robots evolved to travel over flat ground [37.30]. If a robot possesses fewer mechanical DOFs, one can conclude that it has a simpler control policy, because there are fewer motors to control. It seems that the robots evolved to travel over ice do so in the following manner: the complex shapes of their appendages cause the appendages to *reach* down into the crevices between the ice without explicit control; the simple control policy then simply sweeps the appendages back and forth, horizontally, to in effect *skate* along the tops of the ice. In contrast, robots evolved to travel over flat ground must somehow push back, reach up, and pull forward – using several mechanical DOFs – to move forward.

One could conclude from these experiments that the robots evolved to travel over ice perform more morphological computation than those evolved to travel over flat ground: the former robots have more complex bodies but simpler control policies than the latter robots, yet both successfully move in their environments. Much more work is required to generalize this result to different robots, behaviors, and environments, but this initial work suggests that evolutionary robotics may be a unique tool for studying the phenomenon of morphological computation.

37.3 Sensory–Motor Coordination in Evolving Robots

The actions performed by embodied and situated agents inevitably modify the agent–environmental relation and/or the environment. The type of stimuli that an agent will sense at the next time step at t_{+1} crucially depends, for example, on whether the agent turns left or right at the current time t. Similarly, the stimuli that an agent will experience next at time t_{+1} when standing next to an object depend on the effort with which it will push the object at time t. This implies that actions might play direct and indirect adaptive roles. Actions playing a direct role are, for example, foraging or predator escaping behaviors that directly impact on the agent's own survival chances. Action playing indirect roles consists, for example, in wandering through the environment to spot interesting sensory information (e.g., the perception of a food area that might eventually afford foraging actions) or playing a fighting game with a conspecific that might enable the agent to acquire capacities that might later be exploited to deal with aggressive individuals. The possibility to self-select useful sensory stimuli through action is referred with the term sensory–motor coordination.

Together with morphological computation, sensory–motor coordination constitutes a fundamental property of embodied and situated agents and one of most important characteristic that can be used to differentiate these systems from alternative forms of intelligence. In the following sections, we illustrate three of the key roles that can be played by sensory–motor coordination:

i) The discovery of parsimonious behavioral strategies
ii) The access and generation of useful sensory information through action and active perception
iii) The constraining and channeling of the learning process during evolution and development.

37.3.1 Enabling the Discovery of Simple Solutions

Sensory–motor coordination can be exploited to find solutions relying on more parsimonious control policies than alternative solutions not relying, or relying less, on this principle. An example is constituted by a set of experiments in which a Khepera robot [37.31] endowed with infrared and speed sensors, has been evolved for the ability to remain close to large cylindrical objects (food) while avoiding small cylindrical objects (dangers). From a passive perspective, that does not take into account the possibility to exploit sensory–motor coordination, 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 [37.32]. On the other hand, robots evolved for the ability to perform this task tend to converge on a solution relying on a rather simple control policy: the robots begin to turn around objects as soon as they approach them and then discriminate the size of the object on the basis of the sensed differential speed of the left and right wheels during the execution of the object-circling behavior [37.33]. In other words, 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.

Another related experiment in which a Khepera robot provided solely with infrared sensors was adapted for finding and remaining close to a cylindrical object, while avoiding walls, demonstrates how sensory–motor coordination can be exploited to solve tasks that require the display of differentiated behavior in different environmental circumstances, without discriminating the contexts requiring different responses [37.32, 34]. Indeed, evolved robots manage to avoid walls, find a cylindrical object, and remain near it simply by moving backward or forward when their frontal infrared sensors are activated or not, respectively, and by turning left or right when their right and left infrared sensors are activated, respectively (providing that the turning speed and the move forward speed is appropriately regulate on the basis of the sensors activation). Indeed, the execution of this simple control rule combined with the effects of the robot's actions lead to the exhibition of a move-forward behavior far from obstacles, an obstacle avoidance behavior near walls, and an oscillatory behavior near cylindrical objects (in which the robot remains near the object by alternating forward and backward and/or turn-left and turn-right movements). The differentiation of the behavior observed during the robot/wall and robot/cylinder interactions can be explained by considering that the execution of the same action produces different sensory effects in interaction with different objects. In particular, the execution of a turn-left action at time t elicited by the fact that the right infrared sensors are more activated than the left sensors near an object leads to the perception of: (i) a similar sensory stimulus eliciting a similar action at time t_{+1}, ultimately producing an object avoidance behavior near a wall object, (ii) a different sensory stimulus (in which left infrared sensors can become more activated than the left infrared sensors) eliciting a turn-right action at time t_{+1} ultimately producing an oscillatory behavior near the cylinder.

Examples of clever use of sensory–motor coordination abound in natural and artificial evolution. A paradigmatic example of the use of sensory–motor coordination in natural organisms are the navigation capabilities of flying insects that are based on the optic flow, i. e., the apparent motion of contrasting objects in the visual field caused by the relative motion of the agent [37.35]. Houseflies, for example, use this solution to navigate up to 700 body lengths per second in unknown 3D environment while using quite modest processing resources, i. e., about 0.001% of the number of neurons present in the human brain [37.36]. Examples in the evolutionary robotics literature include wheeled robots performing navigation tasks ([37.32], see below), artificial fingers and humanoid robotic arms evolved for the ability to discriminate between object varying in shapes [37.20, 37], and wheeled robots able to navigate visually by using a pan-tilt camera [37.38].

37.3.2 Accessing and Generating Information Through Action

A second fundamental role of sensory–motor coordination consists in accessing and/or generating useful sensory information though action. Differently from experimental settings in which stimuli are brought to the passive agent by the experimenter, in ecological conditions agents need to access relevant information through action. For example, infants access the visual information necessary to recognize the 3D structure of an object by rotating it in the hand and by keeping it at close distance so to minimize visual occlusions [37.39]. The use of sensory–motor coordination for this purpose is usually named *active perception* [37.37, 40, 41].

Interestingly, action can be exploited not only to access sensory information but also to generate it. To understand this aspect, we should consider that through their action agents can elaborate the information they access through their sensory system over time and store the result of the elaboration in their body state and/or in their posture or location. A well-known example of this phenomenon is constituted by depth perception as a result of convergence, i. e., the simultaneous inward movement of both eyes toward each other, to maintain a single binocular percept of a selected object. The execution of this behavior produces a kinesthetic sensation in the eye muscles that reliably correlates with the object's depth.

The careful reader might have recognized that the robot's behavioral discrimination strategies to perceive larger and smaller cylindrical objects, described in the previous section, exploit the same active perception mechanism. For a robot provided with infrared and wheel-speed sensors, the perception of object size necessarily requires a capacity to integrate the information provided by several stimuli. The elaboration of this information however is not realized internally, within the robot's nervous system, but rather externally through the exhibition of the object-circling behavior. It is this behavior that generates the corresponding kinesthetic sensation on the wheel sensors that is then used by the robot to decide to remain or leave, depending on the circumstances.

Examples of clever strategies able to elaborate the required information through action and active perception abound in evolutionary robotics experiments. By carrying out an experiment in which a robot needed to reach two foraging areas located in the northeast and southwest side of a rectangular environment surrounded by walls, *Nolfi* [37.34] observed that the evolved robots developed a clever strategy that allows them to compute the relative length of the two sides of the environment and to navigate toward the two right corners on the basis of a simple control policy. The strategy consists in leaving the first encountered corner with an angle of about 45° with respect to the two sides, moving straight, and then eventually following the left side of the next encountered wall ([37.34] for details). Another clever exploitation of sensory–motor coordination was observed in an experiment involving two cooperating robots that helped each other to navigate toward circular target areas [37.42]. Evolved robots discovered and displayed a behavior solution that allowed them to inform each other on the relative location of the center of their target navigation area despite their sensory system being unable to detect their relative position within the area [37.42].

37.3.3 Channeling the Course of the Learning Process

A third fundamental role of sensory–motor coordination consists in channeling the course of the forthcoming adaptive process.

The sensory states experienced during learning crucially determine the course and the outcome of the learning process [37.43]. This implies that the actions displayed by an agent, that co-determine the agent's forthcoming sensory states, ultimately affect how the agent changes ontogenetically. In other words, the behavior exhibited by an agent at a certain stage of its development constraints and channels the course of the agent's developmental process.

Indeed, evolutionary robotics experiments indicate how the evolution of plastic agents (agents that vary their characteristics while they interact with the environment [37.44]) lead to qualitatively different results with respect to the evolution of nonplastic individuals. The traits evolved in the case of nonplastic individuals are selected directly for enabling the agent to display the required capabilities. The traits evolved in the case of plastic individuals, instead, are selected primarily for enabling the agents to acquire the required capabilities through an ontogenetic adaptation process. This implies that, in this case, the selected traits do not enable the agent to master their adaptive task (agents tend to display rather poor performance at the beginning of their lifetime) but rather to acquire the required capacities through ontogenetic adaptation.

More generally, the behavioral strategies adopted by agents at a certain stage of their developmental process can crucially constrain the course of the adaptive process. For example, agents learning to reach and grasp objects might temporarily reduce the complexity of the task to be mastered by freezing (i. e., locking) selected DOFs and by then unfreezing them when their capacity reaches a level that allows them to master the task in its full complexity [37.45, 46]. This type of process can enable exploratory learning by encompassing variation and selection of either the general strategy displayed by the agent or the specific way in which the currently selected strategy is realized.

37.4 Developmental Robotics for Higher Order Embodied Cognitive Capabilities

37.4.1 Embodied Cognition and Developmental Robots

The previous sections have demonstrated the fundamental role of embodiment and of the agent–environment coupling in the design of adaptive agents and robots capable to perform sensory–motor tasks such as navigation and object discrimination. However, embodiment also plays an important role in higher order cognitive capabilities [37.12], such as object categorization and representation, language learning, and processing, and even the acquisition of abstract concepts such as numbers. In this section, we will consider some of the key psychological and neuroscience evidence of EC and its contribution in the design of linguistic and numerical skills in cognitive robots.

Intelligent behavior has traditionally been modeled as a result of activation patterns across distributed knowledge representations, such as hierarchical networks of interrelated propositional (symbolic) nodes that represent objects in the world and their attributes as abstract, amodal (nonembodied) entities [37.47]. For example, the response *bird* to a flying object with feathers and wings would result from perceiving its features and retrieving its name from memory on the basis of a matching process. Such traditional views were attractive for a number of reasons: They followed the predominant philosophical tradition of logical conceptual knowledge organization, according to which all objects are members of categories and category membership can be determined in an all-or-none fashion via defining features. Also, such hierarchical knowledge networks were consistent with cognitive performance in simple tasks such as speeded property verification, which were thought to tap into the retrieval of knowledge. For example, verifying the statement *a bird has feathers* was thought to be easier than verifying the statement *a bird is alive* because the feature *feathers* was presumably stored in memory as defining the category *bird*, while the feature *alive* applies to all animals and was therefore represented at a superordinate level of knowledge, hence requiring more time to retrieve after having just processed *bird* [37.47]. Finally, it was convenient to computationally model such networks by liking the human mind to an information processing device with systematic input, storage, retrieval, and output mechanisms. Thus, knowledge was considered as an abstract commodity independent of the physical device within which it was implemented.

More recent work called into question several of these assumptions about the workings of the human mind. For example, graded category memberships and prototypicality effects in categorization tasks pointed to disparities between the normative logical knowledge organization and the psychological reality of knowledge retrieval [37.48]. Computational modeling of cognitive processes has revealed alternative, distributed representational networks for computing intelligent responses in perceptual, conceptual, and motor tasks that avoid the neurophysiologically implausible assumption of localized storage of specific knowledge [37.49]. Most importantly, though, traditional propositional knowledge networks were limited to explaining the meaning of any given concept in terms of an activation pattern across other conceptual nodes, thus effectively defining the meaning of one symbol in terms of arbitrary other symbols. This process never referred to a concrete experience or event and essentially made the process of connecting internal and external referents arbitrary. In other words, traditional knowledge representations never make contact with specific sensory and motor modalities that is essential to imbue meaning to the activation pattern in a network. This limitation is known as the grounding problem [37.50] and points to a fundamental flaw in traditional attempts to model human knowledge representations.

A second reason for abandoning traditional amodal models of knowledge representation is the fact that these models cannot account for patterns of sensory and motor excitation that occur whenever we activate our knowledge. Already at the time when symbol manipulation approaches to intelligent behavior had their heyday there was powerful evidence for a mandatory link between intelligent thought and sensory–motor experience: When matching two images of the same object, the time we need to recognize that it is the same object is linearly related to the angular disparity between the two views [37.15]. This result suggests that the mental comparison process simulates the physical object rotation we would perform if the two images were manipulable in our hands. In recent years, there has been both more behavioral and also neuroscientific evidence of an involvement of sensory–motor processes in intelligent thought, leading to the influential notion of action simulation as an obligatory component of intelligent thought (for review, [37.51]).

To summarize, the idea that sensory and motor processes are an integral part of our knowledge is driven by both theoretical and empirical considerations. On the theoretical side, the EC stance addresses the grounding problem, a fundamental limitation of classical views of knowledge representation. Empirically, it is tough for traditional amodal conceptualizations of knowledge to address systematic patterns of sensory and motor biases that accompany knowledge activation.

Amongst the latest development in robotics and computational intelligence, the field of developmental robotics has specifically focused on the essential role of EC in the ontogenetic development of cognitive capabilities. Developmental robotics (also know as epigenetic robotics and as the field of autonomous mental development) is the interdisciplinary approach to the autonomous design of behavioral and cognitive capabilities in artificial agents (robots) that takes direct inspiration from the developmental principles and mechanisms observed in natural cognitive systems (children) [37.18, 52–54]. In particular, the key principle of developmental robotics is that the robot, using a set of intrinsic developmental principles regulating the real-time interaction between its body, brain, and environment, can autonomously acquire an increasingly complex set of sensorimotor and mental capabilities. Existing models in developmental robotics have covered the full range of sensory–motor and cognitive capabilities, from intrinsic motivation and motor control to social learning, language and reasoning with abstract knowledge ([37.18] for a full overview).

To demonstrate the benefits of combining EC with developmental robotics in the modeling of embodied intelligence, the two domains of the action bases of language and of the relationship between space and numerical cognition have been chosen. In Sect. 37.4.2, we will look at seminal examples of the embodied bases of language in psycholinguistics, neuroscience, and developmental psychology, and the corresponding developmental robotics models. Section 37.4.3 will consider EC evidence on the link between spatial and numerical cognition, and a developmental robotics model of embodied language learning.

37.4.2 Embodied Language Learning

In experimental psychology and psycholinguistics, an influential demonstration of action simulation as part of language comprehension was first carried out by *Glenberg* and *Kaschak* [37.55]. They asked healthy adults to move their right index finger from a button in their midsagittal plane either away from or toward their body to indicate whether a visually presented statement was meaningful or not. Sentences like *Open the drawer* led to faster initiation of movements toward than away from the body, while sentences like *Close the drawer* led to faster initiation of movements away from than toward the body. Thus, there was a congruency effect between the implied spatial direction of the linguistic description and the movement direction of the reader's motor response. This motor congruency effect in language comprehension has been replicated and extended (for review, [37.56]). It suggests that higher level cognitive feats (such as language comprehension) are ultimately making use of lower level (sensory–motor) capacities of the agent, as predicted by an embodied account of intelligence.

In parallel, growing cognitive neuroscience evidence has shown that the cortical areas of the brain specialized for motor processing are also involved in language processing tasks; thus supporting the EC view that action and language are strictly integrated [37.57, 58]. For example, *Hauk* et al. [37.59] carried out brain imaging experiments where participants read words referring to face, arm, or leg actions (e.g., *lick*, *pick*, *kick*). Results support the embodied view of language, as the linguistic task of reading a word differentially activated parts of the premotor area that were directly adjacent, or overlapped, with region activated by actual movement of the tongue, the fingers, or the feet, respectively.

The embodied nature of language has also been shown in developmental psychology studies, as in *Tomasello*'s [37.60] constructivist theory of language acquisition and in *Smith* and *Samuelson*'s [37.61] study on embodiment biases in early word learning. For example, *Smith* and *Samuelson* [37.61] investigated the role of embodiment factors such as posture and spatial representations during the learning of first words. They demonstrated the importance of the changes in postures involved in the interaction with objects located in different parts (left and right) of the child's peripersonal space. Experimental data with 18-month old children show that infants can learn new names also in the absence of the referent objects, when the new label is said whilst the child looks at the same left/right location where the object has previously appeared. This specific study was the inspiration of a developmental robotics study on the role of posture in the acquisition of object names with the iCub baby robot [37.62].

The iCub is an open source robotic platform developed as a benchmark experimental tool for cognitive and developmental robotics research [37.63]. It has a total of 53 DOF, with a high number of DOF (32) in the arms and hands to study object manipulation and the role of fine motor skills in cognitive development. This facilitates the replication of the experimental setup of *Smith* and *Samuelson*'s study [37.61]. In the iCub experiments, a human tutor shows two novel objects respectively in the left and right location of a table put in front of the robot. Initially the robot moves to look at each object and learns to categorize it according to its visual features, such as shape and color. Subsequently the tutor hides both objects, directs the robot's attention toward the right side where the first object was shown and says a new word: *Modi*. In the test phase both objects are presented simultaneously in the centre of the table, and the robot is asked *Find the modi*. The robot must then look and point at the object that was presented in the right location. Four different experiments were carried out, as in Smith and Samuelson's child study. Two experiments differ with regards to the frequency of the left/right locations used to show each objects: the Default Condition when each object always appears in the same location, and the Switch Condition when the position of the two objects is varied to weaken the object/location spatial association. In the other two experimental conditions, the object is named whilst in sight, so to compare the relative weighting of the embodiment spatial constraints and the time constraint.

The robot's behavior is controlled by a modular neural network consisting of a series of pretrained Kohonen self-organizing maps (SOMs), connected through Hebbian learning weights that are trained online during the experiment [37.64]. The first SOM is a *color map* as it is used to categorize objects according to their color (average RGB (red-green-blue) color of the foveal area). The second map, the *auditory map*, is used to represent the words heard by the robot, as the result of the automatic speech recognition system. The other SOM is the *body-hub map*, and this is the key component of the robot's neural system that implements the role of embodiment. The body-hub SOM has four inputs, each being the angle of a single joint. In the experiments detailed here only 2 degrees from the head (up/down and left/right motors), and 2 degrees from the eyes (up/down and left/right motors) are used. Embodiment is operationalized here as the posture of eye and head position when the robot has to look to the left and to the right of the scene.

During each experiment, the connection weight linking the color map and the auditory map to the body-hub map are adjusted in real time using a Hebbian learning rule. These Hebbian associative connections are only modified from the current active body posture node. As the maps are linked together in real time, strong connections between objects typically encountered in particular spatial locations, and hence in similar body postures, build up.

To replicate the four experimental conditions of *Smith* and *Samuelson* [37.61], 20 different robots were used in each condition, with new random weights for the SOM and Hebbian connections. Results from the four conditions show a very high match between the robot's data and the child experiment results, closely replicating the variations in the four conditions. For example, in the Default Condition 83% of the trials resulted in the robots selecting the spatially linked objects, whilst in the Switch condition, where the space/object association was weakened, the robots' choices were practically due to chance at 55%. *Smith* and *Samuelson* [37.61] reported 71% of children selected the spatially linked object, versus 45% in the Switch condition.

This model demonstrates that it is possible to build an embodied cognitive system that develops linguistic and sensorimotor capabilities through interaction with the world, closely resembling the embodiment strategies observed in children's acquisition early word learning. Other cognitive robotics models have also been developed which exploit the principle of embodiment in robots' language learning, as in models of compositionality in action and language [37.65–68], in models of the cultural evolution of construction grammar [37.69, 70], and the modeling of the grounding of abstract words [37.71].

37.4.3 Number and Space

Number concepts have long been considered as prototypical examples of abstract and amodal concepts because their acquisition would require generalizing across a large range of instances to discover the invariant cardinality meaning of words such as *two* and *four* [37.72]. Mental arithmetic would therefore appropriately be modeled as abstract symbol manipulation, such as incrementing a counter or retrieving factual knowledge [37.73]. But evidence for an inescapable reference back from abstract number concepts to the sensori-motor experiences during concept acquisition has been present for a long time. Specifically, *Moyer* and *Landauer* [37.74] showed that the speed of deciding which of two visually presented digits represents the larger number depends on their numerical distance, with faster decisions for larger distances. Thus, even in the presence of abstract symbols we seem to refer to analog representations, as if comparing sensory impressions of small and large object compilations.

More recent studies provided further evidence that sensory–motor experiences have a strong impact on the availability of number knowledge. This embodiment signature can be documented by measuring the speed of classifying single digits as odd or even with lateralized response buttons. The typical finding is that small numbers (1, 2) are classified faster with left responses and large numbers (8, 9) are classified faster with right responses [37.76]. This spatial–numerical association response codes, or SNARCs effect, has been replicated across several tasks and extended to other effectors (for review [37.77]), including even attention shifts to the left or right side induced by small or large numbers, respectively [37.78].

Importantly, SNARC depends on one's sensory–motor experiences, such as directional scanning and finger counting habits, as well as current task demands. For example, the initial acquisition of number concepts in childhood occurs almost universally through finger counting and this learning process leaves a residue in the number knowledge of adults. Those who start counting on their left hand, thereby associating small numbers with left space, have a stronger SNARC than those who start counting on their right hand [37.79]. Similarly, reading direction modulates the strength of SNARC. In the original report by *Dehaene* et al. [37.76], it was noticed that adults from a right-to-left reading culture presented with weaker or even reversed SNARC. The notion of a spill-over of directional reading habits into the domain of num-

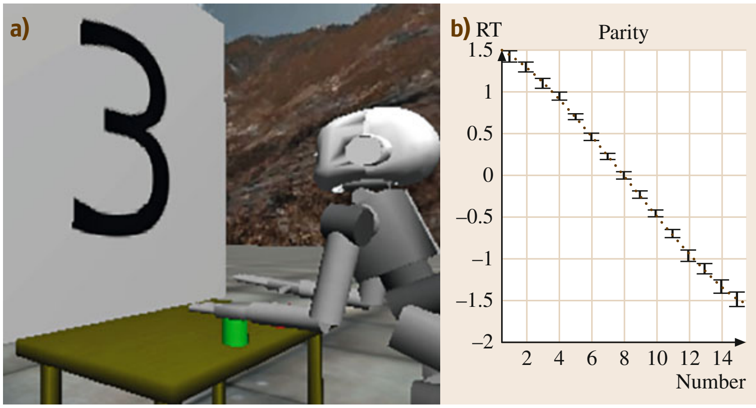

Fig. 37.2a,b **(a)** iCub simulation model of the SNARC (spatial–numerical association response code) effect; **(b)** SNARC effect results, with the difference in reaction times (right minus left hand) is plotted against number magnitude (after [37.75])

ber knowledge was further supported by developmental studies showing that it takes around 3 years of schooling before the SNARC emerges [37.80]. However, more recent work has found SNARC even in preschoolers (for review [37.81], thus lending credibility to the role of embodied practices such as finger counting in the formation of SNARC.

In a recent series of experiments with Russian–Hebrew bilinguals, *Shaki* et al. [37.82–84] (for review [37.85]) documented that both one's habitual reading direction and the most recent, task-specific scanning direction determine the strength of one's SNARC. These findings make clear that SNARC is a compound effect where embodied and situated (task-specific) factors add different weights to the overall SNARC.

SNARC and other biases extend into more complex numerical tasks such as mental arithmetic. For example, the association of larger numbers with right space is also present during addition (the operational momentum or OM effect). Regardless of whether symbolic digits or nonsymbolic dot patterns are added together, participants tend to over-estimate the sum, and this bias also influences spatial behavior [37.86]. More generally, intelligent behavior such as mental arithmetic seems to reflect component processes (distance effect, SNARC effect, OM effect) that are grounded in sensorimotor experiences.

The strong link between spatial cognition and number knowledge permits the modeling of the embodiment processes in the acquisition of number in robots. This has been the case with the recent developmental model developed by *Rucinski* et al. [37.75, 87] to model the SNARC effect and the contribution of pointing gestures in number acquisition. In the first study [37.75], a simulation model of the iCub is used. The robot is first trained to develop a body schema of the upper body through motor babbling of its arms. The iCub is subsequently trained to learn to recognize numbers by associating quantities of objects with numerical symbols as *1* and *2*. In the SNARC test case, the robot has to perform a psychological-like experiment and press a left or right button to make judgments on number comparison and parity judgment (Fig. 37.2b).

The robot's cognitive architecture is based on a modular neural network controller with two main components, following inspiration from a connectionist model of numerical cognition [37.88] and the TRoPICALS cognitive architecture of *Caligiore* et al. [37.89, 90]. The two main components of the neural control system are: (i) *ventral* pathway network, responsible for processing of the identity of objects as well as task-dependent decision making and language processing; and (ii) *dorsal* pathway network, involved in processing of spatial information about locations and shapes of objects and processing for the robot's action.

The *ventral* pathway is modeled, following *Chen* and *Verguts* [37.88], with a symbolic input which encodes the alphanumerical number symbols of numbers from 1 to 15, a mental number line encoding the number meaning (quantity), a decision layer for the number comparison and parity judgment tasks, and a response layer, with two neurons for left/right hand response selection. The *dorsal* pathway is composed of a number of SOMs which code for spatial locations of objects in the robot peripersonal space. One map is associated with gaze direction, and two maps respectively for each of the robot's left and right arms. The input to the gaze map arrives from the 3-dimensional proprioceptive vector representing the robot gaze direction (azimuth, elevation and vergence). The input to each arm position map consists of a 7-dimensional proprioceptive vector representing the position of the relevant arm joints. This dorsal pathway constitutes the core component of the model where the embodied properties of the model are directly implemented as the robot's own sensorimotor maps.

To model the developmental learning processes involved in number knowledge acquisition, a series of training phases are implemented. For the embodiment part, the robot is first trained to perform a process equivalent to motor babbling, to develop the gaze and arm space maps. With motor babbling the robot builds its internal visual and motor space representations (SOMs) by performing random reaching movements to touch a toy in its peripersonal space, whilst following its hand's position. Transformations between the visual spatial map for gaze and the maps of reachable left and right spaces are implemented as connections between the maps, which are learned using the classical Hebbian rule. At each trial of motor babbling, gaze and appropriate arm are directed toward the same point and resulting co-activations in already developed spatial maps is used to establish links between them.

The next developmental training establishes the links between number words (modeled as activations in the ventral input layer) and the number meaning (activations in the mental number line hidden layer). Subsequently the robot is taught to count. This stage models the cultural biases that result in the internal association of *small* numbers with the left side of space and *large* numbers with the right side. As an example of these biases, we considered a tendency of children to count objects from left to right, which is related to the fact that European culture is characterized by left-to-right reading direction [37.91]. In order to model the process of learning to count, the robot was exposed to an appropriate sequence of number words (fed to the ventral input layer of the model network), while at the same time the robot's gaze was directed toward a specific location in space (via the input to the gaze visual map). These spatial locations were generated in such a way that their horizontal coordinates correlated with number magnitude (small numbers presented on the left, large numbers on the right) with a certain amount of Gaussian noise. During this stage, Hebbian learning established links between number word and stimuli location in the visual field.

Finally, the model is trained to perform number reasoning tasks, such as number comparison and parity judgment, which corresponds to establishing appropriate links between the mental number line hidden layer and neurons in the decision layer. Specifically, one experiment focuses on the modeling of the SNARC effect. The robot's reaction time (i. e., amount of activity needed to exceed a response threshold in one of the two response nodes) in parity judgment and number comparison tasks were recorded to calculate the difference between right hand and left hand RTs for the same number. When difference values are plotted against number magnitudes the SNARC effect manifests itself in a negative slope as in Fig. 37.2. As the connections between visual and motor maps form a gradient from left to right, the links to the left arm map become weaker, while those to the right become stronger. Thus, when a small number is presented, internal connections lead to stronger automatic activation of the representations linked with the left arm than that of the right arm, thus causing the SNARC effect.

This model of space and number knowledge was also extended to include a more active interaction with the environment during the number learning process. This is linked to the fact that gestures such as pointing at the object being counted, or the actual touching of the objects enumerated, has been show to improve the overall counting performance in children [37.92]. In the subsequent model by *Rucinski* et al. [37.87], a simpler neural control architecture was used based on the Elman recurrent network to allow sequential number counting and the representation of gestures as proprioceptive states for the pointing gestures. The robot has to learn to produce a sequence of number words (from 1 to 10) with the length of the sequence equivalent to the number of objects present in the scene. Visual input to the model is a one-dimensional saliency map, which can be considered a simple model of a retina. In input, the additional proprioceptive signal was obtained from a pointing gesture performed by the iCub humanoid robot and is used to implement the gestural input to the model in the pointing condition. The output nodes encode the phonetic representation of the 10 numbers.

During the experiment, the robot is first trained to recite a sequence of number words. Then, in order to assess the impact of the proprioceptive information connected with the pointing gesture, the training is divided into two conditions: (i) training to count the number of objects shown to the visual input in the absence of the proprioceptive gesture signal, and (ii) counting though pointing, via the activation of the gesture proprioceptive input. Results show that such a simple recurrent architecture benefits from the input of the proprioceptive gesturing signal, with improved counting accuracy. In particular, the model reproduces the quantitative effects of gestures on the counted set size, replicating child psychology data reported in [37.92].

Overall, such a developmental robotics model clearly shows that the modeling of embodiment phenomena, such as the use of spatial representation in number judgments, and of the pointing gestures for number learning, can allow us to understand the acquisition of abstract concepts in humans as well as artificial agents and robots. This further demonstrates the benefit of the embodied intelligence approach to model a range of behavioral and cognitive phenomena from simple sensory–motor tasks to higher order linguistic and abstract cognition tasks.

37.5 Conclusion

This chapter has provided an overview of the three key principles of embodied intelligence, namely morphological computation, sensory–motor coordination, and EC, and of the experimental approaches and models from evolutionary robotics and developmental robotics. The wide range of behavioral and cognitive capabilities modeled through evolutionary and developmental experiments (e.g., locomotion in different environments, navigation and object discrimination, posture in early word learning and space and number integration) demonstrates the impact of embodied intelligence in the design of a variety of perceptual, motor and cognitive skills, including the potential to model the embodied basis of abstract knowledge as in numerical cognition.

The current progress of both evolutionary and developmental models of embodied intelligence, although showing significant scientific and technological advances in the design of embodied and situated agents, still has a series of open challenges and issues. These issues are informing ongoing work in the various fields of embodied intelligence.

One open challenge in morphological computation concerns how best to automatically design the body plans of robots so that they can best exploit this phenomenon. In parallel to this, much work remains to be done to understand what advantages morphological computation confers on a robot. For one, it is likely that a robot with a simpler control policy will be more robust to unanticipated situations: for example the jamming gripper is able to grasp multiple objects with the same control strategy; a rigid hand requires different control strategies for different objects. Secondly, a robot that performs more morphological computation may be more easily transferred from the simulation in which it was evolved to a physical machine: with a simpler control policy there is less that can go wrong when experiencing the different sensor signals and motor feedback generated by operation in the physical world.

Evolving robots provides a unique opportunity for developing rigorous methods for measuring whether and how much morphological computation a robot performs. For instance, if evolutionary algorithms can be designed that produce robots with similar abilities yet different levels of control and morphological complexity, and it is found that in most cases reduced control complexity implies greater morphological complexity, this would provide evidence for the evolution of morphological computation.

The emerging field of soft robotics [37.93] provides much opportunity for exploring the various aspects of morphological computation because the space of all possible soft robot body plans – with large variations in shape and continuous admixtures of hard and soft materials – is much larger than the space of rigid linkages traditionally employed in *classical* robots.

The design issue, i. e., the question of how systems able to exploit coordinated action and perception processes can be designed, represents an open challenge for sensory–motor coordination as well. As illustrated above, adaptive techniques in which the fine-grained characteristics that determine how agents react to current and previous sensory states are varied randomly and in which variations are retained or discarded on the basis of their effects at the level of the overall behavior exhibited by the agent/s interacting with their environment constitutes an effective method. However, this method might not scale up well with the number of parameters to be adapted. The question of how sensory–motor coordination capabilities can be acquired through the use of other learning techniques that relays on shorter term feedbacks represents an open issue. An interesting research direction, in that respect, consists in the hypothesis that the development of sensory–motor coordination can be induced through the use of task independent criteria such as information theoretic measures [37.94, 95].

Other important research directions concerns the theoretical elaboration of the different roles that morphological computation and sensory–motor coordination can play and the clarification of the relationship between processes occurring as a result of the agent/environmental interactions and processes occurring inside the agents' nervous systems

In developmental robotics models of EC the issues of open-ended, cumulative learning and of the scaling up of the sensory–motor and cognitive repertoires still requires significant efforts and novel methodological and theoretical approaches. Another issue, which combines both evolutionary and developmental approaches, is the interaction of phylogenetic and ontogenetic phenomena in the body/environment/brain adaptation.

Human development is characterized by cumulative, open-ended learning. This refers to the fact that learning and development do not start and stop at specific stages, but rather this is a life-long learning experience. Moreover, the skills acquired in various developmental stages are accumulated and integrated to support the further acquisition of more complex capabilities. One consequence of cumulative, open-ended learning is cognitive bootstrapping. For example in language development, the phenomenon of the vocabulary spurt exist, in which the knowledge and experience from the slow learning of the first 50–100 words causes a redefinition of the word learning strategy, and to syntactic bootstrapping, where children rely on syntactic cues and word context in verb learning to determine the meaning of new verbs [37.96]. Although some computational intelligence models of the vocabulary spurt exist [37.97], robotic experiments on language learning have been restricted to smaller lexicons, not reaching the critical threshold to allow extensive modeling of the bootstrapping of the agent's lexicon and grammar knowledge. These current limitations are also linked to the general issue of the scaling up of the robot's motor and cognitive capabilities and of cross-modal learning. Most of the current cognitive robotics models typically focus on the separate acquisition of only one task or modality (perception, or phonetics, or semantics etc.), often with limited repertoires rarely reaching 10 or slightly more learned actions or words. Thus a truly online, cross-modal, cumulative, open-ended developmental robotics model remains a fundamental challenge to the field.

Another key challenge for future research is the modeling of the interaction of the different timescales of adaptation in embodied intelligence, that is between phylogenetic (evolutionary) factors and ontogenetic (development, learning, maturation) phenomena. For example, maturation refers to changes in the anatomy and physiology of both the child's brain and the body, especially during the first years of life. Maturational phenomena related to the brain include the decrease of brain plasticity during early development, whilst maturation in the body is more evident due to the significant morphological growth changes a child goes through from birth to adulthood (see *Thelen* and *Smith*'s analysis of crawling and walking [37.98]). The ontogenetic changes due to maturation and learning have important implications for the interaction of development with phylogenetic changes due to evolution. Body morphology and brain plasticity variations can be in fact explained as evolutionary adaptations of the species to changing environmental context as with heterochronic changes [37.99]. For example, *Elman* et al. [37.43] discuss how genetic and heterochronic mechanisms provide an alternative explanation of the nature/nurture debate, where genetic phenomena produce architectural constraints of the organism's brain and body, which subsequently control and affects the results of learning interaction. Following this, *Cangelosi* [37.100] has tested the effects of heterochronic changes in the evolution of neural network architectures for simulated robotic agents.

The interaction between ontogenetic and phylogenetic factors has been investigated through evolutionary robotics models. For example, *Hinton* and *Nolan* [37.101] and *Nolfi* et al. [37.102] have developed evolutionary computational models explaining the effects of learning in evolution. The modeling of the evolution of varying body and brain morphologies in response to phylogenetic and ontogenetic requirements is also the goal of the *evo-devo* field of computational intelligence [37.7, 103–105]. These evolutionary/ontogenetic interaction models have, however, mostly focused on simple sensory–motor tasks such as navigation and foraging. Future work combining evolutionary and developmental robotics models can better provide theoretical and technological understanding of the contribution of different adaptation time scales and mechanisms in embodied intelligence.

References

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

37.2 R.A. Brooks: Elephants don't play chess, Robot. Auton. Syst. **6**(1), 3–15 (1990)

37.3 A. Cangelosi: Grounding language in action and perception: From cognitive agents to humanoid robots, Phys. Life Rev. **7**(2), 139–151 (2010)

37.4 H.J. Chiel, R.D. Beer: The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment, Trends Neurosci. **20**, 553–557 (1997)

37.5 F. Keijzer: *Representation and Behavior* (MIT Press, London 2001)

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

37.7 R. Pfeifer, J.C. Bongard: *How the Body Shapes the Way We Think: A New View of Intelligence* (MIT Press, Cambridge 2006)

37.8 C. Paul: Morphology and computation, Proc. Int. Conf. Simul. Adapt. Behav. (2004) pp. 33–38

37.9 C. Paul: Morphological computation: A basis for the analysis of morphology and control requirements, Robot. Auton. Syst. **54**(8), 619–630 (2006)

37.10 R. Pfeifer, F. Iida: Morphological computation: Connecting body, brain and environment, Jpn. Sci. Mon. **58**(2), 48–54 (2005)

37.11 G. Pezzulo, L.W. Barsalou, A. Cangelosi, M.H. Fischer, K. McRae, M.J. Spivey: The mechanics of embodiment: A dialog on embodiment and computational modeling, Front. Psychol. **2**(5), 1–21 (2011)

37.12 D. Pecher, R.A. Zwaan (Eds.): *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking* (Cambridge Univ. Press, Cambridge 2005)

37.13 M. Wilson: Six views of embodied cognition, Psychon. Bull. Rev. **9**, 625–636 (2002)

37.14 L. Meteyard, S.R. Cuadrado, B. Bahrami, G. Vigliocco: Coming of age: A review of embodiment and the neuroscience of semantics, Cortex **48**(7), 788–804 (2012)

37.15 R. Shepard, J. Metzler: Mental rotation of three dimensional objects, Science **171**(972), 701–703 (1972)

37.16 K. Dijkstra, M.P. Kaschak, R.A. Zwaan: Body posture facilitiates the retrieval of autobiographical memories, Cognition **102**, 139–149 (2007)

37.17 L.E. Williams, J.A. Bargh: Keeping one's distance: The influence of spatial distance cues on affect and evaluation, Psychol. Sci. **19**, 302–308 (2008)

37.18 A. Cangelosi, M. Schlesinger: *Developmental Robotics: From Babies to Robots* (MIT Press, Cambridge 2012)

37.19 J. Bongard: The utility of evolving simulated robot morphology increases with task complexity for object manipulation, Artif. Life **16**(3), 201–223 (2010)

37.20 E. Tuci, G. Massera, S. Nolfi: Active categorical perception of object shapes in a simulated anthropomorphic robotic arm, IEEE Trans. Evol. Comput. **14**(6), 885–899 (2010)

37.21 R. Pfeifer, G. Gomez: Morphological computation – Connecting brain, body, and environment, Lect. Notes Comput. Sci. **5436**, 66–83 (2009)

37.22 V. Pavlov, A. Timofeyev: Construction and stabilization of programmed movements of a mobile robot-manipulator, Eng. Cybern. **14**(6), 70–79 (1976)

37.23 S. Hirose, Y. Umetani: The development of soft gripper for the versatile robot hand, Mech. Mach. Theor. **13**(3), 351–359 (1978)

37.24 E. Brown, N. Rodenberg, J. Amend, A. Mozeika, E. Steltz, M.R. Zakin, H. Lipson, H.M. Jaeger: Universal robotic gripper based on the jamming of granular material, Proc. Natl. Acad. Sci. USA **107**(44), 18809–18814 (2010)

37.25 T.J. Allen, R.D. Quinn, R.J. Bachmann, R.E. Ritzmann: Abstracted biological principles applied with reduced actuation improve mobility of legged vehicles, Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. 2 (2003) pp. 1370–1375

37.26 M. Wisse, G. Feliksdal, J. Van Frankkenhuyzen, B. Moyer: Passive-based walking robot, IEEE Robot. Autom. Mag. **14**(2), 52–62 (2007)

37.27 K. Sims: Evolving 3d morphology and behavior by competition, Artif. Life **1**(4), 353–372 (1994)

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

37.29 GECCO 2012 Robot Videos: https://www.youtube.com/playlist?list=PLD5943A95ABC2C0B3

37.30 J.E. Auerbach, J.C. Bongard: On the relationship between environmental and mechanical complexity in evolved robots, Proc. 13th Int. Conf. Simul. Synth. Living Syst. (2012) pp. 309–316

37.31 F. Mondada, E. Franzi, P. Ienne: Mobile robot miniaturisation: A tool for investigation in control algorithms, Proc. 3rd Int. Symp. Exp. Robot. (Kyoto, Japan 1993)

37.32 S. Nolfi: Power and limits of reactive agents, Neurocomputing **49**, 119–145 (2002)

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

37.34 S. Nolfi: Categories formation in self-organizing embodied agents. In: *Handbook of Categorization in Cognitive Science*, ed. by H. Cohen, C. Lefebvre (Elsevier, Amsterdam 2005) pp. 869–889
37.35 J.J. Gibson: *The Perception of the Visual World* (Houghton Mifflin, Boston 1950)
37.36 N. Franceschini, F. Ruffier, J. Serres, S. Viollet: Optic flow based visual guidance: From flying insects to miniature aerial vehicles. In: *Aerial Vehicles*, ed. by T.M. Lam (InTech, Rijeka 2009)
37.37 S. Nolfi, D. Marocco: Active perception: A sensorimotor account of object categorization. In: *From Animals to Animats 7*, (MIT Press, Cambridge 2002) pp. 266–271
37.38 D. Floreano, T. Kato, D. Marocco, S. Sauser: Coevolution of active vision and feature selection, Biol. Cybern. **90**(3), 218–228 (2004)
37.39 H.A. Ruff: Infants' manipulative exploration of objects: Effect of age and object characteristics, Dev. Psychol. **20**, 9–20 (1984)
37.40 R. Bajcsy: Active perception, Proc. IEEE **76**(8), 996–1005 (1988)
37.41 D.H. Ballard: Animate vision, Artif. Intell. **48**, 57–86 (1991)
37.42 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 2010)
37.43 J.L. Elman, E.A. Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi, K. Plunkett: *Rethinking Innateness: A Connectionist Perspective on Development* (MIT Press, Cambridge 1996)
37.44 S. Nolfi, D. Floreano: Learning and evolution, auton, Robots **7**(1), 89–113 (1999)
37.45 N. Bernstein: *The Coordination and Regulation of Movements* (Pergamon, Oxford 1967)
37.46 P. Savastano, S. Nolfi: Incremental learning in a 14 DOF simulated iCub robot: Modelling infant reach/grasp development, Lect. Notes Comput. Sci. **7375**, 369–370 (2012)
37.47 A.M. Collins, M.R. Quillian: Retrieval time from semantic memory, J. Verb. Learn. Verb. Behav. **8**, 240–247 (1969)
37.48 E. Rosch: Cognitive representations of semantic categories, J. Exp. Psychol. Gen. **104**, 192–233 (1975)
37.49 D.E. Rumelhart, J.L. McClelland, P.D.P. Group: *Parallel Distributed Processing: Explorations in the microstructure of Cognition* (MIT Press, Cambridge 1986)
37.50 S. Harnad: The symbol grounding problem, Physica D **42**, 335–346 (1990)
37.51 L.W. Barsalou: Grounded cognition, Annu. Rev. Psychol. **59**, 617–645 (2008)
37.52 M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, C. Yoshida: Cognitive developmental robotics: A survey, IEEE Trans. Auton. Mental Dev. **1**, 12–34 (2009)
37.53 M. Lungarella, G. Metta, R. Pfeifer, G. Sandini: Developmental robotics: A survey, Connect Sci. **15**(4), 151–190 (2003)
37.54 P.Y. Oudeyer: Developmental robotics. In: *Encyclopedia of the Sciences of Learning*, Springer References Series, ed. by N.M. Seel (Springer, New York 2012) p. 329
37.55 A. Glenberg, K. Kaschak: Grounding language in action, Psychon. Bull. Rev. **9**(3), 558–565 (2002)
37.56 M.H. Fischer, R.A. Zwaan: Embodied language – A review of the role of the motor system in language comprehension, Q. J. Exp. Psychol. **61**(6), 825–850 (2008)
37.57 F. Pulvermüller: *The Neuroscience of Language* (Cambridge Univ. Press, Cambridge 2003)
37.58 S.F. Cappa, D. Perani: The neural correlates of noun and verb processing, J. Neurolinguist. **16**(2/3), 183–189 (2003)
37.59 O. Hauk, I. Johnsrude, F. Pulvermüller: Somatotopic representation of action words in human motor and premotor cortex, Neuron **41**(2), 301–330 (2004)
37.60 M. Tomasello: *Constructing a Language* (Harvard Univ. Press, Cambridge 2003)
37.61 L.B. Smith, L. Samuelson: Objects in space and mind: From reaching to words. In: *Thinking Through Space: Spatial Foundations of Language and Cognition*, ed. by K. Mix, L.B. Smith, M. Gasser (Oxford Univ. Press, Oxford 2010)
37.62 A.F. Morse, T. Belpaeme, A. Cangelosi, L.B. Smith: Thinking with your body: Modelling spatial biases in categorization using a real humanoid robot, 2010 Annu. Meet. Cogn. Sci. Soc. (2010) pp. 33–38
37.63 G. Metta, L. Natale, F. Nori, G. Sandini, D. Vernon, L. Fadiga, C. von Hofsten, J. Santos-Victor, A. Bernardino, L. Montesano: The iCub humanoid robot: An open-systems platform for research in cognitive development, Neural Netw. **23**, 1125–1134 (2010)
37.64 A.F. Morse, J. de Greeff, T. Belpaeme, A. Cangelosi: Epigenetic robotics architecture (ERA), IEEE Trans. Auton. Mental Dev. **2**(4), 325–339 (2010)
37.65 Y. Sugita, J. Tani: Learning semantic combinatoriality from the interaction between linguistic and behavioral processes, Adapt. Behav. **13**(1), 33–52 (2005)
37.66 V. Tikhanoff, A. Cangelosi, G. Metta: Language understanding in humanoid robots: iCub simulation experiments, IEEE Trans. Auton. Mental Dev. **3**(1), 17–29 (2011)
37.67 E. Tuci, T. Ferrauto, A. Zeschel, G. Massera, S. Nolfi: An experiment on behavior generalization and the emergence of linguistic compositionality in evolving robots, IEEE Trans. Auton. Mental Dev. **3**(2), 176–189 (2011)

37.68 Y. Yamashita, J. Tani: Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment, PLoS Comput. Biol. **4**(11), e1000220 (2008)

37.69 L. Steels: Modeling the cultural evolution of language, Phys. Life Rev. **8**(4), 339–356 (2011)

37.70 L. Steels: *Experiments in Cultural Language Evolution*, *Advances in Interaction Studies*, Vol. 3 (John Benjamins, Amsterdam 2012)

37.71 F. Stramandinoli, D. Marocco, A. Cangelosi: The grounding of higher order concepts in action and language: A cognitive robotics model, Neural Netw. **32**, 165–173 (2012)

37.72 J. Piaget: *The Origins of Intelligence in Children* (International Univ. Press, New York 1952)

37.73 G.J. Groen, J.M. Parkman: A chronometric analysis of simple addition, Psychol. Rev. **79**(4), 329–343 (1972)

37.74 R.S. Moyer, T.K. Landauer: Time required for judgements of numerical inequality, Nature **215**, 1519–1520 (1967)

37.75 M. Rucinski, A. Cangelosi, T. Belpaeme: An embodied developmental robotic model of interactions between numbers and space, Expanding the Space of Cognitive Science, 23rd Annu. Meet. Cogn. Sci. Soc., ed. by L. Carlson, C. Hoelscher, T.F. Shipley (Cognitive Science Society, Austin 2011) pp. 237–242

37.76 S. Dehaene, S. Bossini, P. Giraux: The mental representation of parity and number magnitude, J. Exp. Psychol. Gen. **122**, 371–396 (1993)

37.77 G. Wood, H.C. Nuerk, K. Willmes, M.H. Fischer: On the cognitive link between space and number: A meta-analysis of the SNARC effect, Psychol. Sci. Q. **50**(4), 489–525 (2008)

37.78 M.H. Fischer, A.D. Castel, M.D. Dodd, J. Pratt: Perceiving numbers causes spatial shifts of attention, Nat. Neurosci. **6**(6), 555–556 (2003)

37.79 D.B. Berch, E.J. Foley, R. Hill, R.P. McDonough: Extracting parity and magnitude from Arabic numerals: Developmental changes in number processing and mental representation, J. Exp. Child Psychol. **74**, 286–308 (1999)

37.80 M.H. Fischer: Finger counting habits modulate spatial-numerical associations, Cortex **44**, 386–392 (2008)

37.81 S.M. Göbel, S. Shaki, M.H. Fischer: The cultural number line: A review of cultural and linguistic influences on the development of number processing, J. Cross-Cult. Psychol. **42**, 543–565 (2011)

37.82 S. Shaki, M.H. Fischer: Reading space into numbers – A cross-linguistic comparison of the SNARC effect, Cognition **108**, 590–599 (2008)

37.83 S. Shaki, M.H. Fischer, W.M. Petrusic: Reading habits for both words and numbers contribute to the SNARC effect, Psychon. Bull. Rev. **16**(2), 328–331 (2009)

37.84 M.H. Fischer, R. Mills, S. Shaki: How to cook a SNARC: Number placement in text rapidly changes spatial-numerical associations, Brain Cogn. **72**, 333–336 (2010)

37.85 M.H. Fischer, P. Brugger: When digits help digits: Spatial-numerical associations point to finger counting as prime example of embodied cognition, Front. Psychol. **2**, 260 (2011)

37.86 M. Pinhas, M.H. Fischer: Mental movements without magnitude? A study of spatial biases in symbolic arithmetic, Cognition **109**, 408–415 (2008)

37.87 M. Rucinski, A. Cangelosi, T. Belpaeme: Robotic model of the contribution of gesture to learning to count, Proc. IEEE ICDL-EpiRob Conf. Dev. (2012)

37.88 Q. Chen, T. Verguts: Beyond the mental number line: A neural network model of number-space interactions, Cogn. Psychol. **60**(3), 218–240 (2010)

37.89 D. Caligiore, A.M. Borghi, D. Parisi, G. Baldassarre: TRoPICALS: A computational embodied neuroscience model of compatibility effects, Psychol. Rev. **117**, 1188–1228 (2010)

37.90 D. Caligiore, A.M. Borghi, R. Ellis, A. Cangelosi, G. Baldassarre: How affordances associated with a distractor object can cause compatibility effects: A study with the computational model TRoPICALS, Psychol. Res. **77**(1), 7–19 (2013)

37.91 O. Lindemann, A. Alipour, M.H. Fischer: Finger counting habits in Middle-Eastern and Western individuals: An online survey, J. Cross-Cult. Psychol. **42**, 566–578 (2011)

37.92 M.W. Alibali, A.A. DiRusso: The function of gesture in learning to count: More than keeping track, Cogn. Dev. **14**(1), 37–56 (1999)

37.93 R. Pfeifer, M. Lungarella, F. Iida: The challenges ahead for bio-inspired 'soft' robotics, Commun. ACM **55**(11), 76–87 (2012)

37.94 M. Lungarella, O. Sporns: Information self-structuring: Key principle for learning and development, Proc. 4th Int. Conf. Dev. Learn. (2005)

37.95 P. Capdepuy, D. Polani, C. Nehaniv: Maximization of potential information flow as a universal utility for collective behaviour, Proc. 2007 IEEE Symp. Artif. Life (CI-ALife 2007) (2007) pp. 207–213

37.96 L. Gleitman: The structural sources of verb meanings, Lang. Acquis. **1**, 135–176 (1990)

37.97 J. Mayor, K. Plunkett: Vocabulary explosion: Are infants full of Zipf?, Proc. 32nd Annu. Meet. Cogn. Sci. Soc., ed. by S. Ohlsson, R. Catrambone (Cognitive Science Society, Austin 2010)

37.98 E. Thelen, L.B. Smith: *A Dynamic Systems Approach to the Development of Cognition and Action* (MIT Press, Cambridge 1994)

37.99 M.L. McKinney, K.J. McNamara: *Heterochrony, the Evolution of Ontogeny* (Plenum, New York 1991)

37.100 A. Cangelosi: Heterochrony and adaptation in developing neural networks, Proc. GECCO99 Genet.

Evol. Comput. Conf., ed. by W. Banzhaf (Morgan Kaufmann, San Francisco 1999) pp. 1241–1248
37.101 G.C. Hinton, S.J. Nowlan: How learning can guide evolution, Complex Syst. **1**, 495–502 (1987)
37.102 S. Nolfi, J.L. Elman, D. Parisi: Learning and evolution in neural networks, Adapt. Behav. **3**(1), 5–28 (1994)
37.103 J. Bongard: Morphological change in machines accelerates the evolution of robust behavior, Proc. Natl. Acad. Sci. USA **108**(4), 1234–1239 (2011)
37.104 S. Kumar, P. Bentley (Eds.): *On Growth, Form, and Computers* (Academic, London 2003)
37.105 K.O. Stanley, R. Miikkulainen: A taxonomy for artifcial embryogeny, Artif. Life **9**, 93–130 (2003)