

6 Why Morphology Matters

Josh Bongard

One can distinguish between traditional and evolutionary robotics (ER) by the way in which each community generates controllers: traditional roboticists hand-design or use learning methods to create control policies, while evolutionary roboticists employ evolutionary algorithms. What further distinguishes these two approaches is that evolutionary algorithms may also be used to optimize robot morphology as well as the control policy. This chapter traces the history of this practice and outlines how we as a community are transitioning from questions regarding *how* to evolve morphology to *why* one should do so. Here I outline seven such reasons: selecting or evolving an appropriate morphology can (1) simplify control, (2) make seemingly difficult tasks easier, (3) increase evolvability, (4) provide new behaviors, (5) facilitate the extraction of information from the environment, (6) generate new research questions, and (7) improve scalability.

6.1 Introduction

Embodied cognition was an intellectual rebellion that entered the field of artificial intelligence (AI) in the 1980s (Brooks 1991a, 1991b), and challenged the prevailing (and still majority) view in AI that intelligence can be replicated in a computer without requiring interaction with the external world (Minsky 1974; Bechtel 1990). Brooks outlined a strong view in which internal modeling was not required for realizing relatively sophisticated behavior such as locomotion over uneven terrain (Brooks 1986) and, later, social interaction (Brooks et al. 1999). Indeed internal processing was minimized as much as possible; instead, emphasis was placed on exploiting the situated and embodied nature of the robot.

This movement has steadily been gaining ground since that time (Clark 1996; Pfeifer 1999; Pfeifer and Bongard 2006), but a large majority of the AI and even robotics community continue to focus extensively on the control side of behavior realization, with minimum attention paid to the robot's morphology. One reason for this emphasis may be that AI and robotics grew out of cybernetics, which in turn was

founded on control theory. Furthermore, the formal, mathematical traditions of the field promote approaches that provide guarantees of convergence or the discovery of optimal solutions, or both. Such guarantees are very difficult or impossible in stochastic optimization processes such as evolutionary algorithms, much less when the optimization process is extended to the morphology of the robot. It may be that this contributes to why evolutionary robotics is not a popular approach among many mainstream robotics and AI practitioners.

Nature however provides abundant examples of organisms with diverse body plans, and an ever-greater diversity of adaptive mechanisms by which those body plans support behavior. Biorobotics (Webb and Consi 2001), a sister field of evolutionary robotics, has demonstrated several successes of replicating both the control and morphological adaptations of individual animals in machine form such that the machines also demonstrate one or more of the animal's behaviors. However, bioroboticists tend to copy the actual body plan of animals, rather than copy the evolutionary mechanisms that produced that body plan in the first place, as is done in evolutionary robotics.

If the right morphology can indeed facilitate behavior, the question then arises as to how to select an appropriate machine body plan for the task environment and desired behavior. A prevailing view in robotics is that choosing such a body plan is much more intuitive than designing controllers: "Humans are much better at designing physical systems than they are at designing intelligent control systems: complex powered machines have been in existence for over 150 years, whereas it is safe to say that no truly intelligent autonomous machine has ever been built by a human" (Nelson, Barlow, and Doitsidis 2009, 22). However, a growing number of examples surveyed in this chapter illustrate that selecting an appropriate body plan is rarely an intuitive process, and that therefore automating the selection process using evolutionary computation may indeed allow for the realization of increasingly intelligent and autonomous machines.

6.1.1 The Role of Morphology in Animal Behavior

As already mentioned, many of the geometric layouts, material properties, and mechanical mechanisms of animals' body plans are shaped by evolution to support particular behaviors. A well-studied example in robotics is human bipedal locomotion. Human legs are structured to support extremely energy-efficient travel by exploiting the passive swing of the leg and loading of the ankle, and therefore support long-distance travel. Indeed several researchers have been able to reproduce this energy-efficient gait in robots through careful hand-tuning of the robot's morphology (McGeer 1990; Collins et al. 2005). However, for many morphological adaptations it is difficult to determine which behaviors they evolved to support. Whether the quadrupedality seen in higher animals or the hexapody seen in insects is a result of historical accident

or serves adaptive behaviors is unknown, and there exist a panoply of hypotheses for why bipedality evolved in early humans (Lovejoy 1980; Morgan 1982; Jablonski and Chaplin 1993; Hunt 1996).

For this reason it is difficult to determine which aspects of an organism's morphology to replicate in a robot, as some aspects may be dictated by physiological constraint or are the result of historical accident. An example of this former constraint is that it seems likely that bipedalism was a simpler path for evolution to take to free up the upper extremities rather than evolving a new pair of limbs, as the quadrupedal body plan is an extremely conserved trait across the animal kingdom. Therefore, although instantiating bipedalism in robots may be desirable so that they can operate in a world built for human body plans, if not done correctly, a legged robot may not be capable of the behavior that legs originally evolved for: energy-efficient transport. Similarly, whether to implement four or five fingers on a robot hand is probably a less important design decision than whether to include an opposable digit, which supports a wider range of grasping strategies than a hand without such a digit (Wilson 1998).

As these two examples illustrate, choosing an appropriate body plan for a robot is not so easy a task as it seems. An alternative strategy therefore is to replicate the evolutionary mechanism that produced the body plan originally: biological evolution. Artificial evolution may then discover a body plan that, while possibly not similar to any found in nature, is well suited to the task environment, desired behavior, and control policy evolved along with it.

6.1.2 "How" versus "Why"

Karl Sims was the first to demonstrate that evolutionary computation could be used to evolve both the morphology and control policies for autonomous agents in a physics-based simulator (Sims 1994). He obtained many agents with morphologies that were and were not biologically familiar. He employed an evolutionary algorithm with a genotype-to-phenotype mapping that included recursion and therefore tended to produce body plans with repeated segments.

This work stimulated a subsequent wave of research that explored different ways of evolving morphology. One research line employed direct encoding schemes (Ventrella 1994; Lipson and Pollack 2000), while another line expanded the recursive Sims-type mapping by incorporating Lindenmayer systems (Prusinkiewicz and Lindenmayer 1990) into the encoding scheme (Hornby and Pollack 2001). Others still pursued more biologically realistic encodings that simulated genetic regulatory networks to grow the body plan (Eggenberger 1997), and later both the body plan and control policy (Bongard 2002) of simulated agents. Examples of robots evolved with this latter approach are shown in figure 6.1q–u).

Yet most of this work focused on *how* to evolve morphology, rather than *why* one should do so: it is imperative for the field of evolutionary robotics to accumulate such




Figure 6.1

A selection of the author's past projects involving embodied cognition. (a–f) Simulated machines with differing morphologies, yet identical sensor, motor, and control systems were evolved to isolate the effect of morphology on evolvability (Bongard and Pfeifer 2002). A physical autonomous machine (g) was reported in Bongard, Zykov, and Lipson 2006 that is capable of creating a simulation of its own morphology (h). The simulation of a brachiating machine (i) was used to prototype passive dynamic behaviors before instantiating them in a physical machine (j) (Frutiger, Bongard, and Iida et al. 2002). The simulation of a machine with both serial and parallel actuated linkages (k) was used to prototype nonintuitive control strategies that were later used on the physical version of the machine (l) (Zykov, Bongard, and Lipson 2004). A simulation framework for prototyping stochastic self-assembling machine modules was developed, and used to discover assembly plans for producing three-dimensional structures such as the tower shown in (m).

Results from the simulation were used to build a macroscale physical prototype of this technology (n) (White et al. 2005). (o) Both mass distribution and control parameters of a simulated bipedal machine were evolved to produce locomotion in Bongard and Paul 2001. (p) Strategies by which multiple machines may share self-models and controllers were investigated in Bongard 2009. (q–u) Evolving the body plans and control circuits for simulated machines. (q,r): Evolution of block pushing behavior (Bongard and Pfeifer 2001). (s–u) Illustrations indicating the combined growth of the body plan during ontogenesis as well as the internal sensor, motor, and control circuits (internal networks). (v) A virtual robot capable of both legged locomotion and object manipulation (Auerbach and Bongard 2009b).

reasons in order to justify the added complexity of these methods. This chapter will explore seven reasons why a machine's body plan should be carefully hand- or automatically designed, beyond the one already mentioned:¹

1. Choosing or evolving an appropriate body plan can simplify control compared to when an inappropriate body plan is employed—*morphology simplifies control*;
2. Seemingly difficult tasks, approached with purely computational methods, become easier if the methods incorporate an appropriate morphology—*morphology eases tasks*;
3. Although incorporating morphological parameters into the evolutionary process increases the dimensionality of the search space, doing so can often improve the probability of finding useful behavior—*morphology increases evolvability*;
4. By exploring the space of robots in which the topology of the robot's body plan may change during behavior, new behaviors not available with a fixed morphology can be realized—*morphology affords new behaviors*;
5. With the right body, a robot can use it to systematically extract useful information from the environment—*morphology supports self-exploration*;
6. With the adoption of certain body plans, research questions not yet explored in artificial intelligence arise—*morphology creates new research questions*; and
7. What constitutes a “good” morphology becomes less intuitive as the complexity of the task increases—*morphology supports scalability*.

6.2 Morphology Simplifies Control

The task environment in which a robot must act dictates much of the control policy that the robot should adopt in order to perform its task successfully. However, the task environment also dictates what kinds of morphology are appropriate: a poorly chosen morphology will require a more complex control policy than a well-chosen morphology. As a simple example, consider a robot that should move forward over flat terrain. Clearly a wheeled robot will require simpler control than a legged robot: a wheeled robot may simply supply constant torque to all of its wheels while a legged robot must orchestrate the motion of its legs. However, for more complex task environments, it is not always so clear how to devise a morphology that will simplify control. Therefore, a growing number of examples have been put forward in the robotics and evolutionary robotics community. One such example, passive dynamic walking machines, has already been mentioned, in which little or no control is required if the body is designed properly. However, it is often the case that human intuition or examples from nature fail to suggest what an appropriate robot body plan should be for a given task. In such circumstances evolutionary algorithms can be used to discover an appropriate body plan for the task at hand.

Lichtensteiger and Eggenberger (1999) described a mobile robot that simulates the facets of an insect's eye, and demonstrated that evolving the distribution of these facets on the robot (which resulted in nonuniform distributions) simplified the visual processing required to calculate the time at which the robot would contact an external object.

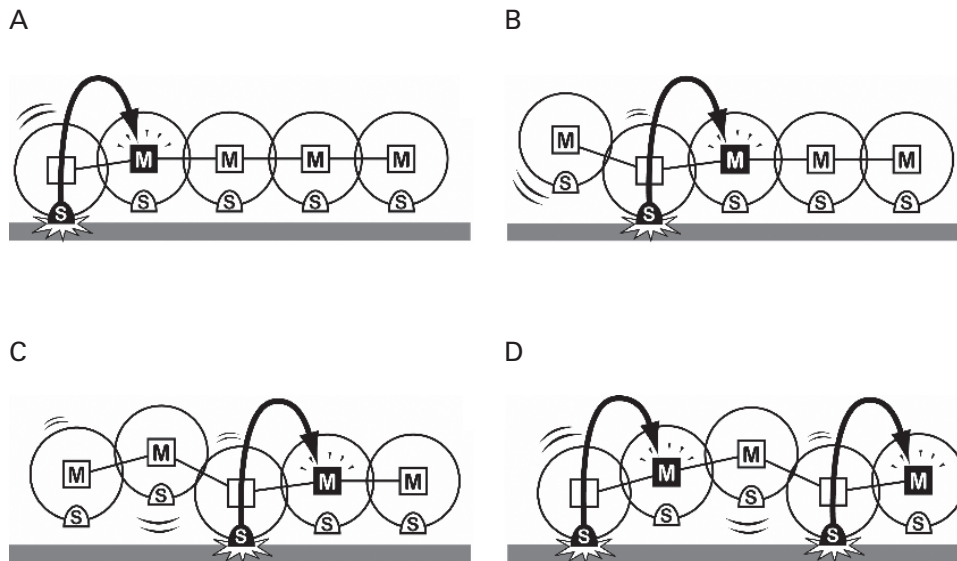
The author developed a method that combines evolution and ontogeny: virtual organisms are grown in the environment in which they must behave, and evolution in turn shapes the form of these growth programs (Bongard and Pfeifer 2001). This was accomplished by simulating genetic regulatory networks (Eggenberger 1997): genomes that are evolutionarily modified may contain noncoding and coding regions. These coding regions are treated as genes, and the parameters within these regions dictate the gene's behavior. Genes produce simulated substances that diffuse through the body of the robot as it grows in its environment. Some of these substances may affect the expression of other genes near the point of the substance's production or, through passive diffusion, the substance may affect gene expression in distal parts of the robot.

Other substances may cause phenotypic change. The robots are composed of spherical modules (figure 6.1q–u), and modules may grow and split in response to substance concentration. The modules also contain neurons and connecting synapses (visualized by the internal networks in figure 6.1s–u); substances may create, destroy, or move neurons throughout the robot's body, and may cause synapses to grow from one neuron, through the body, and attach to other neurons. Additional synapses may connect sensors to neurons, neurons to motors, or sensors directly to motors. Using this approach, few assumptions are made about the form of the robots' body plans and neural controllers, and evolution is free to discover an appropriate body plan and neural controller for the task at hand.

In one experiment in which robots were selected for approaching and pushing large objects in their environment (figure 6.1q,r), two robots were observed to have long front appendages and exhibited a wave-like form of locomotion reminiscent of that seen in inchworms. Both robots were evolutionarily related, and were found to possess very similar neural patterning as illustrated in figure 6.2.

A human engineer asked to design a controller for such a body plan would most likely favor a centralized architecture in which a central timer orchestrated a sequence of motions along the appendage's length. However, artificial evolution here discovered a simpler solution that does not require timing or orchestration, but rather exploits the interactions between the robot's body plan (a series of spherical modules) and the environment (gravity pulls lifted modules back to the ground) and controller (distributed direct sensor-motor reflexes).

This example points to one of the fundamental differences between more formal learning methods and evolutionary algorithms: the former are suited for parametric

**Figure 6.2**

The evolved morphology, control policy, and resultant action of the front appendages of the robots shown in figure 6.1q,r. The front appendage is comprised of a series of morphological modules (large circles) attached to each other by one degree-of-freedom rotational joints. Each joint is motorized (M). The evolved genetic regulatory network places a touch sensor (S) in each module, and connects it by a synapse (arrows) to the motorized joint in the neighboring module. Gravity and the mass distribution of the appendage causes the distal tip to contact the ground, which causes the touch sensor to fire and activate the motor it connects to. This causes the joint to rotate the tip off the ground, thereby causing the second module to contact the ground. This in turn activates the next touch sensor and its associated motor, and so on, producing a wave of motion along the appendage's length without requiring centralized control (Courtesy MIT Press).

optimization in which guarantees of convergence are required, while the latter allow both structural (i.e., topological) and parametric improvement in which such guarantees are more difficult to provide. Currently, popular learning methods support that there are a fixed number of parameters that must be optimized: these parameters may specify the parameters of a controller with a fixed architecture, or aspects of the robot's morphology. However, many evolutionary algorithms do not presuppose the dimensionality of candidate solutions. This is very much the case in indirect genotype-to-phenotype mappings, in which there is dissociation between the complexities of the genotype and phenotype. Indeed the more complex robot shown in figure 6.1r is a descendent of the simpler robot in figure 6.1q, although both were grown from a genotype with the same number of genes (Bongard and Pfeifer 2001).

6.3 Morphology Eases Tasks

One of the surprising results accumulating in the embodied cognition literature is that seemingly challenging tasks when tackled with disembodied algorithms are rendered much simpler when an intelligent agent is allowed to interact with the environment using its body. A particularly striking example comes from the work of Metta and Fitzpatrick (2003), in which a humanoid robot is tasked with object segmentation. This task has primarily been cast as a purely perceptual problem in the computer vision literature in which objects in a cluttered scene should be visually separated (“segmented”) from the background and other objects.

Metta and Fitzpatrick describe a series of experiments in which a humanoid robot moves its arm, and observes the resulting motion in its visual field. When the arm comes in contact with various objects, the robot can simultaneously observe and feel the result of that contact. Different outcomes of this interaction help not only to determine the outline of the object against the background, but also physical characteristics of it. For example, heavy objects halt the arm’s movement. Lighter objects may begin to move, and the resulting flow field indicates the boundary of the object. Complex objects may be more amenable to movement along different axes as a result of their geometry or friction properties, or both, which the robot can discover by poking an object from different directions. The robot thus demonstrates a form of active perception (Noë 2005), in which the object is understood not so much as a result of its external appearance, but rather by the affordances (Gibson 1977) that the object projects: in other words, the ways in which the embodied agent may interact with the object. Indeed active perception is a growing area of study in evolutionary robotics (Gomez and Eggenberger 2007; Tuci, Massera, and Nolfi 2009; Bongard 2009a, 2009b).

Metta and Fitzpatrick also point out that this blending of perceptual and motor processes is observed in the primate brain, and that this blending may therefore have a functional role rather than being an accident of evolution. Finally, they demonstrate that this interaction with the environment can support the development of more complex cognitive abilities such as mimicking a human demonstrator. The main long-term goal of artificial intelligence is, after all, the realization of more complex cognition through the aggregation of simpler competencies; this work, along with others (e.g., Yamashita and Tani 2008) indicates how embodied behavior may simplify this progression.

6.4 Morphology Increases Evolvability

The ability to find good solutions within a search space is dictated primarily by three factors: the dimensionality of the search space, its smoothness, and the degree of

neutrality. Smoothness can be operationally defined as how much phenotypic effect a small genetic mutation causes. Mutational effect can be influenced by the amount of genetic interaction, as in Kauffman's NK fitness landscapes (Kauffman and Johnsen 1991), where N represents the dimensionality of the genotype and K indicates the degree of coupling between genes. For higher K values, a slight genetic mutation will influence the expression of several other genes, magnifying the mutation's phenotypic effect. It has been known for some time in biology that the probability of a mutation conferring a fitness benefit is inversely proportional to the phenotypic magnitude of that mutation (Fisher 1930). Therefore, for evolutionary algorithms with high epistasis, the resulting fitness landscape tends to be "rugged": regions of high fitness are separated from one another by difficult-to-cross regions of low fitness.

The effect of neutrality on evolvability² has also been extensively studied in biology (Kimura 1983) and evolutionary computation (Barnett 1998; Yu and Miller 2002), and shows that populations may discover more high fit regions by undergoing a series of neutral mutations. Neutral mutation in a population can be visualized within the fitness landscape metaphor by envisioning a series of solutions clustering upward along a slope of increasing fitness, but then diffusing outward across a horizontal plateau, and thereby discovering multiple routes upward toward more fit solutions.

It follows from this that increasing the dimensionality of the search space for a given problem does not necessarily make it more difficult for an evolutionary algorithm to discover good solutions, as long as the correspondingly larger search space is smoothed in the process. In other words, a larger yet smoother fitness landscape may be more amenable to evolutionary search than a smaller, more rugged landscape.

Adding additional dimensions to a search may indeed smooth it in the process by introducing what the theoretical biologist Michael Conrad termed "extradimensional bypasses" (Chen and Conrad 1994; Cariani 2002): isolated fitness peaks in the small space may become saddle points in the larger space. This additional genotypic material may not necessarily introduce phenotypic novelty: rather, it may introduce intermediate, highly-fit solutions between two solutions that already existed in the smaller space such that an evolving population more often finds the higher peak by traversing the extradimensional bypass.

Of course, blindly increasing the dimensionality of a search space does not guarantee the creation of bypasses: larger search spaces might just as easily become more rugged than smooth. Seen through the lens of evolutionary robotics, we may take an existing ER experiment in which only the controller is optimized, increase the genotype to also specify some morphological parameters, and then re-evolve robots against the same task. We may expect one of three outcomes. First, the new experimental regime may turn out to be less evolvable than the original experiment, suggesting that the increased search space has either kept the ruggedness of the landscape constant or increased it. Second, the evolutionary algorithm may outperform the original

experiment because it discovers a different body plan than the fixed one used previously. Third, performance may increase but the final, best robot may have the same or very similar body plan to the fixed one from the original experiment.

This last outcome would suggest that the additional evolvable morphological parameters have introduced extradimensional bypasses into the search space: evolution gradually modified the initial, default body plan through a series of intermediate forms—as well as modifying the controllers accompanying those body plans—until eventually evolution converged back on a body plan similar to the original body plan.

This evolutionary trajectory has been observed in several independent studies that employed different evolutionary algorithms, robots, and tasks. In Bongard and Paul (2001) we first evolved a simulated biped with an active torso to walk over flat terrain (figure 6.1o) while keeping the morphology fixed. A standard genetic algorithm was used to evolve a feedforward neural network to optimize the robot's displacement over a fixed time period. In this first set of experiments the radii of the lower and upper legs and arms were kept constant. In the second set of experiments these three morphological parameters (lower leg, upper leg, and arm radii) were evolved along with the robot's neural network controller. In some of these latter experiments evolution discovered body plans different from the default, such as the one shown in figure 6.1o. It can be seen that the radii of the legs and arms are significantly larger than that of the torso, producing a bimodal mass distribution: mass is gathered in the legs to stabilize walking; and torque induced by walking is canceled by the neural network swinging the heavy arms in opposition to the legs. Other experiments produced robots that outperformed those found in the original set of experiments, but had very similar morphologies: in these runs the body plan was observed to change, but converged eventually on the original, default morphology. From these observations we concluded that although the additional morphological parameters had increased the dimensionality of the search space, it had also created new fitness peaks (as illustrated by the robot in figure 6.1o) as well as introduced extradimensional bypasses that allowed evolution to more often climb to fitness peaks that existed in the original experiments but which were more difficult to find.

More recently, I evolved a simulated robot manipulator based on the human arm and hand to grasp, lift, and actively distinguish between several objects. Each experiment series in which only the controller of the arm was evolved was paired with a second experiment series in which the radii and length of each finger's phalanges and the distribution of fingers around the hand were evolved (figure 6.3).

In the fixed, default morphology the fingers are arranged equidistantly around the spherical palm, and the lengths and radii of the phalanges were set to the same value. When the morphology was evolved, initial robots had random controllers but started with the default morphology. Evolution could then place fingers closer or further from each other around the palm, and phalange lengths and radii could differentiate. It was

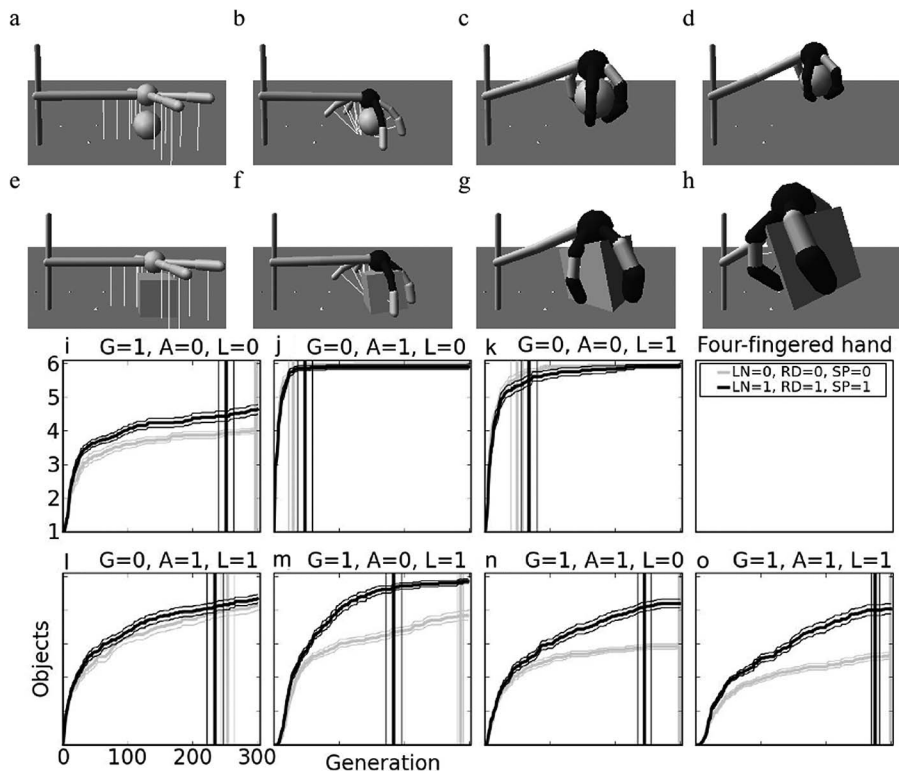


Figure 6.3

Evolving the morphology and controller of an anthropomorphic arm. (a–h) An evolved arm capable of grasping, lifting, and actively distinguishing between objects of different shapes. (a–d) The resultant behavior of the arm when grasping spheres. (e–h) The resultant behavior when the same arm grasps cubes. Note the trajectories of the arm in the two cases are different, indicating it can distinguish between them. The thin white lines indicate range sensors. Black and white finger segments indicate whether the tactile sensor in that segment is firing (black) or not (white). The evolved, differing radii of the phalanges can be seen most clear in (h). (i–o) Fitness improvements over evolutionary time for arms with fixed morphology (gray curves; 30 independent runs each) and evolved morphology (black curves; 30 independent runs each). In the former case phalange length (LN), phalange radii (RD), and spacing (SP) between phalanges was not evolved (LN,RD,SP = 0); in the latter case all three types of morphological parameters were evolved (LN,RD,SP = 1). Evolutionary time is measured in generations along the horizontal axis; the vertical axis indicates the mean number of generations required for the arms to successfully manipulate the object (two of the six objects are shown in [a–h]). The grey and black vertical lines indicate the mean time until the arms successfully evolved to manipulate all six objects. In different regimes, arms were evolved to grasp (G), actively distinguish between (A) and/or lift (L) the objects. Unity and zero indicate whether that aspect of manipulation was selected for or not, respectively. Thick lines indicate means; thin lines indicate one unit of standard error. For more details please refer to Bongard 2009b.

hoped that this would lead to the appearance of specialized digits such as human's opposable thumbs; indeed in several (but not all) runs differentiated digits were observed.

Again, I found that despite the expanded search space, due to the inclusion of the morphological parameters, evolution consistently found fitter behaviors compared to when hand morphology was held fixed. Also, in many runs the shape and distribution of the fingers would drift away from the default morphology and then gradually converge back to the default.

The creation of extradimensional bypasses through the evolution of morphology suggests the morphology should not only be optimized because a body plan different from that envisioned by human engineers may be found. Rather, it is often the *process* of gradually changing morphology that increases performance: there may be particular morphologies in which it is easier to discover a workable but crude controller, but that are not ideal for optimal performance. Subsequent evolution may refine the control and morphology (perhaps returning morphology close to the originally fixed body plan) so that robots may exhibit performance close to this optimal.

For example, it was observed in Bongard and Paul (2001) that early in some evolutionary runs mass was focused in the legs, making the morphology quite stable. This allowed for the discovery of a shuffling gait. Later evolution moved mass into the arms, thereby allowing for longer strides and using the heavy arms to cancel the torque induced by such a gait. This could be viewed as a kind of scaffolding (Wood, Bruner, and Ross 1976), as it is known in psychology, or shaping (Dorigo and Colombetti 1994), as it is known in the robotics literature: typically scaffolding or shaping implies that a learner's environment is structured to facilitate learning; here, the learner's own body provides a gradient for discovering crude behaviors and then refining them as the learner's body plan changes. I refer to this additional kind of scaffolding as *morphological scaffolding*.

6.5 Morphology Affords New Behaviors

In most robotics experiments, the robot's body plan is fixed by the experimenter. In several of the research projects surveyed so far in this chapter, body plans change over evolutionary timescales. In a particular class of robots—modular robots—there is the possibility of morphological change over the lifetime of the robot itself. A modular robot is typically made up of independent units, or modules that can attach to and detach from one another either under active control of the robot, or as response to environmental perturbation. This latter class of modular robots are known as stochastic modular robots (White et al. 2005); a simulated and physical example are shown in figure 6.1m,n. Behaviors exhibited by modular robots are typically denoted using a

self-X nomenclature: such systems may start as independent modules and gradually attach into some desired form (self-assembly [Murata, Kurokawa, and Kokaji 1994; White et al. 2005; Groß et al. 2006]); create a copy of themselves (self-reproduction [Zykov et al. 2005]); start in some desired form, experience module loss or malfunction, and reconfigure back into the original form (self-repair [Tomita et al. 1999; Stoy and Nagpal 2004]); or alter form to exhibit different behaviors (self-reconfiguration [Murata et al. 2002; Park et al. 2008]).

This particular approach to embodied cognition creates a rich new vista of behaviors beyond the reach of traditional robots, or for that matter beyond most biological organisms.³ Using the self-assembling machine and associated simulator shown in figure 6.1m,n, we explored different strategies for stochastic self-assembly (White et al. 2005). One of the main goals of modular robotics is to scale down the units to micro- or nanoscale. Such scales enforce strict constraints on the internal complexity of the units, so one approach is to assume units have no internal power source or actuation machinery: rather, the modules are moved by the flow of their surrounding motion. We further assumed that each side of each cube can be magnetized, and could attach to a powered floor plate if coming into contact with it. Once powered, an attached module can flip the polarity of the magnets on any of its exposed sides, thereby attracting unpowered modules suspended in the medium. If a module attaches to the structure it becomes powered and may in turn attract other modules, allowing for self-assembly.

By specifying a sequence of commands that flip magnet polarity during the self-assembly process, different structures can be created such as a vertical tower (figure 6.1m). However, for more complex structures attracting independent modules to particular surfaces may be difficult. For example if a cage-like structure is self-assembling it may be very improbable for independent modules to ricochet into the internal volume of the cage through random motion. However, by flipping the polarity between two attached modules, substructures can be jettisoned back into the fluid to be attached elsewhere. This allows for the added capability of self-reconfiguration, which can be exploited to overcome this limitation.

In White et al. (2005) we described two programs that specified the assembly of the same structure: in the first program modules were only attached; in the second program, modules were attached and detached during different stages of the assembly process. This latter program was more complex than the former, yet it led to more rapid mean assembly times. The reason for this was that attached modules were detached at strategic times and places. Once released back into the medium, they had low linear and rotational velocity, and were proximal to surfaces that would otherwise have had difficulty attracting an independent module. These surfaces would then often capture the recently released module, and the self-assembly process would continue.

This result demonstrates yet another advantage of building robot systems in which morphology may change: the robot's growing body may serve as its own scaffold, easing the otherwise difficult task of self-assembly. This is another example of the morphological scaffolding concept introduced in the previous section. When evolving robot body plans along with controllers, body plans change over evolutionary time and increase the probability of discovering a robot that can accomplish the desired task. In stochastic self-assembling modular robots, the robot can change the topology of its own body over the lifetime of the system such that it increases the probability of assembling the desired structure.

In another set of experiments from the same work, we investigated assembly times for structures composed of the same number of objects, but having different geometries: we studied structures ranging from a vertical tower (figure 6.1m) to branched and then closed structures. We found that despite the same volume of these structures, the closed structure was assembled more rapidly and consistently than the branched or tower structure. This indicates another importance of considering morphology for such systems: one should not only consider which body plan of a modular robot is appropriate for a given task once assembled, but also which body plan also has the best probability of self-assembling (or self-reconfiguring or self-repairing) in its environment.

Finally, Zykov et al. (2005) described a self-reproducing robot, in which the body plan of the parent robot played an important role in self-reproduction. The parent robot is composed of a series of magnetic modules, bonds to external modules, and deposits them on to the growing child robot, thereby dynamically changing the topology of its body in the process. The child robot's body plan also plays an important role in the self-reproduction process. It rotates itself to accept donated modules from the parent robot at the right place on its body, thereby contributing to a process in which the body plan of both parent and child robot change over time.

In the described experiments the self-assembly, self-reconfiguration, and self-reproduction programs were hand-designed. The finding that self-assembly can be deceptive (i.e., simpler programs do not always produce the desired structure more often or consistently than a more complex program) suggests that human intuition is difficult to apply to this problem, and that an automated process may therefore produce superior self-assembly programs. However, the number and type of commands as well as their timing are not known a priori, so a learning method in which a fixed number of parameters are optimized may not be appropriate. Rather, an evolutionary algorithm such as genetic programming (Koza 1992), in which the structure and parameters of algorithms can be optimized, may be appropriate in this domain in future. Indeed preliminary work (Estévez and Lipson 2007) has investigated evolving self-assembly programs in which the fitness of a program is the rapidity and consistency with which it leads to the desired structure.

6.6 Morphology Supports Self-Exploration

One of the necessary abilities an agent must possess to be considered intelligent is the capacity for learning. Learning has long played a role in AI and robotics, predating evolutionary robotics by several decades. Developmental robotics (Lungarella et al. 2003), another sister discipline of evolutionary robotics (in addition to biorobotics), is dedicated to investigating how a robot can exploit its body to scaffold its own learning, although typically the body does not change during this process. The experiments described in section 6.3 illustrate a typical experiment from that field.

Work within machine learning, a branch of disembodied AI, is often partitioned into passive learning, in which the learner passively receives the raw material for learning from a teacher, and active learning in which the learner must harvest its own material. Much work in that field has investigated strategies how to extract the most informative data for learning from the external world (Seung, Oppen, and Sompolinsky 1992; Baram, Yaniv, and Luz 2004).

When the learner has a physical body, the challenge becomes how to best exploit its body to extract useful information about itself and its environment. This question of course depends on the goal of the learner. In past work (Bongard Zykov, and Lipson 2006) we introduced a physical, autonomous robot that is able to autonomously generate legged gaits for itself and regenerate a new gait if it suffers unanticipated body plan change such as physical damage. In order to accomplish this, the robot evolves simulations that reflect the current topology of its body plan; if it is damaged, it re-evolves these simulations to reflect its changed state.

In most real-world contexts it is infeasible to equip a machine with enough sensors to detect all possible malfunctions. Therefore, it was assumed that the robot would use inferred relationships between motor commands and subsequent sensor signals to infer possible changes in its morphology. The robot therefore performs a short motor sequence (figure 6.4a), records the resulting sensor signature, and then generates a population of simulations, each of which contains a different body plan topology (Figure 6.4b). The robot actuates each simulation with the same motor program it just performed, and compares the resulting sensor data from the simulation against the physical sensor data: the closer the match, the more accurate the simulator must be. Using an internal evolutionary algorithm the robot can then optimize the models to better reflect the topology of its body plan.

However, a single motor program is unlikely to provide sufficient information for it to accurately assess the state of all of its body parts. Therefore the robot must perform several actions to extract sufficient information for self-modeling. But we cannot allow the robot to perform an arbitrary number of actions, as physical motion is expensive in terms of time, power, and the risk of suffering additional damage by performing inappropriate actions in poorly understood environments. Therefore the robot must




Figure 6.4

An example of an autonomous robot that integrates self-modeling and internal behavior generation. The robot executes an action (A) and then creates a set of models to describe the sensorial result of that action (B). It then uses the models to find a new action (C) that will reduce uncertainty in the models when executed. By alternating between modeling and testing, it eventually finds and then uses an accurate model to internally optimize a behavior (D) before executing it in reality (E,F). (Courtesy of *Science* magazine)

employ an active learning methodology that finds motor programs likely to extract the most information from the external environment. If successful, the robot can minimize the number of actions it has to perform.

The robot uses query by committee (Seung, Opper, and Sompolinsky 1992) as the underlying active learning method. A population of initially random self-models is evolved against a single motor program and the resulting sensor data. Model evolution is paused after a short period and a second evolutionary algorithm searches for a motor program that, when supplied to the current model population, causes them to output diverging time series sensor data (figure 6.4c). This informative motor program is executed by the physical robot, the resulting sensor data is recorded, and the models are re-evolved to explain both motor programs. This process continues for a fixed number of cycles or until a sufficiently accurate self-model is found. An accurate self-model is then used in the traditional evolutionary robotics manner: a controller is optimized aboard the simulated robot such it performs the desired task, after which the best controller is executed by the physical robot (figure 6.4d).

In this domain in which the body plan of the physical robot may change unexpectedly,⁴ morphology drives behavior at several levels: the overall learning framework uses the physical body as a vehicle for generating and extracting information from the world, and the simulated morphologies in the self-models influence which motor program will be executed. It was demonstrated in Bongard, Zykov, and Lipson 2006 that this method can be used to automatically create a simulator for a robot rather than hand-design one, and can allow the robot to continuously generate compensating behaviors in the face of unexpected damage, malfunction, or degradation.

This work also goes some way toward reconciling the seemingly divergent philosophies of traditional and embodied AI: while the former stressed purely computational notions such as modeling and planning, the latter demonstrated that robust behavior could be realized through model-free, sensor-to-motor coupling (Brooks 1991a, 1991b). Our robot exhibited several automatically generated gaits such as the one shown in figure 6.1e that result from tight sensor-motor coordination. However, these gaits are the result of a separate process that uses exploratory actions to create internal self-models, and self-models to drive exploratory actions. This suggests that the two branches of AI may not be as irreconcilable as was previously thought: it may be possible to gradually augment the cognitive capabilities of a robot by grounding them in low-level sensor-motor processes. Regardless, this reconciliation is only possible by placing the physical robot's body at the center of cognitive processes, rather than focusing on cognition while marginalizing or discarding the body.

In neuroscience, there is a similar reconciliation between abstract thought and sensor-motor coordination, which is often referred to as mental simulation or motor imagery (Porro et al. 1996; Hesslow 2002). Mental simulation implies that higher brains "run" candidate motor programs and reason about the results. This contrasts

with older, more disembodied theories in which the brain extracts concepts from their sensorimotor components, stores them in mental symbols, and higher-thought processes such as planning and reasoning transform those symbols. The role of symbols in AI and cognitive science raises difficult issues such as how abstract symbols can acquire real-world meaning (Harnad 1990). These issues somewhat dissipate when a closer connection between sensor-motor processes and cognitive processes is forged, such as in the robot just described. The robot's self-models have real-world meaning insofar as they arise from sensor-motor processes, and they allow the robot to determine how it may move. In this sense self-models *afford* (Gibson 1977) possible action: they tell the robot how it may interact with the world.

Although it is not yet known how biological brains simulate action, there is growing evidence that the mental toolkit for this exist. Body images encoded as topographic maps exist in the primate brain, and often exhibit topologies similar to the sensor and motor systems that project to them: the retina projects to the primary visual cortex; the organ of Corti to the primary auditory cortex; and cutaneous receptors to the primary somatosensory cortex (Kandel, Schwartz, and Jessell 2000). Also, both forward and inverse models are found in the cerebellum (Wolpert, Miall, and Kawato 1998; Wolpert and Kawato 1998): like the robot just described, the primate brain can predict future sensor state given a candidate motor primitive by supplying that primitive to a forward model, and inverse models can be used to find an appropriate motor primitive that will produce a desired future sensor state. Although it is unknown how higher brain areas like the cerebral cortex make use of these sensorimotor maps and models, it seems likely that there is no clear dividing line in the brain between sensorimotor representation and simulation and higher cognition.

6.7 Morphology Creates New Research Questions

As just described, how situated action supports the emergence of more complex (and possibly cognitive⁵) behaviors is becoming a hot topic in evolutionary robotics and AI in general. However, fixing a robot's morphology places limits and biases on the kinds of action that the robot can perform, and therefore also on the more complex behaviors that those actions may eventually support. A robot with legs can only exhibit legged locomotion; a wheeled robot with a rigid gripper can only move over flat terrain and grasp objects with a fixed radius. Therefore, given a desired task, a roboticist should select not only appropriate control architecture for the robot, but also an appropriate body plan.

The engineering tradition places an inordinate emphasis on modularity (Suh 1990): presumably, this is primarily due to the reductionist paradigm that emphasizes breaking problems down into separable subproblems, each of which can be more easily solved than the original problem. The general concept of modularity can be broken

down into functional and structural modularity: what are the separate functions that an agent (or its constitutive parts) must perform, and what are the substructures that make up an agent's phenotype? Broadly speaking, structure dictates function: the shape of a protein completely determines its function; the morphology of the primate hand determines its possible grasping strategies. However, it does not follow that a structural module must contribute to only one function: in addition to grasping, the primate hand contributes to a near infinitude of other functions from fine manipulation to whole-body brachiation.

Yet in robotics one tends to observe a one-to-one correspondence between morphological components and function: in a wheeled robot with a gripper, wheels contribute to movement but not grasping, while a gripper contributes to grasping but not locomotion. This seems on the surface to simplify things: the researcher can design a subcontroller for each function that interacts only with the substructure responsible for that function. In this spirit several projects have explicitly evolved modular neural networks in which each neural module contributes to a different function (Brooks 1986; Calabretta et al. 2000). However, this approach is not scalable: humans are capable of a very large number of behaviors, and it is unlikely that there are separate brain structures for each of them. What seems more likely is that similar actions are driven by common neural circuitry. Indeed in the field of evolutionary robotics several researchers have demonstrated that this explicit structural modularity in the robot's neural network is not always necessary (Bongard 2008; Izquierdo and Buhmann 2008; Auerbach and Bongard 2009b).

In recent work (Auerbach and Bongard 2009a) we have explored a simulated robot morphology (figure 6.5) that challenges the structural and functional modularity observed in other morphologies. The robot is composed of a series of segments, each of which supports a pair of cylindrical appendages. Each appendage pair can be actuated, as can intersegmental joints that allow the robot to flex its "spine" within the sagittal plane. This allows the robot to rotate its upper or lower body up- and downward, or to keep its body horizontal during locomotion. In this work the robot's morphology was not evolved; only its controller was. The fitness function selected for robots that locomote toward an object in their environment, grasp the object once reached, and lift it. The robot is equipped with sensors that signal the distance of the sensor from the object at each time step of the simulation: the sensors can be thought of as returning a signal commensurate with ambient sound amplitude, and the object as emitting a continuous sound. The robot's front appendages were also equipped with touch sensors.

Given the body plan and task environment of the robot, there are two classes of behaviors that the robot may adopt to complete the task. In the first class, the robot would raise the front of the body, locomote to the object using the back four appendages, and manipulate the object with the front pair when the object is reached. The




Figure 6.5

Two evolved behaviors for a segmented robot. (a–h) A sample functionally specialized controller. The robot initially raises its front pair of appendages and keeps them raised while it approaches the target object. Once reached, it grasps and lifts the object. (i–p) A functionally generalized controller. The front appendage pair participates in locomotion toward the object. Once reached, the front pair participates in object manipulation.

second strategy involves developing a hexapodal gait such that all six appendages are involved in locomotion toward the object; once reached, the front of the body would rise to allow the front appendage pair to manipulate the object. Both such strategies, if performed rapidly, achieve a very high fitness value. It was found that in different independent trials both strategies were discovered by evolution: figure 6.5a shows an example of the first class of behaviors; figure 6.5b shows an example of the second class.

It was also found that functionally specialized controllers tended to arise much more often than functionally generalized ones, even though both, when optimized, obtain about the same fitness values. This raises the question as to why and how specialization arises. We investigated three hypotheses: (1) functionally specialized controllers are more evolvable (i.e., the slopes of fitness peaks surrounding specialized controllers can be more easily climbed than peaks surrounding generalized

controllers); (2) early discovery of rudimentary specialized controllers is more common, and the degree of specialization is difficult for evolution to tune throughout a run so evolution is forced to refine these behaviors; or (3) functionally specialized controllers more easily allow for active perception (Noë 2005).

The first hypothesis seems unlikely due to the observation that when comparing runs that discover specialized controllers and those that discover functionally specialized controllers, both exhibit relatively rapid, smooth, and consistent fitness improvement over the run. This suggests that both kinds of solutions afford relatively smooth gradients for the population to climb. Many runs exhibited controller phylogenies that gradually changed from specialized to generalized and back again, providing evidence against the second hypothesis.

This leaves the third hypothesis, which seems to be supported by the results. When the front of the body is raised early during an evaluation, the touch sensors fall silent. This silence can be shaped by evolution into a trigger that guides the controller into a cyclic attractor, which produces a rhythmic gait toward the object. If the front appendages come into contact with the object then the touch sensors fire: again, evolution can alter the controller such that this signal shifts the controller into a different dynamic state that causes object grasping and lifting. This is a form of active perception in that through the robot's interaction with the environment—mediated by its morphology—the robot either approaches or manipulates the object.

By choosing this segmented body plan, new research questions arise, namely, how will the responsibility for different behaviors be distributed across the agent's controller *and* morphology? And how will evolution accomplish this, given the opportunity to do so? If we wish to realize robots capable of ever more complex tasks, these robots will have to perform several behaviors separated in time, space, and across their morphological and control substructures. It seems unlikely that each such substructure will support one and only one function, but how to overlap them is nonintuitive. Therefore it will become important to investigate these new research questions by considering different combinations of robot morphology and control. More desirable, however, is to allow evolution to co-optimize both to discover the most appropriate distribution of function across the robot's structure.

6.8 Morphology Supports Scalability

The preceding six lines of argument are unlikely to convince many in the robotics community that formulating an appropriate robot morphology for a given task is nonintuitive. In order to make progress along this front it is necessary to clearly demonstrate under what conditions this intuition breaks down. For this reason I conducted a series of experiments with an anthropomorphic arm (already briefly described in section 6.4 and illustrated in figure 6.3a–h) in which both the morphology and the

controller of the arm were evolved. The arm was exposed to a variety of selection pressures that favored the manipulation of a series of objects in some way. Three fitness functions were formulated that selected for object grasping, lifting, and active categorical perception. Grasping (G) is accomplished by closing the fingers tightly around objects, and lifting (L) is accomplished by maintaining contact with the object while raising the arm. Active categorical perception (A) is accomplished by causing the controller to exhibit common signal signatures among the hidden neurons when grasping objects of the same shape but different signatures when manipulating objects of different shape.

Different selection regimes could be created by evolving a population of robots using multiobjective optimization (Deb 2002), in which each of these competencies is an objective. This allowed for a total of seven regimes in which robots were evolved against only one objective ($G = 1$, $A = 1$, or $L = 1$; figure 6.3i–k), any two objectives ($G + A + L = 2$; figure 6.3l–n), or all three objectives ($G, A, L = 1$; figure 6.3o). Shaping (Dorigo and Colombetti 1994) was also employed in that, at the outset of an evolutionary run, robots were only evolved against a single object. When fitness rose to a prespecified threshold a second object was added to the training pool, and so on up to a total of six objects.

It was found, not surprisingly, that the more objectives a robot population had to satisfy simultaneously, the more difficult it was to evolve a robot that could successfully manipulate all six objects (note the lowered curves in figures 6.3l–o compared to figures 6.3j–k⁶). In addition, it was found that evolving aspects of the hand's morphology along with the controller, compared to just evolving the hand's morphology, did not provide any advantage when only one objective was selected for.⁷ However, for two or three objectives, evolving morphology along with control provided a significant performance advantage.

This result suggests that at least for object manipulation, an intuitive biologically inspired hand design—in which four fingers with more or less equal phalange length, radii, and equal separation between the digits—is perfectly serviceable if the task is relatively simple. However, when the robot must satisfy several goals simultaneously (such as grasping, lifting, and actively distinguishing between objects of different shape), slight changes in this default morphology provide dramatically increased capability. For example, in one independent trial in which all three manipulation capabilities were selected for, the robot shown in figure 6.3a–h was evolved. Note the slight differences in phalange radii within individual digits. This ensures that when the hand closes on an object, different phalanges come in contact with the object depending on the object's shape. This in turn causes different touch sensors to fire for differently shaped objects, stimulating different hidden neuron signatures and satisfying the active categorical perception requirement. Such combinations of a tight grip and active perception are more difficult if phalanges on the same digit have the same radii.

This provides evidence that as roboticists we may indeed be able to formulate perfectly serviceable morphologies for simple tasks, but as the tasks become more difficult, for example by requiring that the robot accomplish several tasks at once, our intuitions as designers break down. Even slight changes to the body plan may greatly simplify the task (as discussed in section 6.3 or allow for simpler controller policies (section 6.2, although which changes provide these advantages may not be immediately clear: they may however be found through evolutionary search.

6.9 Conclusions

This chapter has surveyed the history and current state of an intellectual movement in artificial intelligence centered on the concept of embodiment: in order to exhibit intelligent behavior, the agent must have a body with which to interact with the environment. This movement can be summarized by considering the progression of research questions that have been asked: *what*, *how*, and *why*.

The founders of embodied AI (Brooks 1991a; Clark 1996; Pfeifer 1999; Pfeifer and Bongard 2006) described *what* embodiment is, and what some of the consequences of having a body are. This raised the issue of *how* to choose an appropriate body plan for a given task. Evolutionary robotics provided one answer to this question by showing how evolutionary algorithms could be extended to optimize both the morphology and control policies of autonomous robots (Prusinkiewicz and Lindenmayer 1990; Sims 1994; Ventrella 1994; Eggenberger 1997; Lipson and Pollack 2000; Hornby and Pollack 2001; Bongard 2002). If evolutionary robotics is to be considered a serious engineering or computational discipline, or both, however, we must provide concrete evidence of increased performance to justify the cost of this additional algorithmic complexity. In other words, we must provide reasons for *why* evolutionary algorithms are employed for morphology optimization.

This chapter has described seven such reasons. A robot body plan appropriate to the task at hand can simplify control and make seemingly difficult tasks easier, but it is often difficult to determine manually what morphology would provide these benefits. Allowing evolution to optimize morphology instead requires additional evolvable parameters and therefore increases the dimensionality of the search space, but if done properly this can actually smooth the enlarged search space and increase the probability of finding good solutions. The burgeoning field of modular robotics assumes that the topology of the robot's body plan can change over its lifetime, which generates new behaviors beyond traditional robotics such as self-assembly, self-repair, self-reconfiguration, and even self-reproduction at several size scales. A curious robot with the right morphology can learn intelligently about its own capabilities and limitations, as well as discover unforeseen changes to its body as a result of damage. The adoption of nonstandard morphologies such as segmented body plans raises new

research questions about how to distribute the responsibility for several behaviors across different parts of the robot body. Finally, as we attempt to evolve robots capable of ever more complex behaviors, our intuitions about what form the robot's body must take break down, and we must rely on artificial evolution to explore the rich space of physical forms that support such behaviors—just as biological evolution did.

Notes

1. It is often difficult to decide which aspects of biological body plans to mimic.
2. Evolvability can be loosely defined as the ability of an evolutionary system to consistently discover higher-fit phenotypes.
3. The notable exception is certain slime molds that transition between single- and multicellular configurations.
4. Rather than deliberately, as in the modular robotics examples.
5. “Cognition” is a controversial term: here I use it as shorthand for sufficiently complex behaviors that may cause an outside observer to consider the robot to be, to some degree, cognitive. Of course, for different observers this threshold will be different, rendering “cognition” a subjective term.
6. The exception to this was grasping (figure 6.3i), which benefited when lifting or active perception, or both, was also selected for.
7. Again, with the exception of grasping, in which evolving morphology provided a slight benefit (figure 6.3i).

References

- Auerbach, J., and J. C. Bongard. 2009a. Evolution of functional specialization in a morphologically homogeneous robot. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 89–96. Montreal, Canada.
- Auerbach, J., and J. C. Bongard. 2009b. How robot morphology and training order affect the learning of multiple behaviors. In *Proceedings of the IEEE Congress on Evolutionary Computation*, 39–46. Trondheim, Norway.
- Baram, Y., R. E. Yaniv, and K. Luz. 2004. Online choice of active learning algorithms. *Journal of Machine Learning Research* 5:255–291.
- Barnett, L. 1998. Ruggedness and neutrality: The NKp family of fitness landscapes. In *Proceedings of the Sixth International Conference on Artificial Life*, ed. C. Adami et al., 18–27.
- Bechtel, W. 1990. Connectionism and the philosophy of mind. In *Mind and Cognition: A Reader*, ed. W. G. Lycan, 252–273. Oxford: Basil Blackwell.

- Bongard, J. C. 2002. Evolving modular genetic regulatory networks. In *Proceedings of the 2002 Congress on Evolutionary Computation*, 17–21.
- Bongard, J. C. 2008. Behavior chaining: Incremental behavioral integration for evolutionary robotics. In *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, ed. S. Bullock et al., 64–71.
- Bongard, J. C. 2009a. Accelerating self-modeling in cooperative robot teams. *IEEE Transactions on Evolutionary Computation* 13 (2): 321–332.
- Bongard, J. C. 2009b. The utility of evolving robot morphology increases with task complexity for object manipulation. *Artificial Life* 16 (3): 201–223. doi: 10.1162/artl.2010.
- Bongard, J. C., and C. Paul. 2001. Making evolution an offer it can't refuse: Morphology and the extradimensional bypass. In *Proceedings of the Sixth European Conference on Artificial Life*, ed. J. Keleman and P. Sosik, 401–412. Prague.
- Bongard, J. C., and R. Pfeifer. 2001. Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In *Proceedings of the Genetic and Evolutionary Computation Conference*, ed. L. Spector et al., 829–836.
- Bongard, J. C., and R. Pfeifer. 2002. A method for isolating morphological effects on evolved behaviour. In *Proceedings of the Seventh International Conference on the Simulation of Adaptive Behaviour*, ed. B. Hallam et al., 305–311. Cambridge, MA: MIT Press.
- Bongard, J. C., V. Zykov, and H. Lipson. 2006. Resilient machines through continuous self-modeling. *Science* 314:1118–1121.
- Brooks, R. 1986. A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation* 2 (1): 14–23.
- Brooks, R. A. 1991a. Intelligence without reason. *Artificial Intelligence: Critical Concepts* 3: 569–595.
- Brooks, R. A. 1991b. Intelligence without representation. *Artificial Intelligence* 47:139–159.
- Brooks, R. A., C. Breazeal, M. Marjanovic, B. Scassellati, and M. M. Williamson. 1999. The Cog project: Building a humanoid robot. *Lecture Notes in Computer Science* 1562:52–87.
- Calabretta, R., S. Nolfi, D. Parisi, and G. P. Wagner. 2000. Duplication of modules facilitates the evolution of functional specialization. *Artificial Life* 6 (1): 69–84.
- Cariani, P. A. 2002. Extradimensional bypass. *Bio Systems* 64 (1–3): 47–53.
- Chen, J. C., and M. Conrad. 1994. A multilevel neuromolecular architecture that uses the extradimensional bypass principle to facilitate evolutionary learning. *Physica D. Nonlinear Phenomena* 75:417–437.
- Clark, A. 1996. *Being There*. Cambridge, MA: MIT Press.
- Collins, S., A. Ruina, R. Tedrake, and M. Wisse. 2005. Efficient bipedal robots based on passive-dynamic walkers. *Science* 307 (5712): 1082–1085.

- Deb, K. 2002. *Multi-Objective Optimization Using Evolutionary Algorithms*. New York: Wiley.
- Dorigo, M., and M. Colombetti. 1994. Robot shaping: Developing situated agents through learning. *Artificial Intelligence* 70 (2): 321–370.
- Eggenberger, P. 1997. Evolving morphologies of simulated 3d organisms based on differential gene expression. In *Proceedings of the Fourth European Conference on Artificial Life*, 205–213.
- Estévez, N., and H. Lipson. 2007. Dynamical blueprints: Exploiting levels of system-environment interaction. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 238–244.
- Fisher, R. A. 1930. *The Genetical Theory of Natural Selection*. Oxford: Clarendon Press.
- Frutiger, D. R., J. C. Bongard, and F. Iida. 2002. Iterative product engineering: Evolutionary robot design. In *Proceedings of the Fifth International Conference on Climbing and Walking Robots*, ed. P. Bidaud and F. B. Amar, 619–629. London: Professional Engineering Publishing.
- Gibson, J. J. 1977. The theory of affordances. In *Perceiving, Acting and Knowing: Toward an Ecological Psychology*, ed. R. Shaw and J. Bransford, 67–82.
- Gomez, G., and P. Eggenberger. 2007. Evolutionary synthesis of grasping through self-exploratory movements of a robotic hand. In *Proceedings of the IEEE Congress on Evolutionary Computation*, ed. A. Tay, 3418–3425.
- Groß, R., M. Bonani, F. Mondada, and M. Dorigo. 2006. Autonomous self-assembly in swarm-bots. *IEEE Transactions on Robotics* 22 (6): 1115–1130.
- Harnad, S. 1990. The symbol grounding problem. *Physica D. Nonlinear Phenomena* 42:335–346.
- Hesslow, G. 2002. Conscious thought as simulation of behavior and perception. *Trends in Cognitive Sciences* 6:242–247.
- Hornby, G. S., and J. B. Pollack. 2001. Body-brain coevolution using L-systems as a generative encoding. In *Proceedings of the Genetic and Evolutionary Computation Conference*, 868–875.
- Hunt, K. D. 1996. The postural feeding hypothesis: An ecological model for the evolution of bipedalism. *South African Journal of Science* 92 (2): 77–90.
- Izquierdo, E., and T. Buhrmann. 2008. Analysis of a dynamical recurrent neural network evolved for two qualitatively different tasks: Walking and chemotaxis. In *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, ed. S. Bullock et al., 257–264.
- Jablonski, N. G., and G. Chaplin. 1993. Origin of habitual terrestrial bipedalism in the ancestor of the hominidae. *Journal of Human Evolution* 24:259–280.
- Kandel, E. R., J. H. Schwartz, and T. M. Jessell. 2000. *Principles of Neural Science*. 4th ed. New York: McGraw-Hill.
- Kauffman, S. A., and S. Johnsen. 1991. Coevolution to the edge of chaos: Coupled fitness landscapes, poised states, and coevolutionary avalanches. *Journal of Theoretical Biology* 149 (4): 467–505.

- Kimura, M. 1983. *The Neutral Theory of Molecular Evolution*. Cambridge: Cambridge University Press.
- Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press.
- Lichtensteiger, L., and P. Eggenberger. 1999. Evolving the morphology of a compound eye on a robot. In *Proceedings of the Third Workshop on Advanced Mobile Robots*, 127–134.
- Lipson, H., and J. B. Pollack. 2000. Automatic design and manufacture of robotic lifeforms. *Nature* 406 (6799): 974–978.
- Lovejoy, C. O. 1980. Hominid origins: The role of bipedalism. *American Journal of Physical Anthropology* 52:250.
- Lungarella, M., G. Metta, R. Pfeifer, and G. Sandini. 2003. Developmental robotics: A survey. *Connection Science* 15 (4): 151–190.
- McGeer, T. 1990. Passive dynamic walking. *International Journal of Robotics Research* 9 (2):62–82.
- Metta, G., and P. Fitzpatrick. 2003. Better vision through manipulation. *Adaptive Behavior* 11 (2): 109–128.
- Minsky, M. 1974. *A Framework for Representing Knowledge*. Cambridge, MA: MIT Press.
- Morgan, E. 1982. *The Aquatic Ape: A Theory of Human Evolution*. Souvenir Press.
- Murata, S., H. Kurokawa, and S. Kokaji. 1994. Self-assembling machine. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 441–448.
- Murata, S., E. Yoshida, A. Kamimura, H. Kurokawa, K. Tomita, and S. Kokaji. 2002. M-TRAN: Self-reconfigurable modular robotic system. *IEEE/ASME Transactions on Mechatronics* 7 (4): 431–441.
- Nelson, A. L., G. J. Barlow, and L. Doitsidis. 2009. Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems* 57 (4):345–370.
- Noë, A. 2005. *Action in Perception*. Cambridge, MA: MIT Press.
- Park, M., S. Chitta, A. Teichman, and M. Yim. 2008. Automatic configuration methods in modular robots. *International Journal of Robotics Research* 27 (3–4): 403–421.
- Pfeifer, R. 1999. *Understanding Intelligence*. Cambridge, MA: MIT Press.
- Pfeifer, R., and J. Bongard. 2006. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: MIT Press.
- Porro, C. A., M. P. Francescato, V. Cettolo, M. E. Diamond, P. Baraldi, C. Zuiani, M. Bazzocchi, and P. E. di Prampero. 1996. Primary motor and sensory cortex activation during motor performance and motor imagery: A functional magnetic resonance imaging study. *Journal of Neuroscience* 16 (23): 7688–7698.
- Prusinkiewicz, P., and A. Lindenmayer. 1990. *The Algorithmic Beauty of Plants*. New York: Springer Verlag.

- Seung, H. S., M. Opper, and H. Sompolinsky. 1992. Query by committee. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, 287–294.
- Sims, K. 1994. Evolving virtual creatures. In *Proceedings of SIGGRAPH*, 15–22. Orlando, FL.
- Stoy, K., and R. Nagpal. 2004. Self-repair through scale independent self-reconfiguration. In *Proceedings of the IEEE/RSJ International Conference on Robotics and Systems*, 2062–2067.
- Suh, N. P. 1990. *The Principles of Design*. Oxford: Oxford University Press.
- Tomita, K., S. Murata, H. Kurokawa, E. Yoshida, and S. Kokaji. 1999. Self-assembly and self-repair method for a distributed mechanical system. *IEEE Transactions on Robotics and Automation* 15 (6): 1035–1045.
- Tuci, E., G. Massera, and S. Nolfi. 2009. Active categorical perception in an evolved anthropomorphic robotic arm. In *Proceedings of the IEEE Congress on Evolutionary Computation*, ed. A. Tyrell et al., 31–38.
- Ventrella, J. 1994. Explorations in the emergence of morphology and locomotion behavior in animated characters. In *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, 436–441. Cambridge, MA: MIT Press.
- Webb, B., and T. R. Consi. 2001. *Biorobotics: Methods and Applications*. Cambridge, MA: MIT Press.
- White, P., V. Zykov, J. Bongard, and H. Lipson. 2005. Three-dimensional stochastic reconfiguration of modular robots. In *Proceedings of Robotics: Science and Systems*, 161–168. Cambridge, MA: MIT Press.
- Wilson, F. R. 1998. *The Hand*. New York: Random House.
- Wolpert, D. M., and M. Kawato. 1998. Multiple paired forward and inverse models for motor control. *Neural Networks* 11 (7–8): 1317–1329.
- Wolpert, D. M., R. C. Miall, and M. Kawato. 1998. Internal models in the cerebellum. *Trends in Cognitive Sciences* 2 (9): 338–347.
- Wood, D., J. S. Bruner, and G. Ross. 1976. The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry, and Allied Disciplines* 17 (2): 89–100.
- Yamashita, Y., and J. Tani. 2008. Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment. *PLoS Computational Biology* 4 (11): e1000220.
- Yu, T., and J. Miller. 2002. Finding needles in haystacks is not hard with neutrality. *Lecture Notes in Computer Science*. Berlin: Springer.
- Zykov, V., J. C. Bongard, and H. Lipson. 2004. Evolving dynamic gaits on a physical robot. In *Late Breaking Papers for the 2004 Genetic and Evolutionary Computation Conference*, Seattle, WA.
- Zykov, V., E. Mytilinaios, B. Adams, and H. Lipson. 2005. Self-reproducing machines. *Nature* 435 (7039): 163–164.