Chapter 2
The ‘What’, ‘How’ and the ‘Why’ of Evolutionary Robotics

Josh Bongard

Abstract. The field of embodied artificial intelligence is maturing, and as such has progressed from *what* questions (“what is embodiment?”) to *how* questions: how should the body plan of an autonomous robot be designed to maximize the chance that it will exhibit a desired set of behaviors. In order to stand on its own however, rather than a reaction to classical AI, the field of embodied AI must address *why* questions as well: why should body and brain both be considered when creating intelligent machines? This report provides three new lines of evidence for why the body plays an important role in cognition: (1) an autonomous robot must be able to adapt behavior in the face of drastic, unanticipated change to its body; (2) under-explored body plans raise new research questions related to cognition; and (3) optimizing body plans accelerates the automated design of intelligent machines, compared to leaving them fixed.

2.1 The *What* of Embodiment

Classical artificial intelligence proceeded under the assumption that cognition could be realized in computer programs that were not able to directly sense or affect the physical surroundings of the computer in which they were housed. This approach has led to many successful computer applications, but has shed relatively little light on the nature of cognition. Since the 1980s however, there has been a growing awareness that an agent must be able to act and be acted on by its surroundings, and sense the repercussions of those actions. This requires that an agent be both situated and embodied: it must have the ability to directly sense the world, and have a body with which to act on the world.

Josh Bongard
Morphology, Evolution and Cognition Laboratory, Department of Computer Science, University of Vermont, Burlington, VT 05405, USA
e-mail: `josh.bongard@uvm.edu`

S. Doncieux et al. (Eds.): New Horizons in Evolutionary Robotics, SCI 341, pp. 29–35.
springerlink.com

If a body plan, actuation and sensation are necessary to realize intelligent behavior, the question then arises as to how to choose an appropriate body plan for the desired behavior. Although it is generally agreed that the design of a robot controller for complex behavior is best left to automated optimization, many hold that human intuition can be applied to the design of robot body plans:

> "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." ([6], p. 22).

However, there are many explicit and implicit design decisions that must be made when choosing a robot body plan. As a simple example, a wheeled robot may be appropriate for rapid, efficient travel over flat terrain, but a legged body plan may be more suitable for rough terrain. If a legged body plan is chosen, how many legs should the agent have? Should the spine be flexible or rigid? If flexible, how flexible? This raises the issue then of how to systematically make these decisions.

2.2 The *How* of Embodiment

In order to overcome the infinitude of design decisions that must be made when designing a robot body plan, some researchers in the field of biorobotics design robot morphologies based on animal body plans. However, as has been discovered in the history of engineering many times over, direct copying of nature's designs does not always succeed. This is due to the face that the ecological niche inhabited by the animal and the robot are not usually equivalent: many aspects of the animal's body plan may have evolved for that niche and are therefore not applicable to the robot's niche. For instance, it took the realization that because aircraft are much larger than the birds which originally inspired their design, fixed wings are required rather than flapping wings. This then requires the roboticist to determine, for each aspect of the animal's body plan, which are design innovations that will still serve in the robot's ecological niche, and which are innovations evolved to meet biology-specific constraints. As an example of a biology-specific constraint, it has been hypothesized that the reason why wheels never evolved in animals larger than bacteria is because it would be difficult to provide nutrients across a freely-rotating axle (see for example [4], p. 542).

This difficulty argues against the claim made in [6] that human designers can formulate appropriate robot body plans. Instead, rather than mimic a product of natural selection, one can mimic natural selection itself, and harness it to design a robot's body plan along with its control policy. This is the field of evolutionary robotics, in which an evolutionary algorithm is used to automate the process of robot design.

2.3 The *Why* of Embodiment

However, such resulting algorithms are notoriously complex, and require more than the above verbal defense to justify their use. To this end, three arguments for why one should evolve body plans are given below:

- a physical robot that undergoes physical damage and thus sustains an unanticipated and unobservable change to its body plan can evolve a new description of its topology, and use the evolved model to generate a compensatory controller;
- virtual robots may initially be evolved with few body parts, and then gradually evolved with more body parts to accelerate evolution; and
- as the desired task become more difficult, the ability for evolution to alter robot body plans even slightly increases the probability of evolving a successful robot.

2.4 Why Consider Topological Change to a Robot's Body Plan?

One of the repercussions of considering that the body plans an important role in intelligent behavior leads one to consider how agents should deal with changes to that body. One class of such change is physical damage in which the robot loses one or more body parts. This leads to a change in the topology of the robot's body plan, and typically necessitates a different control policy to recover functionality.

This scenario was investigated in [2], in which we employed a physical quadrupedal robot (Fig. 2.1). This robot begins by evolving a body plan (i.e. its mental model) that accurately reflects the topology of its (initially unknown) physical body plan. This is accomplished by actuating the physical robot with a random control policy and recording the resulting sensor data; actuating a candidate mental model with the same control policy and recording the resulting virtual sensor data; and computing the error of the model as the difference between the physical and virtual sensor data. The evolutionary algorithm then searches the space of virtual body plans for one that minimizes this error within a given time period.

A second evolutionary algorithm then optimizes a controller on the mental model such that the mental model performs the desired behavior, which in that work was forward locomotion. Once such a controller is found, it is executed by the physical robot, which often results in the desired behavior.

Physical damage was then simulated by removing part of one of the robot's leg. As it could not directly sense the damage, the first evolutionary algorithm was restarted, which continues the search for an appropriate body plan. A new candidate body plan, when actuated with a control policy the physical robot has already performed, must reproduce the actions of the damaged robot; this rapidly leads to the evolution of a virtual robot that reflects the damage of the physical robot.

The first evolutionary algorithm is again paused, and the second evolutionary algorithm is restarted. This latter algorithm now searches for a compensating controller that allows the virtual, damaged robot to recover forward locomotion. Once found, this new control policy is executed by the physical, damaged robot; this often results in the recovery of locomotion.

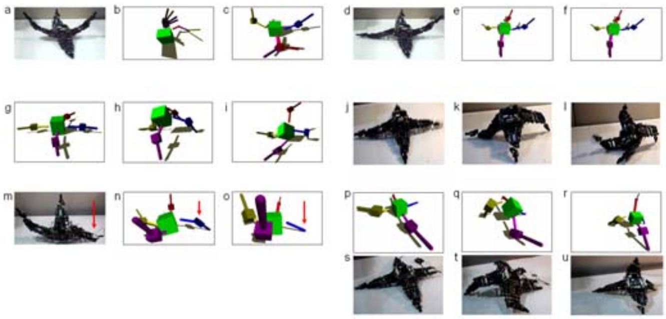

Fig. 2.1 The Resilient Machine Project. The robot performs a random action (**a**). A set of random models (one of which is shown in **b**) is synthesized into approximate models (one of which is shown in **c**). A new action is then synthesized to create maximal model disagreement and is performed by the physical robot (**d**), after which further modeling ensues. This cycle continues for a fixed period or until no further model improvement is possible (**e-f**). The best model is then used to synthesize a behavior (in this case, forward locomotion, the first few movements of which are shown in **g-i**). This behavior is then executed by the physical robot (**j-l**). The robot then suffers damage (the lower part of the right leg breaks off; **m**). Modeling then recommences with the best model so far (**n**), and using the same process of modeling and experimentation, eventually discovers the damage (**o**). The new model is then used to synthesize a new behavior (**p-r**), which is executed by the physical robot (**s-u**), allowing it to recover functionality despite this unanticipated change.

Thus, by considering the fact that a complex physical robot may sustain several physical damage, a solution presents itself in the form of an algorithm that evolves a virtual body plan to match the current state of the physical robot's body plan. This allows for not only the automatic generation of walking controllers, but also the automated recovery of walking after unanticipated damage.

2.5 Why Evolve Robot Body Plans Initially at a Low Resolution?

In attempting to automatically design robots capable of increasingly sophisticated behavior, it is often necessary to smooth the fitness landscape so that. This allows for incremental improvements from primitive behaviors to more complex ones. Several techniques exist for doing this, such as robot shaping (eg. [5]).

Another approach to shaping aside from simplifying the task environment is to simplify the robots' initial body plans. This approach has recently been explored in [1]. In that work purely passive three-dimensional structures were evolved to maximize their displacement when they fell. Initially, structures were grown at 'low resolution': relatively few, large spheres could be used to build the structure. Once structures with high fitness were found, the existing structures in the population were re-grown at 'high resolution': more, smaller spheres were allowed for construction. An evolved low- and high-dimensional structure are shown in Fig. 2.2.

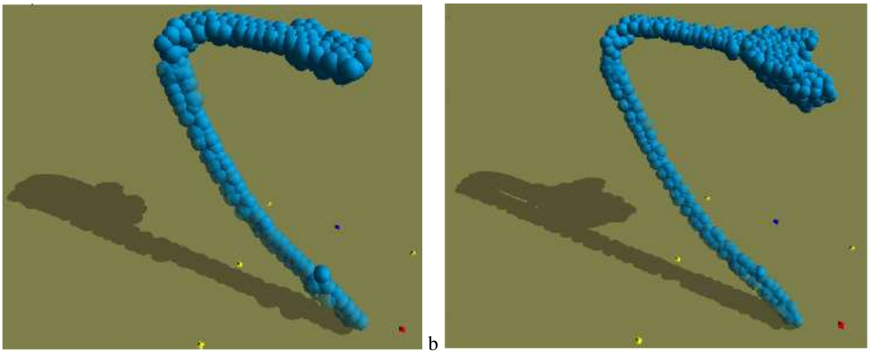

Fig. 2.2 Gradually Increasing Morphological Complexity. a: A structure evolved for maximum displacement when dropped. **a:** The same structure, but regrown at higher resolution.

Several runs in which low-resolution structures were initially evolved and then further evolved at higher resolution were performed, and compared against a set of control runs in which structures were evolved at high resolution throughout. It was found that both runs found structures of about the same fitness, but the variable-resolution runs found such structures in significantly less time.

This time savings was realized for two reasons. First, the low-resolution structures could be simulated and thus evolved more rapidly. Second, the evolutionary method employed increased the probability that a structure regrown at high resolution would have similar behavior—and thus similar fitness—as the structure originally grown from the same genotype at lower resolution.

This latter property was realized by adapting the neuroevolution technique HyperNEAT. HyperNEAT encodes a genome that takes as input a vector p, which indicates a position in a high-dimensional space, and outputs a value that can be used to construct a phenotype. The way in which genomes are encoded ensures that positions near one another output similar values.

In previous work in which the genotype was used to label the weights of a neural network, it was shown that HyperNEAT could evolve smaller neural networks that still functioned when regrown at higher resolution, because the synapses in the larger neural network lay near to their corresponding synapses in the smaller neural network.

This principle was exploited in [1]: during growth, the HyperNEAT genome takes as input a candidate position for placing a sphere, and the output indicates whether the sphere should be placed or not. When a structure is regrown at higher resolution, similar positions are queried as were queried during growth of the low-dimensional structure. This thereby increases the chance that the high-resolution structure will have the same shape—and thus behavior and fitness—as the low-dimensional structure.

2.6 Why Allow Body Plans to Change during Behavior Optimization?

A third reason for evolving robot body plans is that allowing evolution to change the body plan slightly increases the probability of finding a successful robot. This was demonstrated in [3]. For many robot body plans, slight changes will lead to slight changes in behavior, and thus slight changes in fitness. It is well-known that mutations that have slight phenotypic effect has a higher probability of being beneficial, compared to mutations that cause large phenotypic change. Thus, placing some aspects of a robot's body plan under evolutionary control can smooth the fitness landscape, and increase the evolvability of the overall system.

This principle was demonstrated using an anthropomorphic arm robot that was evolved to perform one or more object manipulation tasks: grasping an object, lifting an object, and or distinguishing between different objects. When robots were evolved for only one of these tasks, there was little evolutionary difference between runs in which only the control policy was evolved, and runs in which control and morphology were optimized. As the task became more challenging however by selecting for robots that could grasp, lift and distinguish between objects, runs in which control and morphology were evolved together performed much better than runs in which only control was evolved.

It is hypothesized that when the robots were evolved to perform all three tasks at once, the fitness landscape is more rugged because there are several trade-offs between these different competencies. For example a mutation that improves the robot's grasp may adversely affect its ability to distinguish between different objects: if it grasps all objects more tightly, it may fail to generate different sensor signatures when grasping the different objects and cannot thus distinguish between them.

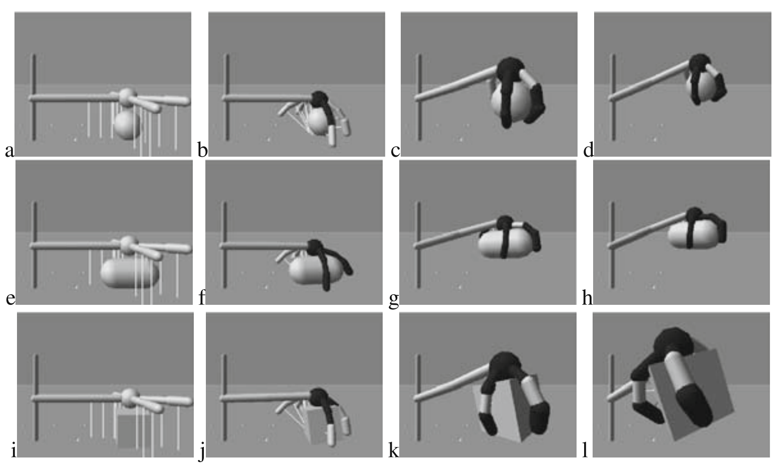

Fig. 2.3 The anthropomorphic arm evolved to lift, grasp and distinguish between different-shaped objects.

However, by allowing evolution to slightly change the radii of the fingers (as seen in Fig. 2.3), when it grasps objects of different shapes, different parts of the fingers may come into contact with the object. Each finger part contains a touch sensor, so grasping different objects may cause different subsets of the touch sensors to fire, and these different patterns can be used to distinguish between the objects. This is visualized in Fig. 2.3, in which finger parts are blackened if the touch sensor in it is firing, and different black-and-white patterns can be observed when the robot grasps different objects. Thus, slight changes to the robot's morphology can smooth the fitness landscape by reducing the tradeoffs between the different behaviors being evolved.

References

[1] Auerbach, J., Bongard, J.: Evolving CPPNs to Grow Three-Dimensional Physical Structures. In: Proceedings of the Genetic and Evolutionary Computation Conference, GECCO (to appear, 2010)

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

[3] Bongard, J.C.: The utility of evolving simulated robot morphology increases with task complexity for object manipulation. Artificial Life (2010), doi:10.1162/artl.2010.Bongard.024

[4] Chrisley, R., Begeer, S.: Artificial intelligence: critical concepts. Taylor & Francis, Abington (2000)

[5] Dorigo, M., Colombetti, M.: Robot shaping: Developing situated agents through learning. Artificial Intelligence 70(2), 321–370 (1994)

[6] Nelson, A., Barlow, G., Doitsidis, L.: Fitness functions in evolutionary robotics: A survey and analysis. Robotics and Autonomous Systems 57(4), 345–370 (2009)