

Evolved Machines Shed Light on Robustness and Resilience

This paper discusses the two-way interaction between brains and bodies, and the consequences for adaptive behavior, along with reviewing research that builds on insights from the neurobiology of these interactions to inform the design of evolving and adaptive robots.

By JOSH BONGARD, *Member IEEE*, AND HOD LIPSON, *Member IEEE*

ABSTRACT | In biomimetic engineering, we may take inspiration from the products of biological evolution: we may instantiate biologically realistic neural architectures and algorithms in robots, or we may construct robots with morphologies that are found in nature. Alternatively, we may take inspiration from the process of evolution: we may evolve populations of robots in simulation and then manufacture physical versions of the most interesting or more capable robots that evolve. If we follow this latter approach and evolve both the neural and morphological subsystems of machines, we can perform controlled experiments that provide unique insight into how bodies and brains can work together to produce adaptive behavior, regardless of whether such bodies and brains are instantiated in a biological or technological substrate. In this paper, we review selected projects that use such methods to investigate the synergies and tradeoffs between neural architecture, morphology, action, and adaptive behavior.

KEYWORDS | Embodied cognition; evolutionary robotics

Manuscript received November 1, 2013; revised February 25, 2014; accepted March 9, 2014. Date of publication April 7, 2014; date of current version April 28, 2014. The work of J. Bongard was supported by the National Science Foundation (NSF) Presidential Early Career Award for Scientists and Engineers (PECASE) under Grant 0953837; by the Defense Advanced Research Projects Agency (DARPA) M3 under Grant W911NF-1-11-0076; by DARPA MSEE under Grant FA8650-11-1-7155; and by Vermont Advanced Computing Core, which is supported by the National Aeronautics and Space Administration (NASA) under Grant NNX06AC88G. The work of H. Lipson was supported by DARPA under Open Manufacturing Grant W911NF-12-1-0449; and by NSF Cyber Enabled Discovery under Grant 0941561.

J. Bongard is with the Department of Computer Science, University of Vermont, Burlington, VT 05405 USA (e-mail: josh.bongard@uvm.edu).

H. Lipson is with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853 USA (e-mail: hod.lipson@cornell.edu).

Digital Object Identifier: 10.1109/JPROC.2014.2312844

I. INTRODUCTION

Embodied cognition [16], [65] (Fig. 1) is a particular paradigm for addressing the challenge of creating autonomous machines, as well as understanding how animals and humans produce adaptive behavior. It dictates that cognition cannot be understood—or realized in machines—without grounding it in behavior. Indeed cognition can be viewed as a collection of particularly subtle ways of exploiting one's interaction with the environment; an interaction that requires neural computation but also use of a body to carry out that interaction. It follows from this that there may exist a large number of pairings of particular mechanical constructs and neural architectures that, when combined, give rise to the behavior of interest. This raises several engineering challenges.

First, if a particular morphology is chosen, the designer may unintentionally constrain himself to a set of unsatisfying neural systems. For example, if a robot must travel over flat ground and a legged morphology is chosen, neural architectures which allow for movement may be much more complex than neural architectures that would allow a wheeled robot to travel successfully in the same environment. As an extreme example, passive dynamic walkers [20], [78] are bipedal robots that travel down declined planes and require the minimum of neural control: their particular morphology and specific environmental niche enables them to convert potential energy into kinetic energy while remaining upright with no control at all.

Second, how does one search the space of all possible body/brain pairings to find those that produce the desired behavior? Such search spaces are high dimensional as parameters describing both the morphological and control subsystems must be optimized. These spaces also contain many local optima: if search is operating in a part of the

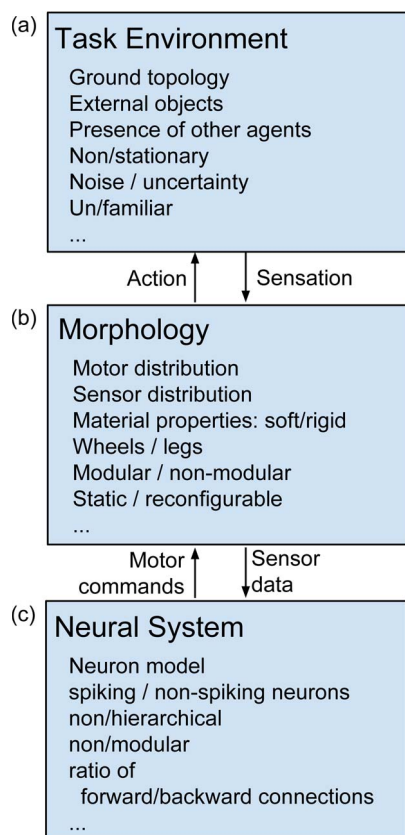


Fig. 1. Embodied cognition framework. In order to design an adaptive machine to operate in some task environment (a) it is imperative to consider aspects of the machine's mechanical construction (b) as well as aspects of its biologically inspired control policy (c). In order to improve the interaction of a machine with its environment one may alter its morphology, controller, or both.

space that contains robots with high centers of mass relative to their polygons of support, slight changes in neural control may cause these robots to become unstable.

Third, one may now consider a machine whose brain and body change while it behaves. A robot may employ a learning algorithm to improve its behavior over time, but if it is a modular robot [91] it may also reconfigure its body, or if it is a soft robot [73] it may alter the volume of different body parts in a simulation of growth. This morphological change may then hinder or ease learning.

Evolutionary algorithms [23] are a class of stochastic optimization methods that are particularly well suited for the challenges raised by embodied cognition. Designed well, they guard against becoming trapped in local optima. They are also more general than learning algorithms, which are constrained to optimizing only the neural subsystem of a robot. Finally, they readily allow for morphological and neurobiological change while a robot behaves (Section V): evolutionary algorithms can be

designed that dictate how a robot's body and brain should change over time, rather than dictating a static morphology and neural control system. The use of evolutionary algorithms to optimize autonomous machines is the subfield of robotics known as evolutionary robotics.

II. EVOLUTIONARY ROBOTICS

In the initial evolutionary robotics experiments [18], [32], neural controllers were optimized for physical robots with fixed morphologies. Neural control policies were generated on a computer, and downloaded serially onto the physical robot. The performance of the robot using a given control policy was recorded and returned to the computer. Once all of the robots in the population were evaluated, some would produce randomly modified offspring and others would die out. By repeating this cycle, robots capable of maze navigation [32] and visual discrimination tasks [18] were evolved.

Soon after, dynamical simulators were employed which allowed the evolutionary algorithm to simulate much larger numbers of virtual robots than could be evaluated in reality. Simulating robots provided another benefit: both the body plan and the neural controllers of the robots could be optimized [74]. Undulation was the primary movement strategy that evolved, which recapitulates what is observed in nature: undulation was the precursor to quadrupedal and then bipedal walking in higher animals. Fig. 2 outlines the basic process of simultaneously evolving the body plan and neural architecture of autonomous robots.

Some important work since the field's inception involved gradually introducing gradients into the search space. In [68], a biped was evolved to walk by applying counteracting forces if it became unstable; as evolution proceeded these forces were weakened and gradually removed. In [82], a passive dynamic biped with no controller was initially evolved to walk down a decline. Once passive walking was achieved, the robots were exposed to flat ground, which forced evolution to gradually incorporate actuation to support the gait evolved in the purely passive regime.

There has also been much work on evolving artificial neural networks with greater biological realism, including networks that allow for fine tuning of temporal dynamics [4] as well as spiking neural networks [88].

More recently, there has been progress on increasing the evolvability of evolutionary algorithms. Evolvability is defined differently by different authors but generally is considered to measure the propensity of a given evolutionary system to produce more fit solutions to a given problem over time [81], [83]. Much success has resulted from incorporating multiobjective optimization into evolutionary algorithms [25], [70]. For a long period there was no good method for introducing sexual recombination into the evolution of artificial neural networks, because




Fig. 2. Basic evolutionary robotics process. At the outset of optimization (a), a population of randomly generated robots are constructed. These robots differ in their morphologies as well as their neural control policies. (R1 = robot 1; circles represent neurons; lines represent synaptic connections between neurons; irregular shapes represent the robot's body plan; squiggles represent flagella; small and large arrows represent slow and fast speed of travel, respectively.) (b) Robots with low fitness (in this example, the slow traveling robots R1 and R3) are then culled from the population. (c) Robots that survive culling produce randomly modified copies of themselves. In this example, R2 produces an offspring R2' that incurs a mutation affecting its neural controller (the red synapse). R4', the offspring of R4 incurs mutations that affect a synapse and neuron (red line and circle, respectively) but also its morphology (a flagellum is added). (d) The behavior of the offspring is now evaluated, and their fitness values are computed. This cycle is repeated until a machine of sufficient fitness is obtained or a fixed amount of computational effort is expended.

combining material from two networks with different architectures rarely produced fit offspring networks. Stanley and Miikkulainen [75], however, introduced a method that enables evolution to alter and thus automatically improve the architectures of artificial neural networks by combining them based on their descent from a common ancestor. It has also been shown [8] that adding aspects of a robot's morphology to an evolutionary algorithm can increase its evolvability.

III. RESILIENT MACHINES

One of the major challenges in the field of evolutionary robotics is known as the "reality gap problem" [46]: once a robot of sufficient capability is evolved in simulation, how can its behavior be transferred seamlessly into a physical robot? Often, evolution specializes behavior to match specific details of the simulation in which the robots are evolved. For example, the temporal dynamics of the artificial neural networks may adapt to the time interval used for numerical integration in the simulator. Once realized in hardware, however, physical constraints may result in slower motors, or technological constraints may lead to slower updating of sensor readings. Early work in addressing the reality gap problem involved wrapping various aspects of the robot simulator in noise, given

knowledge about sources of uncertainty from the robot's sensors, motors, and environment [46].

This and other early approaches, however, require manually altering the simulation in which robots are evolved to better match reality. In [7], we introduced a method for automatically altering the simulation to reflect the current state of the robot: this involved restructuring the body plan of the simulated robots to better reflect the mechanical construction of the physical robot. To accomplish this, the physical robot maintains three separate stochastic optimization methods, as shown in Fig. 3.

These three optimizers together produce several forward and inverse models of the robot. Forward models are known to exist in the brains of higher animals [86] and are somehow employed to predict future sensation, given a candidate action. In the resilient machine project, forward models took the form of simulated robots, as shown in Fig. 3(b). An evolutionary algorithm searched the space of possible robot body plans to find those that could, when actuated using the same motor programs that the physical robot had performed, predict the sensor data generated by that motor program by the physical robot. Note that no assumptions were made about the structure of the robot: it had to learn its own structure through physical experimentation. This is known as structural learning in the motor learning literature [85], and is a particular talent of humans: we must learn to maintain coordination even

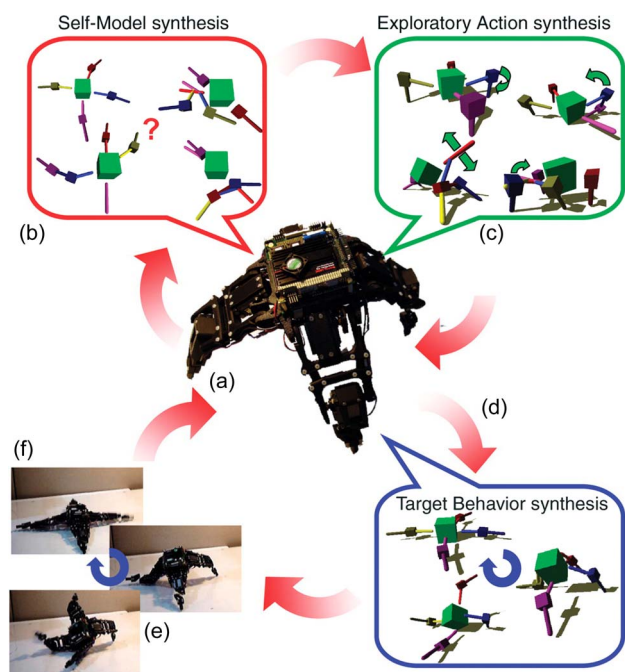


Fig. 3. Resilient machine. A physical robot (a) maintains three separate optimization methods (b)–(d). In the first optimizer (a), the robot performs structural learning [85]: it must search through the space of all possible configurations of its own body [four examples are shown in (a)] to find those that can most closely reproduce the sensor data generated by the physical robot, when actuated with the same motor program that the physical robot has executed. With little sensorimotor data, multiple structures may explain the data. This requires the robot to search for a new action to perform (b): candidate actions not yet performed are represented by green arrows. Once it finds an action likely to yield new information, the physical robot performs it (a). This cycle continues until the robot models converge. The physical robot then optimizes a neural controller using the best model found so far to perform some desired behavior, which in this example is forward locomotion (d). The physical robot then employs this evolved controller to produce the desired behavior in reality (e) and (f). Photo courtesy of the American Association for the Advancement of Science (AAAS); video courtesy of the authors.

when the structure of our body changes because we swing a tennis racket or don skis.

An advantage of not assuming what the structure of the robot should be *a priori* means that the evolutionary algorithm can continue to adapt the structure of the forward models if the morphology of the robot changes. We demonstrated such a simulation by damaging the robot: we separate one of the robot's four legs. The robot was not equipped with any sensors that could directly detect this separation, and the robot could continue to send commands to the dragging motor (although they would not affect the overall movement of the robot). Instead, the robot was only equipped with vestibular sensors that could detect the orientation of the robot's main body. However, with a sufficient number of physical actions, the evolutionary algorithm would eventually

determine that the only forward models that explained the robot's new sensor/action data pairs were models that reflected the damage that the robot had suffered.

The virtual robots that evolve to reflect the morphology of the physical robot can also be viewed as body schema: neural encodings of body structure and dynamics [42]. It is known that the human brain maintains multiple representations of the body, but much controversy remains about how many representations there are and how they are encoded. There is general consensus, however, that the brain maintains a separation between body schema employed for action and those used for perception [24]. In [7], no such separation was made: forward models were constrained to actuable kinematic chains. It would be of interest to conduct further experiments in which the evolutionary algorithm would be free to formulate its own body schema representations appropriate to the robot and its environment. Indeed Massera *et al.* [54] investigated active categorical perception—the ability to distinguish between objects by actively manipulating them—and demonstrated that an evolutionary algorithm could devise its own neural encoding for objects of different shape.

We could then compare the encodings discovered by artificial evolution to those encoded in biological nervous systems. If these artificial and biological encodings differ, this suggests that it does not matter how bodies are represented in artificial or biological nervous systems. If these encodings resemble one another, this suggests that there is a particular way that bodies should be neurally represented to enable successful prediction of action in particular, and thus enable rapid generation of adaptive behavior in general. Indeed, this is but one way evolutionary methods could be employed to study the role of mental representation in complex behavior, which has remained a central and controversial subject in artificial intelligence research for decades [12].

Once the physical robot has evolved a sufficiently accurate forward model, it can use the second optimizer to evolve neural controllers so that the robot performs some desired task [Fig. 3(d)]. In the experiments reported in [7], the physical robot's task was forward locomotion: a population of neural controllers are sequentially tested on the evolved forward model [one is shown actuating the forward model in Fig. 3(d)]. Once a neural controller is evolved that allows the forward model to move sufficiently fast, it is employed by the physical robot [Fig. 3(e) and (f)]. We can view the evolutionary algorithm that produces these neural controllers as an inverse model. An inverse model in biological nervous systems takes as input some desired sensation and produces as output a set of actions that will realize that sensation [85]. This second evolutionary algorithm takes as input a fitness function that maximizes forward acceleration and outputs a neural network; this neural network generates motor commands in either the simulated or physical robot that result in this sensation.

In the event that the robot undergoes damage and the forward models evolve to reflect this damage, the neural controllers evolved by the second optimizer no longer produce the desired behavior in the new forward models. This causes the second optimizer to rapidly replace these failed neural controllers with newly evolved neural controllers that do produce the desired behavior in the damaged robot. Taken together, this results in resiliency: when the physical robot experiences an unanticipated situation for which it has no contingency plan, it diagnoses the situation by fitting new forward models to it and then evolving a qualitatively different behavior to compensate for the new situation. Resiliency—abandoning a behavior that is no longer relevant and generating a new one that is—differs markedly from robustness, which is the ability of an organism or machine to continue operations in the face of noise or change.

The third optimizer searches through the space of possible actions that the physical robot should perform next [Fig. 3(c)], when there is currently insufficient information to distinguish between competing forward models [Fig. 3(b)]. There is relatively well developed theory from the machine learning field that dictates how to do so: the physical robot should perform an action that induces the most disagreement among the predictions of the current forward models [72]. Thus, this third evolutionary algorithm uses prediction disagreement as the fitness function for deciding which candidate actions to delete and which to reproduce. If one observes the physical robot from a distance, it seems to perform a series of random actions, despite the fact that the selected actions are decidedly nonrandom. This method may prove useful for developmental psychologists, who theorize about the amount—if any—of randomness in the seemingly random movements infants make to learn about their own bodies and the world around them [40], [69].

In summary, one evolutionary algorithm searches the space of possible forward models for those that most closely resemble the physical robot [Fig. 3(b)]; a second searches the space of neural controllers that enable an evolved forward model to perform the desired task [Fig. 3(d)]; and a third searches the space of candidate actions that will generate the most information about the robot's current state.

A. Neuronal Replicator Theory

Each of these three optimization methods maintains its own population of forward models, neural controllers, and candidate actions. Within each population, individuals compete to explain the physical robot's state, generate desired behavior, and induce disagreement among the forward models, respectively. There has been long debate within neuroscience about whether and how patterns of neural activity or neural circuits themselves compete against one another in some form of a Darwinian process [14], [22], [29]. Recently, it has been shown [31] that

temporal dynamics within one set of neurons can migrate to another set, and that selection can cause certain patterns to proliferate at the expense of others. Although no evidence has yet been found for such Darwinian processes occurring in actual nervous systems, their presence has not yet been ruled out, nor is it clear what role such competition might play in the generation of adaptive behavior.

B. Deep Learning

Recently, a theory of general brain function has been proposed and become increasingly popular: hierarchical predictive coding [17], [37]. In short, it has been argued that the brain continuously makes predictions about arriving sensory inputs, and these predictions are organized hierarchically: cortical regions responsible for abstract cognition may form high-order predictions that span different perceptual systems such as visual, auditory, and tactile sensations (e.g., “I am seeing and holding a hat”); as these predictions project “forward” toward specific perceptual systems the predictions become more specific (e.g., “If I rub my fingers I will feel the softness of the hat's material”).

If there is a mismatch between predicted and actual sensation, the resulting error propagates upward to restructure predictions generated by higher brain regions, triggers actions that correct the error [35], and/or attracts attention to this unexpected stimulus. Part of the allure of this hypothesis is that it is intuitive; it explains the prevalence of forward connections in the brain (projections from cortical regions toward perceptual systems); it unifies various aspects of perception, action, and attention; and computational methods based on this concept, known as deep belief networks, have shown superior performance in pattern recognition tasks [41]. Again, there is much controversy about whether this is the primary function of the brain [17] and, if it is, what form do predictions take at different levels in this hierarchy. One could imagine a neural system that combines the advantages of multiple forward models as demonstrated in the resilient machines project (Section III), Darwinian competition between forward models of various levels of detail and locality at different levels in a hierarchy (Section III-A), and different actions designed to disambiguate between predictions made by these forward models at different levels of granularity.

C. Multiple Resilient Machines

In [10], we expanded the resilient machines approach to robot swarms. As shown in Fig. 4, each member of a team of robots constructs models of itself and chooses particular actions designed to accelerate that process. Periodically each robot exports its best self-model and imports self-models generated by its peers. If performed by a group of sufficiently similar machines, each robot is able to generate a model of self more rapidly than if it had done so individually. This is enabled by incorporating sexual

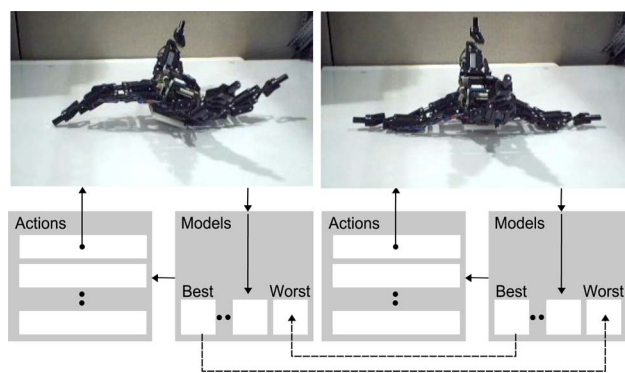


Fig. 4. Robot pair creating forward models together. Each robot maintains its own self-modeling engine (Fig. 3). Periodically each robot exports its best self model to its peer and overwrites its worst self-model with the best self-model received from its peer. Each robot decides independently which action to perform to improve its own current set of self-models. Image courtesy of IEEE.

recombination into the evolutionary process that generates self-models: a chance crossing may combine one correctly modeled part of the body from an imported model with another correctly modeled part of the body from a native self-model.

Collective construction of self-models can also accelerate the recovery from injury. If one robot becomes damaged, it may perform actions to diagnose that damage and then evolve a compensating neural controller to recover movement despite its injury. If robots are sufficiently similar, it is likely that they will experience similar wear and tear, or damage will occur at similar weak points in the robots’ construction. If one robot suffers damage, recovers, and broadcasts its self-models to its peers, the evolutionary algorithms generating self-models in its peers will automatically discard the damaged self-model because this model does not fit their sensorimotor experiences. If, however, a similarly damaged peer imports this damaged self-model, it will automatically bypass the damage diagnosis stage: the damaged model already explains the new sensorimotor experiences obtained by the damaged machine.

This ability to directly share experiences obviously has no equivalent in the biological realm. However, humans do have the ability to combine language and physical demonstration to communicate to a peer how a new activity or body state may feel, and how to fight or channel that experience to master a new skill. For example, an instructor may counsel a student to fight the urge to lean forward when first riding a snowboard because leaning back will recenter the student’s center of mass over their polygon of support.

D. Social Resilient Machines

In addition to modeling self, it is of great utility for a social organism or machine to also model others [56]. For

example, such models may help the observer to predict the physical capabilities and limitations of others and—along with predictions about the sensory repercussions of one’s own actions—make predictions about whether a contemplated social interaction will succeed. Indeed much work in the field of social robotics is concerned with modeling self and others [39], [48], [62].

In [51], we reported on a student robot that attempts to detect, model, and then imitate other robots that are sufficiently similar to itself (Fig. 5). It first creates models of self using the method described above. It then, upon detecting another machine, constructs multiple kinematic chains of differing structure to describe it. It then samples pixels from the video stream coming from its camera in a manner analogous to how it determines which actions to perform when modeling itself: it “looks” at each of the models of the candidate teacher robot, and searches for pixels in each mental image that differ. For example, if one teacher model takes the form of a long robot lying close to the ground but another teacher model represents a tall, upright robot, many of the lower pixels in the former




Fig. 5. Self- and other modeling. A student robot (a) observes other robots in its environment in an attempt to determine which are sufficiently similar to imitate. Although robot (b) has a similar visual appearance and the same degrees of freedom, the student robot determines that it is sufficiently structurally different from itself to make imitation impossible. The student robot then models robot (c), determines that it is sufficiently similar to it, and learns to imitate its movements. Image courtesy of Neural Networks.

image will be dark (because the candidate teacher robot is dark) but will be light in the latter image (assuming the teacher robot is moving against a light background). If the physical student robot then queries a low pixel from its actual camera and detects darkness, this lends support to the hypothesis that the other robot is lying on the ground while weakening the hypothesis that the robot is tall and thin. By continuously evolving multiple models of the other robot and repeatedly sampling only one pixel from its video stream the student robot can rapidly infer its kinematic structure. By performing this process repeatedly as the candidate teacher robot is moving, the student robot can also infer the changes in its joint angles.

Finally, armed with a model of self and a model of the other, the student robot attempts to determine whether it is possible to imitate the other. It does so by evolving a neural controller that actuates its self-model. In this evolutionary algorithm, the fitness of any neural controller is determined to be how close it can cause the self-model to move like the teacher robot. If a controller evolves that reproduces the teacher robot's movements sufficiently closely in the self-model, then that teacher is deemed to be imitable; otherwise, the student robot continues its search for an appropriate teacher. If the physical student robot now employs this evolved controller, it will imitate the physical robot in its field of view.

There are two points of contact between this particular approach to social robotics and neuroscience. First, the visual method employed here suggests a possible link between active perception (selecting a subset of available sense data through action [66]), the neuronal replicator theory, and the hierarchical predictive coding hypothesis: organisms may indeed learn multiple forward models of different levels of detail; models at each level of a hierarchy may compete with one another to explain previous experience; but models at each level may also collectively determine which perceptual streams to sample from, or which actions to perform to collect the required sensory input. Second, it is known that there are certain brain regions that activate if an animal performs an action or observes another animal performing that action. Neurons in these regions are colloquially referred to as mirror neurons [36], but much controversy surrounds the role (if any) such neurons play in social skills such as imitation and empathy [43]. It may be that mirror neurons represent some linkage between forward models of self and forward models of others that have been genetically encoded or learned. However, despite some preliminary evidence relating mirror neurons to forward models of self [58], more such evidence is required to support this hypothesis.

IV. MODULARITY

It is well known that both structural modularity (i.e., modular wiring [57]) and functional modularity (i.e.,

discrete sets of activation patterns [3]) exist in the brains of higher animals, but much work remains to determine the evolutionary and developmental causes that give rise to them. We must determine what such causes are because naive optimization of neural controllers for robots rapidly leads to increasingly densely interconnected networks [9]. This dense connectivity in turn slows optimization, because any slight change in a partially optimized network causes functional change to propagate throughout the network, usually leading to a decrease in performance. Also, we cannot simply copy the modularity observed in biological nervous systems into artificial neural networks, because the imported modularity may not match the functional demands of the robot's environment. However, if we can create an evolutionary algorithm that exerts selection pressures that favor functional and modular networks, we can automatically derive modular robot neural controllers and thus obtain all the benefits of modularity such as robustness and the ability to adapt to novel situations.

It has been argued that metabolic cost of connections between neurons can account for modularity [19], however this criterion is not sufficient: a sparse but nonmodular network will contain many fewer connections than a densely interconnected modular network. It has also been shown [52] that evolutionary change can favor the evolution of modularity in a computational model.

Wagner has argued [84] that evolution will produce modular phenotypes if directional selection acts on one part of the phenotype and stabilizing selection acts on another part. Directional selection is an evolutionary force that causes a structure or behavior of an organism to change over time; stabilizing selection favors organisms that maintain a given structure or behavior. Espinosa-Soto and Wagner [30] instantiated this theoretical argument in a computational model in which networks, meant to represent genetic regulatory networks, were modeled as dynamical systems: some parts of the network were challenged, over evolutionary time, to maintain an earlier evolved attractor landscape (stabilizing selection). Other parts were challenged to add new attractors as evolution proceeded (directional selection). This led to the evolution of increasingly modular networks, without explicitly selecting for modularity.

Kashtan and Alon [49] also demonstrated the spontaneous evolution of modularity using what they termed "modularly-varying goals." Such goals are reminiscent of the directional/stabilizing selection approach to modularity in that some subgoals within the overall goal change over evolutionary time while other subgoals remain constant. This goal structure comes to be reflected in increasingly modular logic circuits: relatively independent subcircuits evolve to compute results for each of the unchanging subgoals, while sparse connections between the subcircuits change to track changes in the time-varying subgoals.

A. Modularity in Robot Neural Controllers

However, in all of these previous computational approaches to modularity [19], [49], [52], [84], embodiment was not taken into account: the logic circuits or neural networks did not control the movements of a robot. A notable exception is the work of Yamashita and Tani [89], in which it was shown that *functional* modularity—the ability of a neural network to exhibit independent, stable dynamical states—could become organized into a hierarchy: neurons with slow temporal dynamics could trigger collections of fast neurons to switch between different stable patterns, which correspond to specific robot behaviors. In [9], we introduced a method following Espinosa-Soto’s approach in which action played a central role in the spontaneous evolution of *structural* modularity (densely connected neural modules with sparse connectivity between modules). Neural controllers were evolved to drive a robot arm to grasp differently sized objects placed at different positions (Fig. 6).

At the outset of an evolutionary trial, neural controllers were only selected for their ability to drive the robot to grasp a large object placed on its left [Fig. 6(a)]. Evolution proceeded until such a network was found, after which time subsequent neural controllers were selected for their ability to generate successful behavior in two environments [Fig. 6(a) and (b)]. Once a network evolved that enabled success in both environments, its descendents were evolved to produce successful behavior in three

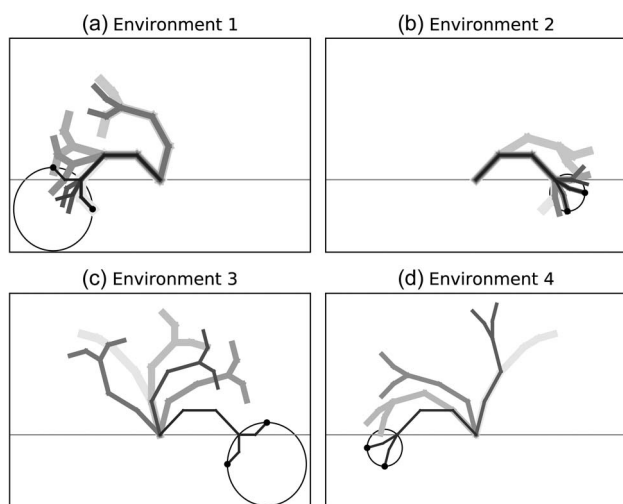


Fig. 6. Robot arm evolved for object grasping. The robot arm, when exposed to a large object on its left (a), rotates to the object and places its two fingertips (black dots) on the circumference of the object. Thick light lines represent the robot’s position at the outset of its motion; thin black lines represent its final position. The same robot’s behavior when exposed to a small object on its right (b); a large object on its right (c); and a small object on its left (d). This robot successfully “grasps” all four objects. (a) Environment 1. (b) Environment 2. (c) Environment 3. (d) Environment 4. (Reproduced from [9].)

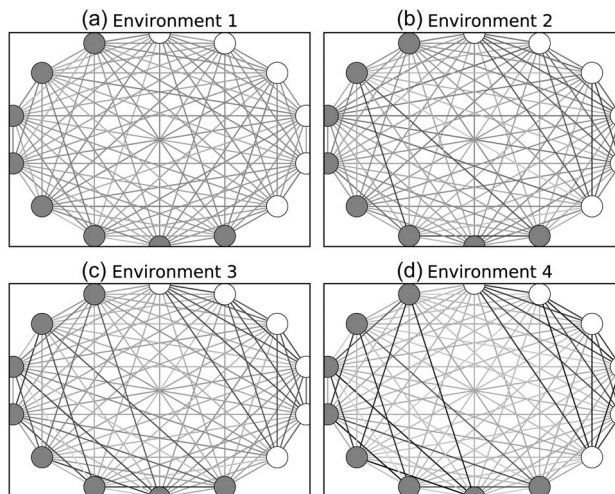


Fig. 7. Spontaneous evolution of modularity in robot neural controllers. (a) Mean connectivity across neural controllers that enable the robot to succeed in environment 1 [Fig. 6(a)]. Gray circles represent hand motor neurons, white circles represent arm motor neurons. The lines represent the probability of a synaptic between each pair of these neurons across the evolved networks: lighter and darker lines indicate a lower and higher probability of that connection occurring in any given network, respectively. (b) Mean connectivity across evolved networks that enable the robot to succeed in environments 1 and 2 [Fig. 6(a) and (b)]. (c) and (d) report the mean connectivity across networks that enable the robot to perform in the first three environments and all four environments, respectively. (a) Environment 1. (b) Environment 2. (c) Environment 3. (d) Environment 4. (Reproduced from [9].)

environments [Fig. 6(a)–(c)]. Finally, neural controllers were evolved to enable object grasping in all four environments [Fig. 6(a)–(d)].

Many trials were conducted, and successful neural controllers at the end of each of these four phases were collected together. It was found that among neural controllers evolved against just the first environment, no modularity was present: the density of synaptic connections between arm and hand motor neurons was not appreciably different from the density of connections among arm motor neurons, or among hand motor neurons [Fig. 7(a)]. However, there was significant modularity present among the neural controllers that produced successful behavior in all four environments: there was significantly greater connectivity among hand motor neurons and among arm motor neurons than there was between hand and arm motor neurons [Fig. 7(b)].

One can account for the increasing modularity by applying the directional/stabilizing selection argument to the movements of the robot’s arm and hand. When a familiar object [such as the object in Fig. 6(a)] appears in an unfamiliar location [as in Fig. 6(c)], one possible solution is for the hand to perform the same grasp, but for the arm to act different (rotate right instead of left). In effect, stabilizing selection is acting on the hand while

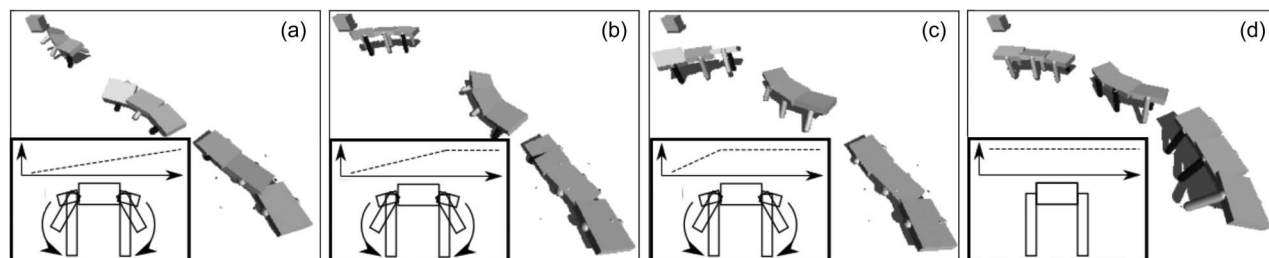


Fig. 8. Incorporating development into the evolution of robot behavior. A population of simulated robots are evolved to travel to a target object placed in their environment (small cube). At the outset of an evolutionary trial, the robots' morphology changes while they behave: they transition from an infant, legless form into an adult legged form while moving (a). The inset graph indicates how the robot's morphology changes: the robot gradually "grows" legs, and those legs become increasingly vertical relative to the body. The inset graph shows how the robot's center of mass changes during the evaluation period. Evolution continues until a neural controller is found that successfully enables the robot to reach the target object. When this occurs, the robots' developmental trajectory is altered: subsequent neural controllers must control a robot that grows from the infant into the adult form over the first two-thirds of the evaluation period and then maintains the upright legged for the remaining third of the evaluation period (b). Once successful behavior is regained, the infant to adult form is accelerated further (c) and evolution continues. In the final stage, the infant form is discarded completely (d) and evolution continues with the existing neural controllers until they produce successful legged locomotion to the target. (Reproduced from [11].)

directional selection is acting on the arm. Alternatively, in comparing the successful behaviors in environments 2 and 3 [Fig. 6(b) and (c)], stabilizing selection acts on the arm while directional selection acts on the hand.

Importantly, neural dynamics were found to play an important role in the evolution of modularity. For example, it was possible for a successful neural controller to fall into a point or cyclic attractor of neural activity. A point attractor results in the robot coming to rest with a successful grasp; a cyclic attractor results in the robot alternatively grasping and releasing the object. If selection did not explicitly select for point attractors, modularity would not evolve. It is currently unknown why this is so, nor what the general relationship is between neural dynamical behavior and conditions that favor the evolution of modularity. Additionally, it is not clear how robots of greater complexity should be evolved to favor the hierarchical and modular architecture ("modules within modules") observed in biological nervous systems [57].

V. MORPHOLOGICAL SCAFFOLDING

A majority of work in biomimetic or neuromorphic computation focuses on learning or other neural change which leads to improved function as the robot or device interacts with its environment (for examples, refer to [33] and [53]). However, very little work has focused on the role of morphological change in the acquisition of adaptive behavior. One notable example comes from developmental robotics [1] (a sister field to evolutionary robotics), which takes inspiration from human cognitive development.

Developmental roboticists have shown how altering a robot's body can provide a solution to Bernstein's problem in motor control, which states that there are many ways to perform a movement for an animal with a large number of

mechanical degrees of freedom [5]. They demonstrated that initially reducing a robot's number of mechanical degrees of freedom can facilitate the learning of motor control. Once the robot achieves some competency, more of its degrees of freedom are released and coopted by the developing neural controller [6]. A popular approach to Bernstein's problem is to use optimal control: among the many (possibly infinite) movements that may result in success, the movement that minimizes some cost function is chosen as the optimal way to achieve the given task [26], [79]. However, optimal control approaches usually consider the musculoskeletal system of the animal (or mechanical system of the robot) fixed, despite the observation that animals with large numbers of degrees of freedom undergo significant morphological change over their lifetimes. Indeed humans in particular exhibit both high morphological complexity and extreme morphological change.

In a recent study [11], we investigated the role that metamorphosis might play in the evolution of neural controllers for robots. We began by formulating a control experiment, in which artificial neural networks were evolved for a hexapod such that it would locomote toward a target object placed in its environment [Fig. 8(d)]. In this study, the topology of the neural network controller was held constant and the evolutionary process only optimized the strength of the synaptic connections between the robot's sensor neurons, its sensor neurons and its interneurons, and its interneurons and its motor neurons. The evolutionary algorithm terminated once a neural network evolved that brought the robot sufficiently close to the target object. We ran several such trials and recorded the average time it took evolution to discover such controllers.

We then performed a second set of trials in which we introduced metamorphosis. At the outset of the evolutionary

trial additional mechanical degrees of freedom were added to the robot such that it would gradually reconfigure from a legless segmented robot [lower right robot in Fig. 8(a)], while being controlled by a given neural controller, into an upright legged hexapod [top left in Fig. 8(a)]. The evolutionary algorithm continues until a neural controller is found that produces successful behavior in this metamorphosing robot [Fig. 8(a)]. When this occurs, the robot's developmental trajectory is altered. The robot now transitions from the infant legless form into the adult legged form over the first two thirds of its evaluation period [Fig. 8(b)]. The neural controllers previously evolved are reevaluated on this altered robot, which typically reduces the quality of behavior they produce. Evolution thus continues until it rediscovers a neural controller that works for the altered robot. Once success is again achieved, evolution of the neural controllers continues on a robot that metamorphoses from the infant form to the adult form during the first third of its evaluation period [Fig. 8(c)]. When success is again achieved, the infant form is discarded altogether [Fig. 8(d)]: the evolved controllers are reevaluated on the robot which now begins with and maintains a hexapedal form through its evaluation period. Evolution continues until a neural controller evolves that generates the desired behavior in this morphologically fixed robot. We conducted several trials using this approach and measured the average time it took to progress through all four evolutionary stages.

Despite the three prepended evolutionary stages with metamorphosing robots, these experimental trials took less time than the control trials. The reason for this is that optimizing controllers for dynamically unstable machines like legged robots is difficult: there are few gradients in the search space. However, a legless robot creates a much smoother search space: any slight change to a given neural controller may slightly increase or decrease the travel speed of such a robot. By forcing these robots to change into legged robots while they move, however, we force the optimization method to find controllers that enable movement in both robot forms rather than specialized controllers that only work for the legless robot. Then, by gradually removing the infant form over evolutionary time, we gradually remove the gradients from the search space. However, the infant robot form has brought search into the neighborhood of neural controllers that work for the legged robot.

Not only were controllers evolved for the hexapod more rapidly if metamorphosis was employed, but also the resulting controllers were more robust. The successful neural controller returned at the end of each control and experimental evolutionary trial was reevaluated on the hexapedal robot, but under novel environmental conditions. The robot was exposed to small, random external impulse forces, simulating mild collisions or wind. It was found that the robot was better able to continue locomotion despite these novel perturbations when

equipped with the neural controllers from the experimental trials than when it was equipped with neural controllers from the control trials. This observation can be explained by the fact that controllers evolved across robots with differing morphologies must maintain legged locomotion over a wider range of sensor-motor contingencies than controllers evolved for a single robot. In these experiments, this evolved robustness against internal change (metamorphosis) translated, after optimization, to robustness against external perturbation.

Although metamorphosis has been explored with robots previously [61], [90], the role that such change has on the ability to generate adaptive and robust behavior was hitherto unknown. In the biological realm, there is some understanding about how motor control changes during metamorphosis [21]. However, the way in which behaviors obtained prior to metamorphosis facilitate the acquisition of behaviors after metamorphosis is largely unknown, from a neurobiological perspective: for example, do the behaviors that a human infant learns (crawling and then scooting) somehow prepare it for learning different yet related behaviors (walking and running) as it transitions into an adult form? Such questions are related to scaffolding, a concept originating in psychology: how can a caregiver structure the environment of a learner such that it learns a behavior more rapidly than if it attempted to learn it in an unstructured environment [87]? This idea has carried over into robotics in which the robot's environment is often scaffolded [28]. Our work in [11], however, can be viewed as a form of *morphological scaffolding*: the robot's own body may allow it to rapidly acquire a behavior; subsequent body changes may then increase the robustness of that behavior when exhibited in the final, fixed adult form.

Finally, the relationship between the robustness of neural control and metamorphosis is little explored. Humans excel in adapting behavior to suit environmental change, and clearly much of that ability can be attributed to learning. However, some of this plasticity may also be caused by the radical changes we experience in growing from infants into adults. Robots may serve as a uniquely suited tool for investigating this connection.

VI. MORPHOLOGICAL COMPLEXITY

There is a large body of work in the neuroscience literature that makes use of information theory [27]. Some work uses information theoretic measures to gauge how much information about a stimulus or behavior is encoded by a population of neurons [67]. Such metrics may also be employed to characterize the complexity of nervous systems [13]. Indeed it has been argued that such measurements indicate that the cerebral cortex exhibits higher complexity than other brain regions, which may reflect the ability of higher organisms to successfully grapple with a complex environment [80].

Similarly, evolutionary biologists have explored the evolution of morphological complexity. However, much analysis has focused on macroscopic trends toward increasing complexity [15]; whether such rises in complexity are the result of selection or simply occur because of drift through the space of possible phenotypes [55], [59]; or investigation of the genetic underpinnings of phenotypic complexity [34]. Relatively little work, however, has focused on determining the adaptive advantage of morphological complexity. An exception to this is the work of Passy [64], who determined that increased morphological complexity in the form of internal variation may give an organism an advantage in changing environments: Passy found that the fractal dimension of different diatom species correlates with the variability of their particular environmental niche.

One can often detect a “neurochauvinist” bias in artificial intelligence and robotics research: more work is often dedicated to modeling neural systems than in designing the bodies into which these neural models will be housed. This chauvinism is evidenced by the greater amount of work on neural complexity than on morphological complexity in the AI and robotics literature. Dating back to the earliest computational neural models, it was demonstrated that artificial neural networks required a minimum level of complexity—in this case, the inclusion of interneurons—in order to correctly classify data sets that are not linearly separable [45]. More recent work has investigated the relation between neural complexity and adaptive behavior [47], [71]. However, only a few studies have investigated the relationship between morphological complexity and behavior [2], [38].

If adaptive behavior can best be understood by considering how an animal’s brain (or robot’s neural controller) mediates the interaction between its body and its environment [16], [65], it follows that the body will constrain as well as afford particular ways in which this interaction can be exploited. From this, it follows that for any given task, different pairs of robot body plans and neural controllers will result in the desired behavior. Finally, some of these pairs will contain morphologies of greater complexity and neural controllers of less complexity than other pairs. As described in Section I, there are many legged robots that are capable of locomotion, but the passive dynamic walkers provide an example of a brain/body pairing with zero control complexity: these walkers have no controllers at all [20], [78].

For these reasons, it would be helpful to understand the relationships between morphological, neural, and environmental complexity. Such an understanding might help to determine, for example, the minimal amount of morphological and control complexity required for success in a given task environment. In a recent work [2], we explored one of these relationships, that between morphological and environmental complexity.

To do so, we evolved robots in simple and complex environments. The simple environment required the robot




Fig. 9. Robots whose morphologies and neural controllers were evolved in different environments. If robots are evolved to move over flat, high-friction terrain, relatively simply shaped robots arise (top left). If, however, robots are evolved to travel over low-friction objects (blue bars), they evolve more complexly shaped body plans which enable them to reach within the crevices between blocks and propel themselves forward. (Reproduced from [2].)

to locomote across a high-friction, flat ground plane [Fig. 9(a)]. The complex environment was composed of a series of low-friction blocks sitting atop the high-friction ground plane [Fig. 9(b)–(d)]. This cluttered environment is not only more complex but it is more challenging: robots evolved to travel over these objects had to evolve the ability to reach between the blocks and push or pull themselves forward.

We then employed an evolutionary algorithm [76] that enabled us to evolve the 3-D shape of the robot, along with parameters describing a neural controller of fixed topology: thus evolution could complexify (or simplify) the robot’s morphology but not its neural complexity. (In future work, we intend to explore the evolutionary tradeoff between morphological and neural complexity.) The robots comprised triangular meshes, which facilitated our ability to measure the morphological complexity of a given robot. We measured complexity using shape entropy [63], which equates the complexity of a 3-D shape with the variation in local curvature on its surface. (It has been shown that high values of shape entropy correspond with human observers’ attribution of complexity [77].)

After evolving robots in the simple and complex environments, we compared the average shape entropies of both sets of robots. When no cost was placed on morphological complexity, no significant difference in morphological complexities between the two sets of robots evolved. However, when the robots were evolved in both environments again, and a cost was placed on complexity,




Fig. 10. Evolution of morphological complexity. Robots evolved to travel in complex, challenging environments show more of an increase in morphological complexity (blue curve) than robots evolved to move in simpler environments (black curve). In this example, the robots evolved in the complex environment also evolved to move slightly further (red curve) than those evolved in the simpler environment (green curve). Dotted lines denote one unit of standard error. (Reproduced from [2].)

robots in the simple environment were simpler [Fig. 9(a)] than the robots evolved in the complex environment [Fig. 9(b)–(d)]. This different rate in the evolution of morphological complexity is depicted in Fig. 10. Although this situation did not arise in every pairing of a simple and complex environment studied in [2], the inverse was never observed: there were no simple and complex environment pairings in which more complex robots evolved in the simpler environment. It may not be surprising that such a correlation does exist between morphological and environmental complexity in this simple domain—or that this correlation may not hold in other task environments—but this was the first quantitative determination of such a relationship.

Many modern evolutionary methods bias search toward increasingly complex solutions as optimization proceeds, regardless of what is selected for, including the method we used in this work [76]. So, we had to rule out that the observed increases in complexity were the result of bias, rather than because they were useful for the task at hand. To do so we ran several evolutionary trials with the exact same parameter settings but with selection turned off: robots were selected for culling or reproduction randomly, rather than their ability to locomote. We then measured evolutionary changes in morphological complexity within these control trials. We found that morphological complexity does indeed increase, but it increases at a much lower rate than within the populations evolving for locomotion over flat ground (black line in Fig. 10) or cluttered terrains (blue line in Fig. 10). This indicates that,

despite the fact that there is an inherent bias in this particular evolutionary simulation toward greater complexity over time, the greater complexity observed in the robots evolved to move through the more challenging environment is a result of selection and not random chance. In a similar manner, it is often difficult to determine which of these two forces produced a given form of complexity observed in nature. Thus, much work remains to understand the links between morphological, neural, and environmental complexity in biological organisms, as well as for machines. An understanding of such relationships would enable us to build both the most capable machine and the most simple machine for a given task environment.

VII. SELF-REFLECTION: FROM BODY TO MIND

In Section III, we described how a robotic system can become more resilient by acquiring an ability to create a forward model of its own body. That forward model can predict the mechanical consequence of various motor commands. The ability to create such a *self-model*, or *self-image*, allows the robot to try out various candidate actions on the model before carrying them out in physical reality, thereby eliminating unnecessary risk and energy expenditure. A robot or an animal can then use this ability to quickly optimize actions through virtual trial and error, to detect when something is wrong by noticing that predictions do not match reality, and to adapt to unforeseen situations.

From an evolutionary perspective, an innate ability of an individual to model its body could incur an advantage to survival under a changing environment where experimentation is costly or risky, especially under a long life span where the body or the environment is likely to change. But this advantage comes at a price. First, the self-modeling process can go wrong, leading to pathologies such as “phantom limbs.” But more importantly, the self-modeling machinery itself requires sophisticated learning and adaptation processes. Would there be other uses for such capability to justify the costs?

Just as the self-modeling capacity can help individuals to create models of their own body, the same self-modeling capability can be used to allow systems to model their own *behavior*. While body and brain are often anatomically distinct, from a purely algorithmic perspective, both the body and the brain can be seen, at a high level, as self-contained input–output (I/O) systems. The body has motor actuators as outputs and mechanical sensors as inputs; the brain issues output command signals and receives input signals. Modeling the brain should, in principle, not be fundamentally different than modeling the body.

It is quite possible that somewhere in our evolutionary heritage, the same mental processes that allowed some

individuals to model their own mechanical body eventually were used also to allow those individuals to model their own behavior. Imagine, for example, an individual that can predict how they would behave or react in a future situation, without actually being in that situation. They could then predict what situation would be more favorable, and take the action to engender that situation.

The ability to model one's own thinking, or "think about thinking," is a metacognitive capacity that is one step away from forward modeling of one's own mechanical body. In robotic terms, it is in fact a dynamical system-identification task where a dynamical model is created from observations. Often the term "self-reflection" is used to refer to modeling one's own behavior, in contrast to the term "self-modeling" used to describe the process of modeling one's own body. Fundamentally, however, these two processes are very similar learning processes, with very similar benefits and similar implementation challenges.

To implement self-reflection, we need to think of a robot, or an agent, as having two "minds"—one mind is being modeled, and the other mind is doing the modeling. Evidence has been presented that suggests that the human brain is indeed composed of multiple "mind modules" [44], [57], where some are capable of modeling others but not *vice versa*. For example, many animal brains contain a hippocampus, responsible for some spatial memory aspects. Humans and some other animals also have a cerebral cortex, which could be seen as a separate "mind" responsible for learning at a higher level. The cerebral cortex can think about the hippocampus (as we are doing right now while reading this text), but not the other way round. Minsky initially proposed that a brain could actually be composed of multiple subbrains, each serving as the external world for the next brain in a cascade [Fig. 11(a)].

We recently attempted to implement a variant of Minsky's architecture in a two-mind robot. Imagine a robotic system containing only two "brains." Let us think of one brain as observing the second brain, and the second brain as being observed by the first brain [Fig. 11(b)]. The observed mind is a simple, low-level brain that processes input signals and produces output signals according to some fixed behavior. The simple brain is evolved and optimized for some basic survival task, such as chasing blue dots and avoiding red dots using simple sensors (red/blue, left/right) and simple actuators (left/right wheels) [Fig. 11(d)]. The higher level brain does not have access to the inner workings of the lower brain, but the high-level brain can observe the lower level brain's I/O in various situations—i.e., the signals it receives and the signals it generates. This is conceptually similar to the way that you can consider your own spatial-memory behavior in various actual and hypothetical situations, even though you cannot actually access the inner workings of your own hippocampus.

What would be the advantages of this two-mind architecture? The advantages would be similar to the benefits gained from the ability to model one's own

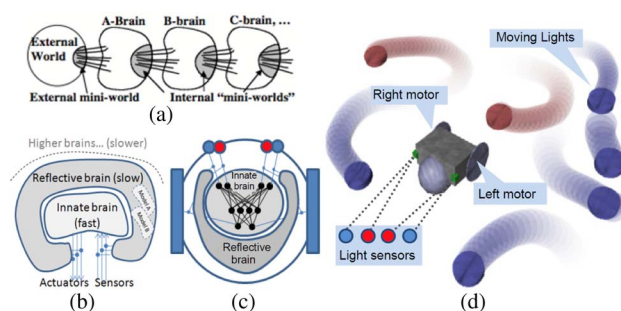


Fig. 11. Self-modeling cognitive architecture. (a) Minsky's brain chain, from [60]. (b) Proposal of nested brains architecture, from [92]. (c) Robot contains an innate brain, and a reflective brain which can model and modulate the I/O of the innate brain. (d) Schematic of robot and environment consisting of moving sources of blue and red light. (Reproduced from [60] and [92].)

mechanics. An individual with introspective self-reflection ability would be able to predict and optimize its own behavior in various situations without incurring the costs and risks of experiencing them out in physical reality. A self-reflective agent would be able to detect that something has changed in the lower level brain, or that the environment or tasks have changed to the point where the lower level brain's behavior is not optimal anymore and needs to be adapted. The higher level brain might even intervene and modulate the I/O behavior so that the lower level brain's behavior is useful in a new context.

To try this out, Zagal and Lipson [92] constructed a simple robot with two brains. The lower level brain was a simple neural network evolved to seek blue lights and avoid red lights, by turning the left and right wheels of a wheeled robot. The second onboard brain could only monitor the inputs and outputs to the lower brain. The higher level brain then formed a model of the lower brain by collecting data on the lower brain's I/O signals. The higher level brain could then detect anomalies in the behavior (e.g., due to damage in the lower level brain) and often recover at least partially by modulating the inputs and the outputs. More importantly, when the environment changed and now red lights needed to be chased rather than be avoided, the higher level brain could modulate the inputs so that the robot would still function appropriately. In other words, self-reflection allowed the robot to adapt to major internal and external changes.

One could challenge this thinking by asking why a robotic system would ever need to model its own behavior, since that behavior is known *a priori* to the engineers who designed the robot in the first place. While this argument is correct for most conventional robots today, it is likely that, in the future, robotic systems would not be so easily understood by their designers. Future robotic minds are likely to be generated indirectly through machine learning processes; they are likely to be increasingly complex and

opaque; and they might need to operate in changing environments where the original engineers are not available.

Taking the idea of self-reflection one step further, it is likely that individuals that evolved the capacity to model their own body and brain, were eventually able to model the body and brain of other agents as well. For example, some primates are able to model and anticipate the actions and behavior of other primates purely by observation. This ability, known as Theory of Mind [50], underlies much of human–human interaction and social behavior, and is key to cooperation, competition, deception, and manipulation observed in many social environments. While these concepts are very difficult to test and validate in human and animal behavior, they can be induced in simple robotic systems, offering a new window to this complex yet important aspect of cognition and self-awareness.

VIII. CONCLUSION

Much work in biomimetic computation is dedicated to creating detailed neural models in order to understand the function of biological nervous systems. Some of this work connects models to robots to determine how the model affects the machine’s ability to behave adaptively. However, with the exception of the projects described herein, few studies have investigated the interaction between morphological and neurobiological change. This change may occur at different time scales. We described work in which a robot’s morphology may change over the course of optimization (Sections VI and V), or, it may change over the course of the evaluation of a single neural controller due to growth (Section V) or damage (Section III).

Even if a robot’s morphology does not change, its conception of its own body—or those of others—may. We

described a machine that can spontaneously optimize a self-image (Section III) and use that to internally rehearse behaviors before attempting them in reality. Similarly, a team of robots may collectively construct self-models and share them to accelerate the process (Section III-C). Social robots may construct models of self and models of others in order to learn from them (Section III-D) or predict their intentions (Section VII).

Embodied cognition, coupled with an evolutionary approach to the improvement of autonomous machines, provides a framework for incorporating biological inspiration at different temporal and spatial scales. Temporally, change may occur within populations over evolutionary time; individual robots may develop from infant into adult forms; learning may induce neurobiological change; and models of self and others may form as the robot interacts with its environment as well as other machines. Spatially, evolutionary methods allow for local or coordinated change to the robot’s body as well as its brain; learning methods are constrained to improve the robot’s control policy.

Finally, we may investigate feedback loops between these different spatial and temporal processes. The studies reported here all employ neural models that are much simpler than those used elsewhere. For instance, in all of the work described here, no learning occurs: synaptic connections between neurons are fixed by the evolutionary algorithm. Future work is needed to incorporate more detailed neural models into evolving machine populations. This would allow us to explore how short-term change, such as learning, is affected by change not only at longer time scales, but also by change across the machine’s various mechanical, material, motoric, sensory, and control subsystems. ■

REFERENCES

- [1] M. Asada, K. Hosoda, Y. Kuniyoshi, H. Ishiguro, T. Inui, Y. Yoshikawa, M. Ogino, and C. Yoshida, “Cognitive developmental robotics: A survey,” *IEEE Trans. Autom. Mental Develop.*, vol. 1, no. 1, pp. 12–34, May 2009.
- [2] J. E. Auerbach and J. C. Bongard, “Environmental influence on the evolution of morphological complexity in machines,” *PLoS Comput. Biol.*, vol. 10, no. 1, 2014, e1003399.
- [3] H. C. Barrett and R. Kurzban, “Modularity in cognition: Framing the debate,” *Psychol. Rev.*, vol. 113, no. 3, pp. 628–647, 2006.
- [4] R. D. Beer, “Parameter space structure of continuous-time recurrent neural networks,” *Neural Comput.*, vol. 18, no. 12, pp. 3009–3051, 2006.
- [5] N. Bernstein, *The Coordination and Regulation of Movements*. Oxford, U.K.: Pergamon Press, 1967.
- [6] L. Berthouze and M. Lungarella, “Motor skill acquisition under environmental perturbations: On the necessity of alternate freezing and freeing of degrees of freedom,” *Adapt. Behav.*, vol. 12, no. 1, pp. 47–64, 2004.
- [7] J. Bongard, V. Zykov, and H. Lipson, “Resilient machines through continuous self-modeling,” *Science*, vol. 314, pp. 1118–1121, 2006.
- [8] J. C. Bongard, “The utility of evolving simulated robot morphology increases with task complexity for object manipulation,” *Artif. Life*, 2010, DOI: 10.1162/artl.2010. Bongard.024.
- [9] J. C. Bongard, “Spontaneous evolution of structural modularity in robot neural network controllers,” in *Proc. Gen. Evol. Comput. Conf.*, 2011, pp. 251–258.
- [10] J. C. Bongard, “Accelerating self-modeling in cooperative robot teams,” *IEEE Trans. Evol. Comput.*, vol. 13, no. 2, pp. 321–332, Apr. 2009.
- [11] J. Bongard, “Morphological change in machines accelerates the evolution of robust behavior,” *Proc. Nat. Acad. Sci. USA*, vol. 108, no. 4, pp. 1234–1239, 2011.
- [12] R. A. Brooks, “Intelligence without representation,” *Artif. Intell.*, vol. 47, pp. 139–160, 1991.
- [13] E. Bullmore and O. Sporns, “Complex brain networks: Graph theoretical analysis of structural and functional systems,” *Nature Rev. Neurosci.*, vol. 10, no. 3, pp. 186–198, 2009.
- [14] W. H. Calvin, *The Cerebral Code: Thinking a Thought in the Mosaics of the Mind*. Cambridge, MA, USA: MIT Press, 2000.
- [15] S. B. Carroll, “Chance and necessity: The evolution of morphological complexity and diversity,” *Nature*, vol. 409, no. 6823, pp. 1102–1109, 2001.
- [16] A. Clark, *Being There: Putting Brain, Body, World Together Again*. Cambridge, MA, USA: Bradford Books, 1998.
- [17] A. Clark, “Whatever next? Predictive brains, situated agents, the future of cognitive science,” *Behav. Brain Sci.*, vol. 36, no. 3, pp. 181–204, 2013.
- [18] D. Cliff, P. Husbands, and I. Harvey, “Explorations in evolutionary robotics,” *Adapt. Behav.*, vol. 2, no. 1, pp. 73–110, 1993.
- [19] J. Clune, J.-B. Mouret, and H. Lipson, “The evolutionary origins of modularity,” *Proc. Roy. Soc. B, Biol. Sci.*, vol. 280, no. 1755, 2013, DOI: 10.1098/rspb.2012.2863.
- [20] S. Collins, A. Ruina, R. Tedrake, and M. Wisse, “Efficient bipedal robots based on passive-dynamic walkers,” *Science*, vol. 307, no. 5712, pp. 1082–1085, 2005.

- [21] C. Consoulas, C. Duch, R. J. Bayliss, and R. B. Levine, "Behavioral transformations during metamorphosis: Remodeling of neural and motor systems," *Brain Res. Bull.*, vol. 53, no. 5, pp. 571–583, 2000.
- [22] F. Crick, "Neural edelmanism," *Trends Neurosci.*, vol. 12, no. 7, pp. 240–248, 1989.
- [23] K. A. De Jong, *Evolutionary Computation: A Unified Approach*. Cambridge, MA, USA: MIT Press, 2006.
- [24] F. De Vignemont, "Body schema and body image—Pros and cons," *Neuropsychologia*, vol. 48, no. 3, pp. 669–680, 2010.
- [25] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, Apr. 2002.
- [26] J. Diedrichsen, R. Shadmehr, and R. B. Ivry, "The coordination of movement: Optimal feedback control and beyond," *Trends Cogn. Sci.*, vol. 14, no. 1, pp. 31–39, 2010.
- [27] A. G. Dimitrov, A. A. Lazar, and J. D. Victor, "Information theory in neuroscience," *J. Comput. Neurosci.*, vol. 30, no. 1, pp. 1–5, 2011, DOI: 10.1007/s10827-011-0314-3.
- [28] M. Dorigo and M. Colombetti, "Robot shaping: Developing situated agents through learning," *Artif. Intell.*, vol. 70, no. 2, pp. 321–370, 1994.
- [29] G. M. Edelman, *Neural Darwinism: The Theory of Neuronal Group Selection*. New York, NY, USA: Basic Books, 1987.
- [30] C. Espinosa-Soto and A. Wagner, "Specialization can drive the evolution of modularity," *PLoS Comput. Biol.*, vol. 6, no. 3, 2010, e1000719.
- [31] C. Fernando, R. Goldstein, and E. Szathmáry, "The neuronal replicator hypothesis," *Neural Comput.*, vol. 22, no. 11, pp. 2809–2857, 2010.
- [32] D. Floreano and L. Keller, "Evolution of adaptive behaviour in robots by means of Darwinian selection," *PLoS Biol.*, vol. 8, no. 1, 2010, e1000292.
- [33] N. Franceschini, "Small brains, smart machines: From insects to robots, back again," *Proc. IEEE*, vol. 102, no. 5, May 2014, DOI: 10.1109/JPROC.2014.2312916.
- [34] M. Freeling and B. C. Thomas, "Gene-balanced duplications, like tetraploidy, provide predictable drive to increase morphological complexity," *Genome Res.*, vol. 16, no. 7, pp. 805–814, 2006.
- [35] K. Friston, "The free-energy principle: A unified brain theory?" *Nature Rev. Neurosci.*, vol. 11, no. 2, pp. 127–138, 2010.
- [36] V. Gallese and A. Goldman, "Mirror neurons and the simulation theory of mind-reading," *Trends Cogn. Sci.*, vol. 2, no. 12, pp. 493–501, 1998.
- [37] D. George and J. Hawkins, "Towards a mathematical theory of cortical micro-circuits," *PLoS Comput. Biol.*, vol. 5, no. 10, 2009, e1000532.
- [38] G. Gómez, M. Lungarella, P. E. Hotsz, K. Matsushita, and R. Pfeifer, "Simulating development in a real robot: On the concurrent increase of sensory, motor, neural complexity," in *Proc. 3rd Int. Conf. Develop. Learn.*, 2004, pp. 119–122.
- [39] M. A. Goodrich and A. C. Schultz, "Human-robot interaction: A survey," *Found. Trends Human-Comput. Interaction*, vol. 1, no. 3, pp. 203–275, 2007.
- [40] A. Gopnik and L. Schulz, "Mechanisms of theory formation in young children," *Trends Cogn. Sci.*, vol. 8, no. 8, pp. 371–377, 2004.
- [41] G. E. Hinton, S. Osindero, and Y. Teh, "A fast learning algorithm for deep belief nets," *Neural Comput.*, vol. 18, pp. 1527–1554, 2006.
- [42] M. Hoffmann, H. Marques, A. Arieta, H. Sumioka, M. Lungarella, and R. Pfeifer, "Body schema in robotics: A review," *IEEE Trans. Autom. Mental Develop.*, vol. 2, no. 4, pp. 304–324, Dec. 2010.
- [43] M. Iacoboni, "Imitation, empathy, mirror neurons," *Annu. Rev. Psychol.*, vol. 60, pp. 653–670, 2009.
- [44] H. Imamizu, T. Kuroda, S. Miyauchi, T. Yoshioka, and M. Kawato, "Modular organization of internal models of tools in the human cerebellum," *Proc. Nat. Acad. Sci. USA*, vol. 100, no. 9, pp. 5461–5466, 2003.
- [45] B. Irie and S. Miyake, "Capabilities of three-layered perceptrons," in *Proc. IEEE Int. Conf. Neural Netw.*, 1988, pp. 641–648.
- [46] N. Jakobi, P. Husbands, and I. Harvey, "Noise and the reality gap: The use of simulation in evolutionary robotics," in *Advances in Artificial Life*. Berlin, Germany: Springer-Verlag, 1995, pp. 704–720.
- [47] T. D. Jorgensen, B. P. Haynes, and C. C. F. Norlund, "Pruning artificial neural networks using neural complexity measures," *Int. J. Neural Syst.*, vol. 18, no. 5, pp. 389–403, 2008.
- [48] T. Kanda, D. F. Glas, M. Shiomi, H. Ishiguro, and N. Hagita, "Who will be the customer?: A social robot that anticipates people's behavior from their trajectories," in *Proc. 10th Int. Conf. Ubiquitous Comput.*, 2008, pp. 380–389.
- [49] N. Kashtan and U. Alon, "Spontaneous evolution of modularity and network motifs," *Proc. Nat. Acad. Sci. USA*, vol. 102, no. 39, pp. 13773–13778, 2005.
- [50] K.-J. Kim and H. Lipson, "Towards a theory of mind in simulated robots," in *Proc. 11th Annu. Conf. Companion Gen. Evol. Comput. Conf.*, 2009, pp. 2071–2076.
- [51] K. N. Krishnanand, J. C. Bongard, and A. N. Meltzoff, "Self discovery enables robot social cognition: Are you my teacher?" *Neural Netw.*, vol. 23, pp. 1113–1124, 2010.
- [52] H. Lipson, J. B. Pollack, and N. P. Suh, "On the origin of modular variation," *Evolution*, vol. 56, no. 8, pp. 1549–1556, 2002.
- [53] W. Maass, "Noise as a resource for computation and learning in networks of spiking neurons," *Proc. IEEE*, vol. 102, no. 5, May 2014, DOI: 10.1109/JPROC.2014.2310593.
- [54] G. Massera, A. Cangelosi, and S. Nolfi, "Evolution of prehension ability in an anthropomorphic neurobotic arm," *Front. Neurobot.*, vol. 1, 2007, DOI: 10.3389/neuro.12.004.2007.
- [55] D. W. McShea and R. N. Brandon, *Biology's First Law: The Tendency for Diversity and Complexity to Increase in Evolutionary Systems*. Chicago, IL, USA: Univ. Chicago Press, 2010.
- [56] A. N. Meltzoff, "'Like me': A foundation for social cognition," *Develop. Sci.*, vol. 10, no. 1, pp. 126–134, 2007.
- [57] D. Meunier, R. Lambiotte, A. Fornito, K. D. Ersche, and E. T. Bullmore, "Hierarchical modularity in human brain functional networks," *Front. Neuroinf.*, vol. 3, 2009, DOI: 10.3389/neuro.11.037.2009.
- [58] R. C. Miall, "Connecting mirror neurons and forward models," *NeuroReport*, vol. 14, no. 16, pp. 1–3, 2003.
- [59] T. Miconi, "Evolution and complexity: The double-edged sword," *Artif. Life*, vol. 14, no. 3, pp. 325–344, 2008.
- [60] M. Minsky, *Society of Mind*. New York, NY, USA: Simon & Schuster, 1988.
- [61] S. Murata, E. Yoshida, A. Kamimura, H. Kurokawa, K. Tomita, and S. Kokaji, "M-tran: Self-reconfigurable modular robotic system," *IEEE/ASME Trans. Mechatron.*, vol. 7, no. 4, pp. 431–441, Dec. 2002.
- [62] C. L. Nehaniv and K. Dautenhahn, *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge, U.K.: Cambridge Univ. Press, 2007.
- [63] D. L. Page, A. F. Koschan, S. R. Sukumar, B. Roui-Abidi, and M. A. Abidi, "Shape analysis algorithm based on information theory," in *Proc. IEEE Int. Conf. Image Process.*, 2003, vol. 1, pp. 229–232.
- [64] S. I. Passy, "Environmental randomness underlies morphological complexity of colonial diatoms," *Funct. Ecol.*, vol. 16, no. 5, pp. 690–695, 2002.
- [65] R. Pfeifer and J. Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA, USA: MIT Press, 2006.
- [66] F. Pulvermüller and L. Fadiga, "Active perception: Sensorimotor circuits as a cortical basis for language," *Nature Rev. Neurosci.*, vol. 11, no. 5, pp. 351–360, 2010.
- [67] R. Q. Quiroga and S. Panzeri, "Extracting information from neuronal populations: Information theory and decoding approaches," *Nature Rev. Neurosci.*, vol. 10, no. 3, pp. 173–185, 2009.
- [68] T. Reil and P. Husbands, "Evolution of central pattern generators for bipedal walking in a real-time physics environment," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 159–168, Apr. 2002.
- [69] R. Saegusa, G. Metta, G. Sandini, and S. Sakka, "Active motor babbling for sensorimotor learning," in *Proc. IEEE Int. Conf. Robot. Biomimetics*, 2009, pp. 794–799.
- [70] M. Schmidt and H. Lipson, "Age-fitness Pareto optimization," in *Genetic Programming Theory and Practice VIII*. New York, NY, USA: Springer-Verlag, 2011, pp. 129–146.
- [71] A. K. Seth and G. M. Edelman, "Environment and behavior influence the complexity of evolved neural networks," *Adapt. Behav.*, vol. 12, no. 1, pp. 5–20, 2004.
- [72] H. S. Seung, M. Opper, and H. Sompolinsky, "Query by committee," in *Proc. 5th Workshop Comput. Learn. Theory*, 1992, pp. 287–294.
- [73] R. F. Shepherd, F. Ilievski, W. Choi, S. A. Morin, A. A. Stokes, A. D. Mazzeo, X. Chen, M. Wang, and G. M. Whitesides, "Multigait soft robot," *Proc. Nat. Acad. Sci. USA*, vol. 108, no. 51, pp. 20400–20403, 2011.
- [74] K. Sims, "Evolving 3D morphology and behaviour by competition," *Artif. Life*, vol. 4, pp. 28–39, 1994.
- [75] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evol. Comput.*, vol. 10, no. 2, pp. 99–127, 2002.
- [76] K. O. Stanley, "Compositional pattern producing networks: A novel abstraction of development," *Genetic Programm. Evolvable Mach.*, vol. 8, no. 2, pp. 131–162, 2007.
- [77] S. R. Sukumar, D. L. Page, A. F. Koschan, and M. A. Abidi, "Towards understanding what makes 3D objects appear simple or complex," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2008, DOI: 10.1109/CVPRW.2008.4562975.
- [78] R. Tedrake, T. Zhang, and H. Seung, "Learning to walk in 20 minutes," in *Proc. 14th Yale Workshop Adapt. Learn. Syst.*, pp. 1–6, 2005.

- [79] E. Todorov and M. I. Jordan, "Optimal feedback control as a theory of motor coordination," *Nature Neurosci.*, vol. 5, no. 11, pp. 1226–1235, 2002.
- [80] G. Tononi, O. Sporns, and G. M. Edelman, "A measure for brain complexity: Relating functional segregation and integration in the nervous system," *Proc. Nat. Acad. Sci. USA*, vol. 91, no. 11, pp. 5033–5037, 1994.
- [81] L. G. Valiant, "Evolvability," *J. ACM*, vol. 56, no. 1, 2009, DOI: 10.1145/1462153.1462156.
- [82] E. Vaughan, E. A. Di Paolo, and I. Harvey, "The evolution of control and adaptation in a 3D powered passive dynamic walker," in *Proc. 9th Int. Conf. Simul. Synthesis Living Syst.*, 2004, pp. 139–145.
- [83] G. P. Wagner and L. Altenberg, "Perspective: Complex adaptations and the evolution of evolvability," *Evolution*, vol. 50, no. 3, pp. 967–976, 1996.
- [84] G. P. Wagner, M. Pavlicev, and J. M. Cheverud, "The road to modularity," *Nature Rev. Gen.*, vol. 8, no. 12, pp. 921–931, 2007.
- [85] D. M. Wolpert, J. Diedrichsen, and J. R. Flanagan, "Principles of sensorimotor learning," *Nature Rev. Neurosci.*, vol. 12, no. 12, pp. 739–751, 2011.
- [86] D. M. Wolpert and M. Kawato, "Multiple paired forward and inverse models for motor control," *Neural Netw.*, vol. 11, no. 7, pp. 1317–1329, 1998.
- [87] D. Wood, J. S. Bruner, and G. Ross, "The role of tutoring in problem solving," *J. Child Psychol. Psychiatry*, vol. 17, no. 2, pp. 89–100, 1976.
- [88] S. G. Wysoski, L. Benuskova, and N. Kasabov, "Evolving spiking neural networks for audiovisual information processing," *Neural Netw.*, vol. 23, no. 7, pp. 819–835, 2010.
- [89] Y. Yamashita and J. Tani, "Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment," *PLoS Comput. Biol.*, vol. 4, no. 11, 2008, e1000220.
- [90] M. Yim, W.-M. Shen, B. Salemi, D. Rus, M. Moll, H. Lipson, E. Klavins, and G. S. Chirikjian, "Modular self-reconfigurable robot systems [grand challenges of robotics]," *IEEE Robot. Autom. Mag.*, vol. 14, no. 1, pp. 43–52, Mar. 2007.
- [91] M. Yim, Y. Zhang, and D. Duff, "Modular robots," *IEEE Spectrum*, vol. 39, no. 2, pp. 30–34, 2002.
- [92] J. C. Zagal and H. Lipson, "Towards self-reflecting machines: Two-minds in one robot," in *Advances in Artificial Life. Darwin Meets von Neumann*. New York, NY, USA: Springer-Verlag, 2011, pp. 156–164.

ABOUT THE AUTHORS

Josh Bongard (Member, IEEE) received the B.Sc. degree in computer science from McMaster University, Hamilton, ON, Canada, in 1997, the M.Sc. degree in evolutionary and adaptive systems from the University of Sussex, Brighton, Sussex, U.K., in 1999, and the Ph.D. degree in computer science from the University of Zurich, Zurich, Switzerland, in 2003.

He served as a Postdoctoral Associate at Cornell University, Ithaca, NY, USA. He is currently an Associate Professor at the Computer Science Department, University of Vermont, Burlington, VT, USA.

Dr. Bongard received the Presidential Early Career Award for Scientists and Engineers (PECASE) in 2011. In 2006, he was named a Microsoft New Faculty Fellow, as well as one of the top 35 innovators under the age of 35 by MIT's *Technology Review Magazine*.



Hod Lipson (Member, IEEE) received the B.Sc. degree in mechanical engineering and the Ph.D. degree in mechanical engineering in computer-aided design and artificial intelligence in design from the Technion—Israel Institute of Technology, Haifa, Israel, in 1989 and 1998, respectively.

He is currently an Associate Professor with the Mechanical and Aerospace Engineering and Computing and Information Science Schools, Cornell University, Ithaca, NY, USA. He was a Postdoctoral Researcher with the Department of Computer Science, Brandeis University, Waltham, MA, USA. He was a Lecturer with the Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA, USA, where he was engaged in conducting research in design automation. His current research interests include computational methods to synthesize complex systems out of elementary building blocks and the application of such methods to design automation and their implication toward understanding the evolution of complexity in nature and in engineering.

