



Morphological change in machines accelerates the evolution of robust behavior

Josh Bongard¹

University of Vermont, Burlington, VT 05405

Edited by Terrence J. Sejnowski, Salk Institute for Biological Studies, La Jolla, CA, and approved December 1, 2010 (received for review October 14, 2010)

Most animals exhibit significant neurological and morphological change throughout their lifetime. No robots to date, however, grow new morphological structure while behaving. This is due to technological limitations but also because it is unclear that morphological change provides a benefit to the acquisition of robust behavior in machines. Here I show that in evolving populations of simulated robots, if robots grow from anguilliform into legged robots during their lifetime in the early stages of evolution, and the anguilliform body plan is gradually lost during later stages of evolution, gaits are evolved for the final, legged form of the robot more rapidly—and the evolved gaits are more robust—compared to evolving populations of legged robots that do not transition through the anguilliform body plan. This suggests that morphological change, as well as the evolution of development, are two important processes that improve the automatic generation of robust behaviors for machines. It also provides an experimental platform for investigating the relationship between the evolution of development and robust behavior in biological organisms.

evolutionary computation | robotics | evolutionary robotics | biomechanics | locomotion

Using robots to study adaptive behavior is of interest as a basic intellectual pursuit, but it may also lead to machines that assist or replace humans in unstructured or outdoor environments. To date, however, limited success has been achieved in realizing machines that continually perform simple yet robust behaviors in unstructured environments. It is contended here that this is due to overemphasis on the proximate mechanisms (1) of adaptive behavior—copying specific morphological and neuro-morphological detail from organisms of interest into robots (2) in the hopes of replicating their behavior—and too little emphasis on the ultimate mechanisms of behavior—replicating the ontogenetic processes and selection pressures that gave rise to the behavior initially.

To demonstrate the value of incorporating the evolution of development (3) into robotics, here I show that the automated discovery of robot behaviors can be accelerated if their body plans progress from an infant to an adult form over their lifetime (Fig. 1*B*)—and this ontogenetic change itself changes over successive generations of robots (Fig. 1*D* and *F*) such that later robots exhibit only the adult form during their lifetime (Fig. 1*H*)—than if robots maintain the adult form throughout behavior optimization. Moreover, robots evolved with this evolution-of-development method were found to be more robust to environmental perturbation than robots evolved without it because the former robots' ancestors assumed a range of body plans and thus had to maintain the behavior over a wider range of sensor-motor relationships.

The method introduced here follows a biologically inspired approach to robotics in which populations of simulated robots undergo evolutionary pressure favoring some desired behavior (4, 5). In some approaches, only the robot's controller is optimized (6, 7) after which the best controller can be downloaded onto a physical robot (8). Alternatively, the robot's body plan can be evolved as well (9–12), after which the best body plan can be automatically manufactured (10). However, there are only two cases reported in the literature (13, 14) in which a simulated

robot's body was allowed to change while it was behaving. It was not clear though whether this ontogenetic morphological change facilitated the evolution of behavior.

Evolving Behavior

The system introduced here allows a robot's body plan to change while it behaves: Robots progress from anguilliform body plans (15) into upright, legged body plans while behaving in a simulated environment (Figs. 1*B* and Fig. 2*N*). Moreover, this morphological change itself changes over phylogenetic time such that later robots exhibit only the legged form.

This is accomplished by starting with an initial population of random neural-network controllers. Each robot is evaluated for a fixed time period using one of the controllers, and its ability to perform phototaxis is measured. The next generation is generated by probabilistically discarding controllers that produced poor behavior, creating randomly perturbed copies of the remaining controllers, and evaluating the new controllers. Generations elapse in this way until a controller evolves that successfully moves the robot to a light source placed in its environment during a fixed time period (Fig. 1*B*). Once this successful controller is found, each controller is reevaluated on a robot with an altered ontogeny: Morphological change is accelerated such that the robot begins with the anguilliform body plan and progresses to the adult, upright form over the first two-thirds of its evaluation period (Fig. 2*O*). Over the remaining third of the evaluation period it maintains the upright form. Usually, this degrades the behavior of the current controllers, so evolution continues until successful phototaxis is rediscovered (Fig. 1*D*).

Ontogeny is again altered such that robots now reach the upright form after the first third of the evaluation period (Fig. 2*P*). Evolution again proceeds until phototaxis is rediscovered (Fig. 1*F*). In the fourth and final phase of evolution, each robot begins with the upright, adult form and maintains it throughout its evaluation in the simulated environment (Fig. 2*Q*). The experiment ends when an upright robot capable of phototaxis is discovered (Fig. 1*H*).

Impact of Morphological Change on Behavior

This process was found to produce successfully phototacting robots with the upright stance more rapidly, and the resulting behaviors were more robust, compared to four baseline algorithm variants in which morphological change is lessened or disabled. In all five variants, it is assumed that the final goal of the system is the automated discovery of a controller for an upright legged robot such that it locomotes toward a light source placed in its simulated environment (Fig. 1*H*). Thus, in all five algorithm variants, in the final generations controllers are evolved on robots that maintain the upright, legged form throughout behavior evaluation (Fig. 2*A, E, I, M, and Q*). Once such a controller is found

Author contributions: J.B. designed research, performed research, analyzed data, and wrote the paper.

The author declares no conflict of interest.

This article is a PNAS Direct Submission.

¹E-mail: josh.bongard@uvm.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1015390108/-DCSupplemental.

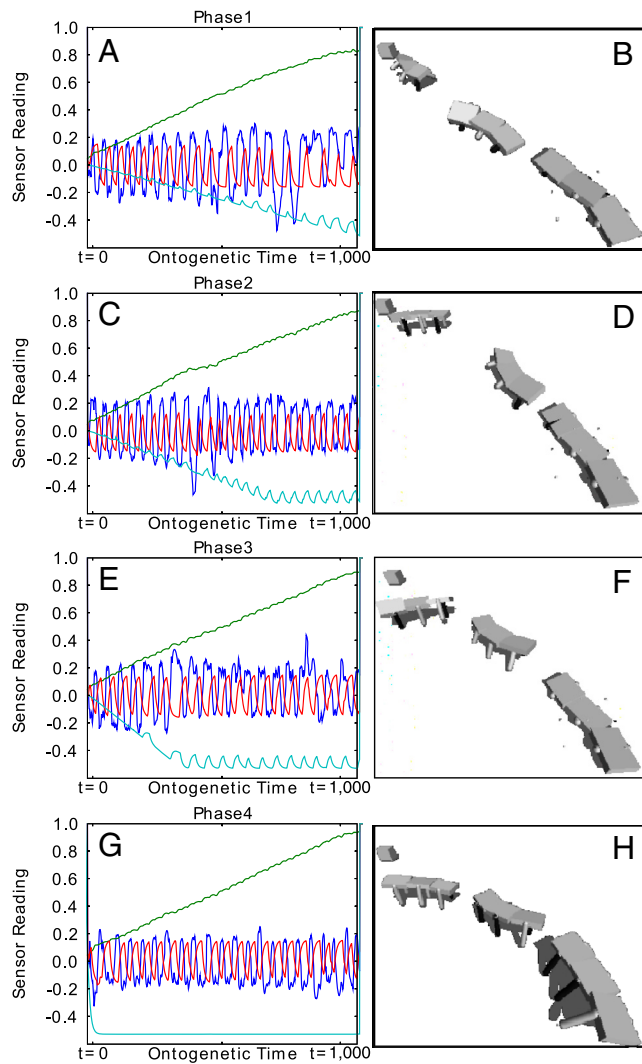


Fig. 1. Results from a typical experiment. (A) Sensor signals obtained from the first robot to evolve that successfully performs phototaxis in the first phase of evolution. (The blue line reports signals from one of two vestibular sensors in the robot's front segment, the green line reports signals from one of three photoreceptors located in the front segment, the red line reports signals from one of four proprioceptive sensors in the spine, and the light blue line reports signals from one of six proprioceptive sensors at the base of each leg.) The robot's behavior and morphological change is shown in B: It begins in the legless form (*Lower Right*); as it locomotes, it grows limbs (*Center*) and finally reaches the target object at the end of its "lifetime," at which point it has assumed the fully upright, legged form (*Upper Left*). The sensor signals (C) and behavior (D) of the successful robot found during the second phase: this robot reaches the fully upright, legged form two-thirds of the way through its lifetime. The sensor signals (E) and behavior (F) of the successful robot found during the third phase: this robot reaches the adult form one-third of the way through its lifetime. The sensor signals (G) and behavior (H) of the successful robot found during the fourth phase: this robot never assumes the infant, legless form; it maintains the legged form throughout evaluation.

in any of these trials, it can be downloaded onto a physical robot with an upright quadrupedal form as shown in Fig. 2 U–W.

No Morphological Change. In the first baseline variant, the robots evolve using the same controller architecture, objective function, and optimization algorithm as before, but now no robots experience any morphological change: Each robot begins with the upright legged form and maintains it until the end of its evaluation period (Fig. 2A). Evolution continues until a robot is discovered that successfully exhibits phototaxis.

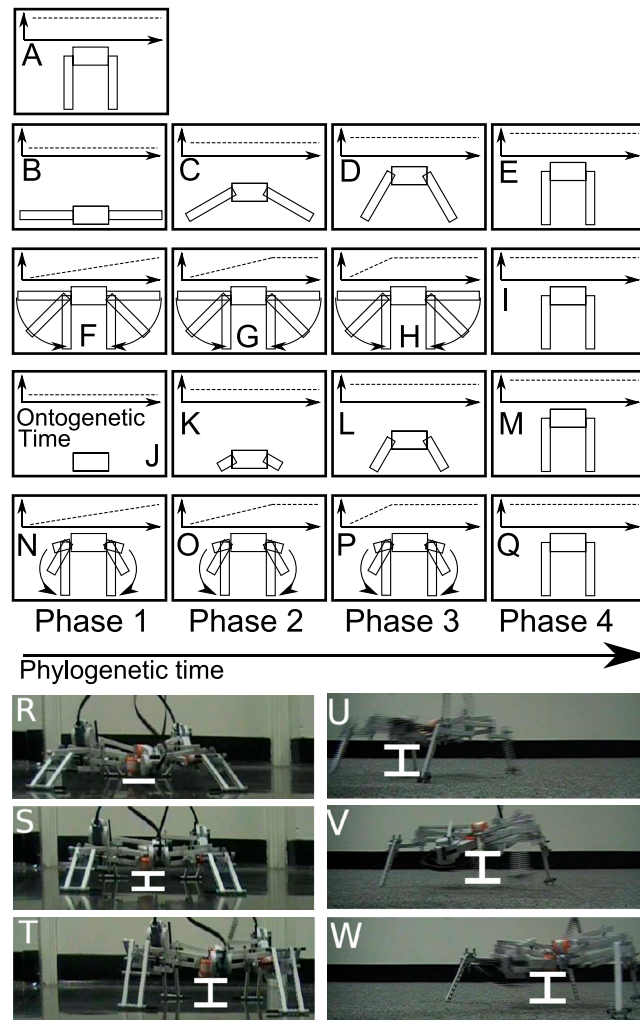


Fig. 2. Allowing morphological change during the evolution of robot behavior. (A) The first baseline algorithm variant in which all robots were evaluated in a fixed, upright stance. (*Inset*) Graph indicates how the height of the robot's center of mass changes during evaluation. (B–E) The second baseline variant, in which the robot's body plan changes parametrically from one evolutionary phase to the next but not while moving. (F–I) The third baseline variant in which ontogenetic morphological change is incorporated. (J–M) The fourth baseline variant in which anguilliform robots in the first phase evolve into upright legged robots in the final phase, but there is no morphological change during robot evaluation. (N–Q) The proposed method in which topological morphological change is combined with ontogenetic morphological change. (R–T) A physical robot was constructed that undergoes ontogenetic morphological change (as in F). This robot walks from left to right while transitioning from a prone to an upright stance. Once such a gait is found for the physical robot, the additional motors required to achieve ontogenetic morphological change can be removed, and the robot exhibits the same gait with a fixed, upright stance (U–W).

Parametric Change. In the second baseline variant, only parametric morphological change was allowed. Topological morphological change is defined here as the growth of legs from an initially anguilliform body plan (Fig. 1). Parametric morphological change is defined as the angle between the robot's legs and its trunk, but the number of body parts and motorized degrees of freedom remain constant. In this variant robots begin with legs, but they are aligned horizontally such that the trunk of the robot is in contact with the ground plane. The controllers in the robots may then actuate the trunk and sweep the legs through the sagittal plane, thereby mimicking a reptilian body plan (Fig. 2B). In the initial generations of this variant, robots maintain this body plan throughout the evaluation period. When a successful robot is dis-

covered, in the second evolutionary phase the legs are clamped to a descending 30° angle relative to the trunk, producing a splayed body plan (Fig. 2C). In the third phase the legs are clamped to a descending 60° angle (Fig. 2D); in the fourth phase the legs are vertical (Fig. 2E).

Ontogenetic Parametric Change. In the third baseline variant, topological body change is replaced with parametric body plan change as in the second baseline variant, but ontogenetic morphological change is allowed: In the first evolutionary phase, legged robots transition from the prone to the upright stance as they behave (Fig. 2F); in the second stage they progress from the prone to the upright stance over the first two-thirds of the evaluation period (Fig. 2G); in the third stage, over the first third of the evaluation period (Fig. 2H); and in the fourth stage, they begin and maintain the upright stance (Fig. 2I).

Phylogenetic Change. In the fourth baseline variant, the body plan of any robot does not change during evaluation of its behavior, but the body plan itself may change topologically over evolutionary time: Fixed-form anguilliform robots in early generations are incrementally replaced by fixed-form legged robots in later generations. In the first evolutionary phase robots maintain the anguilliform form throughout evaluation (Fig. 2J); in the second phase robots are equipped with legs one-third of the length of the legs in the adult form, the legs are clamped at a descending 30° angle relative to the trunk, and the robot maintains that form throughout its evaluation (Fig. 2K); in the third phase robots are equipped with legs two-thirds the length of the adult legs, which are clamped at a descending 60° angle (Fig. 2L); and in the fourth phase the robots maintain the adult, upright legged form (Fig. 2M).

Results

A total of $2 \times 5 \times 5 \times 100 = 5,000$ independent evolutionary trials were conducted. The effect of morphological change on the evolution of behavior was measured across two robots, five amounts of morphological change, and five task environments.

Trials were conducted with a quadrupedal and hexapedal robot. These two forms were used as they are common choices when constructing physical robots (16, 17), and similarity of results across both forms would indicate the results are not specific to any given body plan. For each robot, five sets of trials were conducted using the experimental and four baseline algorithm variants. For each robot and algorithm variant, five sets of trials were conducted using increasingly challenging task environments. In the first task environment the light source was placed 20 m directly in front of the robot; in the second through fifth task environments the light source was placed 20 m away and $\pi/16$, $2\pi/16$, $3\pi/16$, and $4\pi/16$ radians to its front and left, respectively, thereby increasingly challenging robots to turn during movement while maintaining balance (Fig. 1B, D, F, and H). For each robot, algorithm variant, and task environment, 100 independent evolutionary trials were conducted. Each trial began with 400 random neural-network controllers and progressed until either an upright legged robot successfully performed phototaxis or 30 h of central processing unit time elapsed.

A Typical Experiment. Fig. 1 reports results from a typical experiment using the hexapedal robot when its body plan changed topologically and ontogenetically. Phototaxis was discovered in the first phase of evolution in which each robot began with an anguilliform body plan and gradually progressed to the legged body plan over the evaluation period (Fig. 1B). The sensor signatures for this successful behavior (Fig. 1A) indicate that the spine drives the gait (red line), which is to be expected because the robot has no legs for walking at the outset of the evaluation. However, as the legs grow, they contribute to locomotion (light blue line).

Later in the same evolutionary trial successful phototaxis is rediscovered for the robot, which progresses from the infant to the adult form over the two-thirds of the evaluation period (Fig. 1C and D). Much of the original gait is retained in the new gait as indicated by the similar motions of the body (Fig. 1A and C, red and blue lines). However, the proprioceptive sensor signature at one of the shoulders (Fig. 1A and C, light blue line) differs between the gaits as this sensor captures the initial growth of the legs (initial downward slope of light blue line) as well as its motion pattern once grown (oscillations in the line). This same pattern was observed when evolution transitioned from the second to the third phase and then rediscovered phototaxis in the new robot body plan (Fig. 1E and F). This trial terminated when the gait illustrated in Fig. 1G and H was discovered. Like the previous three gaits, the spine is actuated similarly, but some of the six legs do not move (no oscillation in the light blue line). This indicates that the initial undulatory gait employing the spine discovered initially has been adapted but not discarded during evolution, but the evolved motion primitives in the legs have changed.

Time to Successful Behavior Discovery. It was found that for both robots (Fig. 3), phototaxis becomes more difficult as the target object is displaced further from the robot's initial midline: The time to discovery of a controller capable of producing phototaxis in the robot increased along with the turning angle the robot must make while behaving.

It was also found that introducing morphological change during behavior evolution may increase or decrease this time to discovery: If robots progress over evolutionary time from the anguilliform to the legged body plan (Fig. 3B and D) but any one robot maintains the same body plan during the evaluation period (fourth baseline variant), the task becomes more difficult compared to only evolving robots with the upright legged form. This can be explained by the fact that when one evolutionary phase terminates and a new one begins, the evolved neural networks experience a sudden change in motor-sensor relationships: When the robot with the new body plan is actuated in the same way the previous robot was, different sensor signals are received. This leads to a large overall decrease in the fitness of the evolved controllers (Fig. S1), and many generations are required to reevolve phototaxis in the new body plan.

Impact of Ontogenetic Body Plan Change. However, when ontogenetic body plan change was added to this variant (the experimental variant), successful phototaxis was discovered significantly faster in the most challenging task environment for the quadruped (Fig. 3B), and in all five task environments for the hexapod (Fig. 3D) compared to the algorithm variant in which no morphological change was allowed. In this approach, there are no novel motor-sensor relationships experienced by controllers reevaluated on robots with accelerated body plan change. The controllers successfully maintained locomotion for all gradations between the infant and adult body plan in the previous phase and must do so for the robots in the new phase. There is still increased challenge in the new phase, however, as the controller must maintain the robot's balance for the upright legged form for a longer period of time than it had to in the previous phase. Thus there is a less drastic decrease in overall fitness in the controller population when transitioning to a new phase (Fig. S1) compared to when ontogenetic body plan change is disabled. This slight fitness disruption yields a shorter time period to rediscovery of phototaxis.

Impact of Topological Body Plan Change. In addition to ontogenetic body plan change, topological body plan change accelerates the time to discovery of successful behavior compared to parametric body plan change. This acceleration is slight yet not significant for the quadruped for the last two task environments (black bars in Fig. 3A and B). However, for the hexapedal robot topological

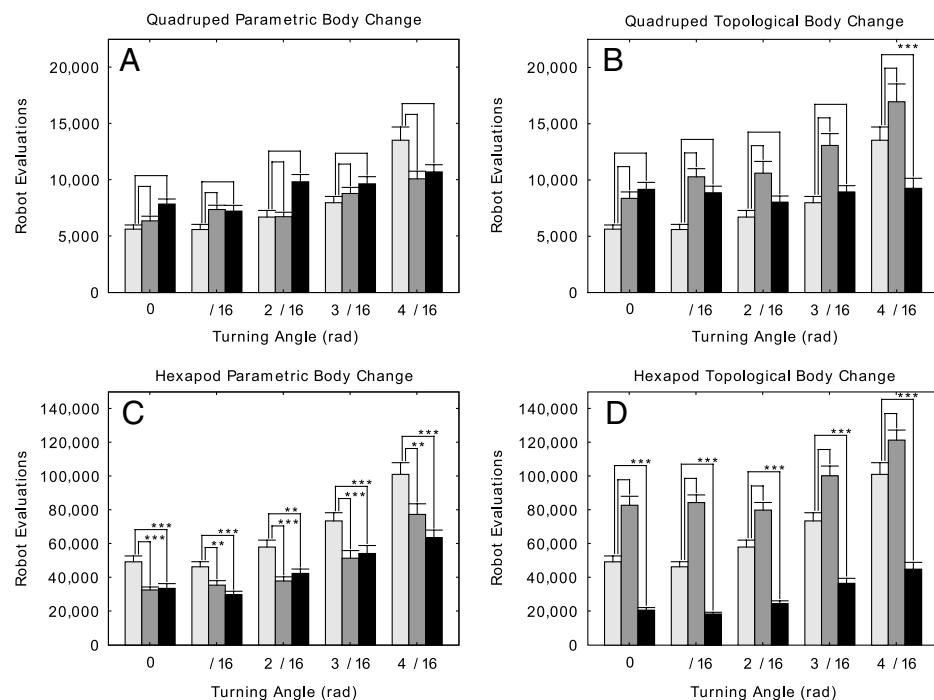


Fig. 3. How morphological change affects the time to discovery of the desired behavior in the quadruped (A and B) and hexapod (C and D) robot. Light gray bars indicate the number of controllers that had to be evaluated when no morphological change was allowed. Dark gray bars indicate the discovery time when the robots' body plans did not change during a robot's lifetime, but did change over evolutionary time. Black bars indicate the discovery time when body plans changed during each robot's lifetime, and also over evolutionary time. The dark gray and black bars in A and C report the impact of changing the robot's body plans parametrically; the dark gray and black bars in B and D report the impact of changing the robot's body plans topologically. Asterisks report statistically significant differences between no morphological change and topological, ontogenetic morphological change. Error bars report one unit of standard error of the mean.

change accelerates behavior discovery significantly compared to parametric change for all five task environments (Fig. 3 C and D): The p values obtained from the Mann-Whitney U test produced $p = 4.3 \times 10^{-6}$, $p = 1.7 \times 10^{-6}$, $p = 1.2 \times 10^{-8}$, $p = 8.7 \times 10^{-4}$, and $p = 4.4 \times 10^{-5}$ for the five increasingly challenging task environments, respectively.

This advantage can be explained as follows: When parametric body plan change is employed, evolution must shape controllers that orchestrate the spine and all four (or six) legs in the initial generations. When topological change is employed, however, only the spine motors must be orchestrated initially: the legs, while growing, have relatively little impact on the displacement and heading of the robot. Subsequent evolution aligns the legs' motions with the trunk's motion. This result corroborates observations from learning theory (18, 19) that manually clamping some degrees of freedom in a robot with a fixed body plan, and then gradually releasing them during learning, can accelerate behavior acquisition (20).

It was found that in general (Fig. 3) as the task thus became more challenging, the benefit of employing topological and ontogenetic morphological change during behavior optimization became more pronounced, compared to not allowing morphological change, for both robots. In short, the infant body plans scaffold behavior acquisition: Maintaining balance does not become an issue until robots in later evolutionary phases must begin and maintain an upright form during locomotion. It is known that environmental and social scaffolding accelerate the rate of learning in humans (21) and robots (22). Here, the robot's own body scaffolds behavior acquisition.

Impact of Different Developmental Trajectories. It is possible that the advantage conferred by morphological change is an artifact of the particular evolution-of-development trajectory employed: Robots grow from an infant to an adult form gradually, but this change is accelerated such that later robots transition increasingly rapidly, and the final robots begin evaluation with the adult form. To test this a different developmental trajectory was formulated: Initial robots transition from the infant to adult form as before, but in the second phase robots begin with a form one-third of the way from the infant to the adult form and progress to the full adult form (Fig. S2 J and N). In the third phase they begin with

a form two-thirds of the way to the adult form (Fig. S2 K and O), and in the fourth phase begin and end in the adult form.

This trajectory was tested on both robots, in all five task environments, and using parametric (Fig. S2 I–L) and topological (Fig. S2 M–P) change. These trials were contrasted with another set of trials in which robots maintain the infant form throughout evaluation in the first phase (Fig. S2 A and E), robots transition from the infant to the adult form in the second phase (Fig. S2 B and F), robots begin evaluation with a form midway between the infant and adult form and transition to the adult form in the third phase (Fig. S2 C and G), and finally begin evaluation with the adult form in the fourth phase.

It is predicted that this latter developmental trajectory should perform poorly as controllers will suffer a large decrease in fitness when they experience the adult form for the first time in the second phase. The former trials, however, should perform well as controllers experience all forms of the body plan during the first phase. Fig. S3 shows that this is indeed the case. As for the first developmental trajectory, when topological and ontogenetic change is allowed using the new developmental trajectory, and robots in the first phase experience the infant and adult forms, phototaxis is discovered more rapidly for the final, upright robot (Fig. S3, green bars) than it is if robots do not experience the adult form during the first phase (Fig. S3, red bars) or no morphological change is allowed (Fig. S3, light gray bars).

Robustness of Evolved Behaviors. It was found that not only did the right kind of morphological change accelerate the discovery of successful behavior, but the discovered controllers were more robust to novel environmental perturbations than controllers discovered when morphological change was disabled. This was accomplished by taking the final controllers from each trial, which successfully produced phototaxis in the upright robot, and reevaluating them in the same robot but applying slight random perturbations to the robot while it behaved. One hundred of these noisy reevaluations were performed for each of the controllers, and the mean difference in performance was measured as a percentage decrease in the overall fitness of the noisily evaluated robots (Fig. 4).

Robustness can arise as a result of modularity (23), but no increased structural (24) or functional (25) modularity was detected in the controllers evolved within the morphologically changing

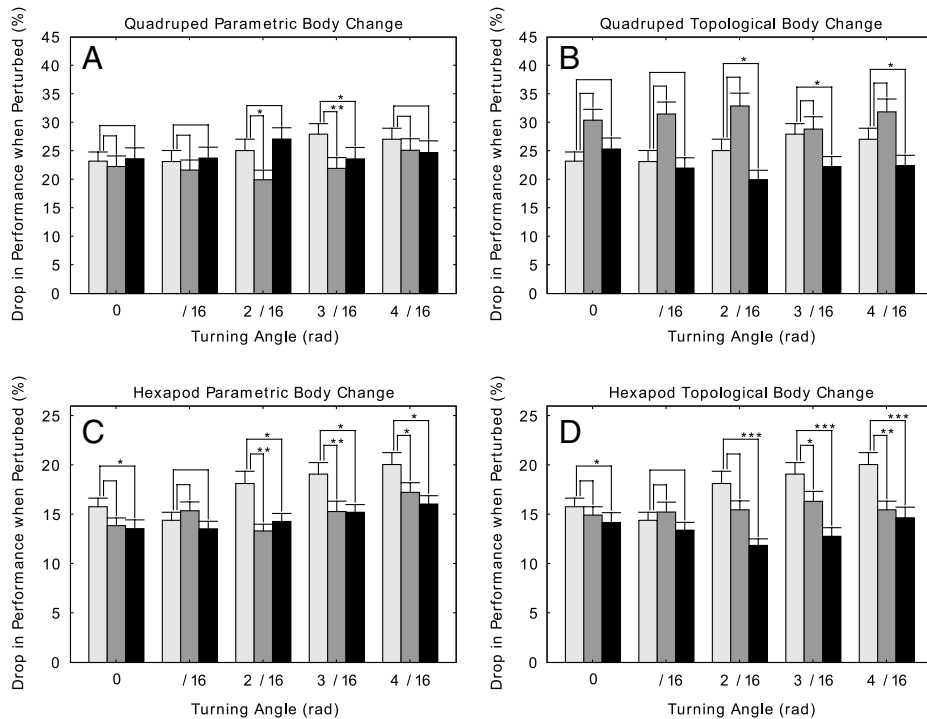


Fig. 4. How morphological change affects the robustness of the discovered behaviors. Bars report results from the same set of trials described in Fig. 3. The final robot capable of phototaxis from each independent experiment was reevaluated 100 times in the same simulated environment in which it evolved, but now exposed to small random external perturbations. The reduction in its ability to reach the light source was computed as the percent difference between the original distance it traveled and its new distance traveled during the perturbation.

robots compared to the morphologically static robots. Despite this, it was observed that the spine tends to be actuated more and the limbs less in the morphologically changing robots, which results from the fact that only the spine exists to generate locomotion in the infant legless form, and this behavior is retained in the descendent legged robots. (This can be seen by comparing the gaits shown in *Movies S1* and *S2*, which result from evolution with no morphological change, with the gaits shown in *Movies S3–S7*, which result from evolution with morphological change.) The different ways the spine and legs are employed in both cases suggests that there may be a form of functional modularity evolving across the robot's body plan: The spine becomes, and remains, responsible for propulsion while the legs adopt the function of retaining balance. This observation, however, requires support from more detailed investigation.

Aside from the possibility of modularity, it is conjectured that robustness was increased because controllers evolved in ancestral robots with different body plans had to maintain the desired behavior over a wider range of sensor-motor relationships than controllers evolved in robots with fixed body plans. In the ancestral robots the wide range was caused by changes in the body plan; in the reevaluated robots, this wide range was caused by external environmental perturbations.

Robustness to external perturbation is of extreme importance for any robots deployed in uncontrolled environments. Much work (26–28) has thus been dedicated to automatically producing robust robot behaviors; however, in all cases more evaluations had to be performed during training: Robustness arises in robots that must maintain a given behavior in different task environments. Here, robust controllers were obtained using fewer evaluations, compared to behaviors obtained from robots that experience no morphological change.

Conclusions

Here I have demonstrated a hitherto unknown relationship between development, evolution, morphology, and the neural control of behavior: In a population of simulated organisms evolving toward a desired behavior, that behavior is achieved more rapidly—and the resulting behavior is more robust—if the organism's neural-network controllers evolve in a succession of body plans

that change over evolutionary time and also over the lifetime of a single organism. In contrast, realizing adaptive behavior in machines has to date focused on dynamic controllers but static morphologies: Neural-network based (25, 29) or other learning methods (30–32) allow for changes to the robot's controller but not its body plan. This is an inheritance from traditional artificial intelligence in which computer programs were developed that had no body with which to affect, and be affected by, the world (33, 34). In the few projects that have investigated robot morphological change (9–14) or the evolution of physical structures in general (35, 36), the goal has been to discover structures unseen in the biological world and unintuitive to an engineer, and thus not likely to arise through manual design. However, the relationship between robust behavior and morphological change had not been investigated.

This work suggests that the automated design, manufacture (10), and deployment (8) of adaptive machines should allow for neurological and morphological change at ontogenetic and phylogenetic time scales. One class of robots that allows for morphological change during behavior is modular robotics (37), particularly at the meso- and microscale (38) where greater topological and material change is possible than if only a few macroscopic modules are employed.

In the biological domain it has been theorized that environmental fluctuations as well as mixtures of directional and stabilizing selection can favor the evolution of developmental robustness (23). This dynamic has been observed in biological organisms (39) as well as in computational models (24, 40, 41). Further it has been argued that, through genetic assimilation, this developmental robustness can lead to higher evolvability (42) in biological organisms: Genetic robustness can eventually supplant developmental robustness (43). Thus the work here provides another example that robustness and evolvability can be increased through the evolution of development.

Finally, this work also provides a unique tool for investigating whether lifetime morphological change provides behavioral opportunities for biological organisms. Currently it is difficult to determine whether morphological change is a response to biological constraints or whether the change itself provides some survival benefit to the animal in question (44). New computational

methods such as the one presented here that simulate phylogenetic and ontogenetic change to body plans and neural networks could serve as unique tools for investigating such hypotheses.

Materials and Methods

This material describes in more detail the simulated robots and their controllers.

The Simulated Robots. A 10-degree-of-freedom (DOF) quadrupedal and 16-DOF hexapedal robot were employed in this work. Both have actuated spines and employ the same sensor modalities. The trunk of the quadrupedal robot is composed of a front and back segment; the segments are attached by a 2-DOF actuated rotational joint that allows the segments to yaw and roll relative to one another. The segments may yaw through $[-50^\circ, 50^\circ]$ and roll through $[-20^\circ, 20^\circ]$. The wider range of motion through the coronal plane helps the robot to turn when the light source is placed to its left. Each of the four legs is attached to the front or back segment by a 2-DOF actuated rotational joint. One DOF sweeps the leg through a transverse plane centered at the leg. The joint's range is defined as $[(-10+x)^\circ, (10+x)^\circ]$, where x is the angle of the leg relative to body. If the leg is absent or is horizontal to the body, $x = 0^\circ$; if the leg is vertical, $x = -90^\circ$. For those experimental regimes in which the robot undergoes ontogenetic morphological change, x changes as the robot behaves. For the second DOF in the leg, if the leg is horizontal the joint sweeps $[-50^\circ, 50^\circ]$ through the robot's coronal plane; if the leg is vertical, it sweeps through the sagittal plane; for other angles of the leg it sweeps through a plane intermediate between the coronal and sagittal plane. The hexapedal robot has a different mechanical construction to ensure the results were not specific to a given body plan. The hexapod's spine is composed of three segments, with each pair connected by two, 2-DOF actuated rotational joints. Both joints allow the segments to yaw and roll relative to one another, but unlike the quadruped each DOF can sweep through $[-30^\circ, 30^\circ]$. Each of the six legs is controlled by a 1-DOF rotational joint that sweeps the leg $[(-10+x)^\circ, (10+x)^\circ]$ through a transverse plane centered at the leg. As for the quadruped robot, x is the joint's set point that

can change from $x = 0^\circ$ if the leg is absent or horizontal through to $x = -90^\circ$ if the leg is vertical. Both robots are equipped with proprioceptive sensors at each DOF, a vestibular and photosensor at their centroids, and a second and third photosensor at the base of their front left and right legs, respectively. Each foot is also equipped with a binary tactile sensor.

The Controllers. Each of the two robots is equipped with a continuous time recurrent neural network (45). For each robot, a fully connected network is employed: Each degree of freedom is assigned a motor neuron, each motor neuron is connected with a synapse to every other motor neuron, and every sensor is connected to every motor neuron. At each time step of the simulator each motor neuron is updated according to

$$\tau_i y'_i = -y_i + \sum_{j=1}^m w_{ji} \sigma(y_j - \theta_i) + \sum_{j=1}^s n_{ji} s_j, \quad [1]$$

where m is the number of motor neurons in the robot ($m = 10$ and $m = 16$ for the quadruped and hexapod, respectively), s is the number of sensor neurons in the robot ($s = 5 + 10$ and $s = 5 + 16$ for the quadruped and hexapod, respectively), τ_i is the time constant associated with motor neuron i , y_i is the value of neuron i (with a range in $[0.0001, 1.0]$), $\sigma(x) = 1/(1 + e^{-x})$ is an activation function that brings the value of neuron i back into $[0, 1]$, w_{ji} is the weight of the synapse connecting neuron j to neuron i (with a range in $[-16, 16]$), θ_i is the bias of neuron i (with a range in $[-4, 4]$), n_{ji} is the weight of the synapse connecting sensor j to neuron i (with a range in $[-16, 16]$), and s_j is the value of sensor j . The optimizer was found to be insensitive to different settings for the ranges, so these ranges were adopted from previous reported values (45). This yielded 270 and 624 free parameters to be optimized by evolution for the quadruped and hexapod, respectively. Movies of the simulated and physical robots can be found in [SI Text](#).

ACKNOWLEDGMENTS. This work was supported by National Science Foundation Grant CAREER-0953837.

- Tinbergen N (1963) On aims and methods of ethology. *Z Tierpsychol* 20:410–433.
- Datteri E, Tamburrini G (2007) Biorobotic experiments for the discovery of biological mechanisms. *Philos Sci* 74:409–430.
- Hoekstra H, Coyne J (2007) The locus of evolution: Evo devo and the genetics of adaptation. *Evolution* 61:995–1016.
- Nolfi S, Floreano D (2000) *Evolutionary Robotics* (MIT Press, Boston, MA).
- Harvey I, et al. (2005) Evolutionary robotics: A new scientific tool for studying cognition. *Artif Life* 11:79–98.
- Smith T, Husbands P, O'Shea M (2003) Local evolvability of statistically neutral Gas-Net robot controllers. *Biosystems* 69:223–243.
- Floreano D, Dürr P, Mattiucci C (2008) Neuroevolution: From architectures to learning. *Evol Intell* 1:47–62.
- Bongard J, Zykov V, Lipson H (2006) Resilient machines through continuous self-modeling. *Science* 314:1118–1121.
- Sims K (1994) *Evolving 3D Morphology and Behaviour by Competition Artificial Life IV* (MIT Press, Boston, MA), pp 28–39.
- Lipson H, Pollack JB (2000) Automatic design and manufacture of artificial lifeforms. *Nature* 406:974–978.
- Bongard JC (2002) Evolving modular genetic regulatory networks. *Proceedings IEEE 2002 Congress on Evolutionary Computation* (Inst of Electrical Engineers, London), pp 1872–1877.
- Hornby G, Pollack JB (2002) Creating high-level components with a generative representation for body-brain evolution. *Artif Life* 8:223–246.
- Ventrella J (1998) Designing emergence in animated artificial life worlds. *Virtual Worlds* (Springer, Berlin), pp 143–155.
- Komosin M, et al. (2003) The Framsticks system: Versatile simulator of 3D agents and their evolution. *Kybernetes* 32:156–173.
- Ijspeert AJ, Crespi A, Ryzcko D, Cabelguen JM (2007) From swimming to walking with a salamander robot driven by a spinal cord model. *Science* 315:1416–1419.
- Nelson G, Blankespoor K, Raibert M (2006) Walking BigDog: Insights and challenges from legged robotics. *J Biomech* 39:S360.
- Saranli U, Buehler M, Koditschek D (2001) RHex: A simple and highly mobile hexapod robot. *Int J Robot Res* 20:616–631.
- Bernstein N (1967) *The Co-ordination and Regulation of Movements* (Pergamon Press, London).
- Sporns O, Edelman G (1993) Solving Bernstein's problem: A proposal for the development of coordinated movement by selection. *Child Dev* 64:960–981.
- Lungarella M, Metta G, Pfeifer R, Sandini G (2003) Developmental robotics: A survey. *Connect Sci* 15:151–190.
- Wood D, Bruner J, Ross G (1976) The role of tutoring in problem solving. *J Child Psychol Psych* 17:89–100.
- Dorigo M, Colombetti M (1994) Robot shaping: Developing situated agents through learning. *Artif Intell* 70:321–370.
- Wagner GP, Altenberg L (1996) Complex adaptations and the evolution of evolvability. *Evolution* 50(3):967–976.
- Espinosa-Soto C, Wagner A (2010) Specialization can drive the evolution of modularity. *PLoS Comput Biol* 6:e1000719.
- Yamashita Y, Tani J (2008) Emergence of functional hierarchy in a multiple timescale neural network model: A humanoid robot experiment. *PLoS Comput Biol* 4:e1000220.
- Jakobi N (1997) Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adapt Behav* 6:325–368.
- DiPaolo A (2000) Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. *Proc Simulation of Adaptive Behavior* (MIT Press, Cambridge, MA), pp 440–449.
- Abdallah C, Dawson DM, Dorato P, Jamshidi M (2002) Survey of robust control for rigid robots. *IEEE Contr Syst Mag* 11:24–30.
- Lewis F (2002) Neural network control of robot manipulators. *IEEE Expert* 11:64–75.
- Abbeel P, Ng A (2004) Apprenticeship learning via inverse reinforcement learning. *Proceedings of the 21st International Conference on Machine Learning* (ACM Press, New York), pp 1–8.
- Thrun S, Burgard W, Fox D (2005) *Probabilistic Robotics* (MIT Press, Cambridge, MA).
- Collins S, Ruina A, Tedrake R, Wisse M (2005) Efficient bipedal robots based on passive-dynamic walkers. *Science* 307:1082–1085.
- Clark A (1999) An embodied cognitive science? *Trends Cogn Sci* 3:345–351.
- Pfeifer R, Lungarella M, Iida F (2007) Self-organization, embodiment, and biologically inspired robotics. *Science* 318:1088–1093.
- Lohn JD, Hornby GS, Linden DS (2008) Human-competitive evolved antennas. *AI EDAM* 22:235–247.
- Koza J, Al-Sakran S, Jones L (2008) Automated ab initio synthesis of complete designs of four patented optical lens systems by means of genetic programming. *AI EDAM* 22:249–273.
- Yim M, et al. (2007) Modular self-reconfigurable robot systems. *IEEE Robot Autom Mag* 14:43–52.
- White P, Zykov V, Bongard J, Lipson H (2005) Three dimensional stochastic reconfiguration of modular robots. *Proceedings of Robotics: Science and Systems* (MIT Press, Cambridge, MA), pp 161–168.
- Borenstein E, Ruppel E (2006) Direct evolution of genetic robustness in microRNA. *Proc Natl Acad Sci USA* 103:6593–6598.
- Lipson H, Pollack JB, Suh NP (2002) On the origin of modular variation. *Evolution* 56(8):1549–1556.
- Kashtan N, Alon U (2005) Spontaneous evolution of modularity and network motifs. *Proc Natl Acad Sci USA* 102:13773–13778.
- Kirshner M, Gerhart J (1998) Evolvability. *Proc Natl Acad Sci USA* 95:8420–8427.
- Ance LW, Fontana W (2000) Plasticity, evolvability, and modularity in RNA. *J Exp Zool* 288:242–283.
- Ruff C, Holt B, Trinkaus E (2006) Who's afraid of the big bad Wolff? "Wolff's law" and bone functional adaptation. *Am J Phys Anthropol* 129:484–498.
- Beer R (2006) Parameter space structure of continuous-time recurrent neural networks. *Neural Comput* 18:3009–3051.