

Morphological and Environmental Scaffolding Synergize when Evolving Robot Controllers

Artificial Life/Robotics/Evolvable Hardware

Josh C. Bongard
Department of Computer Science
University of Vermont
josh.bongard@uvm.edu

ABSTRACT

Scaffolding—initially simplifying the task environment of autonomous robots—has been shown to increase the probability of evolving robots capable of performing in more complex task environments. Recently, it has been shown that changes to the body of a robot may also scaffold the evolution of non trivial behavior. This raises the question of whether two different kinds of scaffolding (environmental and morphological) synergize with one another when combined. Here it is shown that, for legged robots evolved to perform phototaxis, synergy can be achieved, but only if morphological and environmental scaffolding are combined in a particular way: The robots must first undergo morphological scaffolding, followed by environmental scaffolding. This suggests that additional kinds of scaffolding may create additional synergies that lead to the evolution of increasingly complex robot behaviors.

Categories and Subject Descriptors

I.2.9 [Computing Methodologies]: Artificial Intelligence—Robotics

General Terms

Experimentation, Algorithms, Reliability

Keywords

Evolutionary Robotics, Shaping, Evolutionary Algorithms

1. INTRODUCTION

Scaffolding (or shaping) is a well-known phenomenon in learning theory, in which the learner's environment is initially simplified and then gradually made more difficult as learning progresses. This concept has its roots in psychology [26, 25] but has since become a common tool in robotics [12, 21, 4, 10]. In all previous approaches however, robot shaping involved simplifying the task environment: this might involve

making the task easier, or exposing robots to a gradually increasing number of task environments.

In recent work [5] it was shown that for legged robots, the robot's own body may serve as a vehicle for scaffolding the evolution of behavior. This was achieved by initially evolving neural network controllers for robots such that they perform phototaxis: they must move toward a light source placed in their environment in a fixed period of time. Four sequential phases of evolution were performed within any one run. First, robots gradually progress from an anguilliform (eel-shaped) legless body plan into a legged body plan as they behave (Figs. 1a, 2a).

In the second phase they progress from the legless to the legged form more rapidly (Figs. 1b, 2b). This distorts the fitness landscape slightly such that the previously-successful behavior is now only mediocre, thus triggering a short burst of evolution until a successful controller is re-discovered. Once such a controller is found the third phase of evolution begins in which the robot progresses from the legless to the legged form over the first third of its lifetime (Figs. 1c, 2c). In the fourth phase, robots begin and maintain the legged form throughout their lifetime (Figs. 1d, 2d).

It was found [5] that this kind of morphological scaffolding makes the behavior optimization process more evolvable: successful controllers are evolved for the final, legged form more rapidly than if controllers are only evolved within populations of robots that never change body plan. Moreover, the final controllers produced from the evolutionary runs in which robots under body plan change were more robust than those produced from the runs in which no body plan change occurred: when the controller were re-evaluated in slightly different environments the behavior was maintained more often for the former group than for the latter.

It was hypothesized (although not yet proven) that this increased robustness was the result of the fact that the controllers had to support phototaxis across a lineage of different body plans which generated slightly different sensor-motor relationships. This then allowed the controller to operate well when different environments, rather than the robot's own body plan, induced sensor-motor relationships different from those experienced during evolution.

In that work, however, only morphological scaffolding was employed: the robots' environment did not change during evolution. Here we investigate how two different kinds of scaffolding may be combined, and whether they increase

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '11, July 12–16, 2011, Dublin, Ireland.

Copyright 2011 ACM 978-1-4503-0557-0/11/07 ...\$10.00.

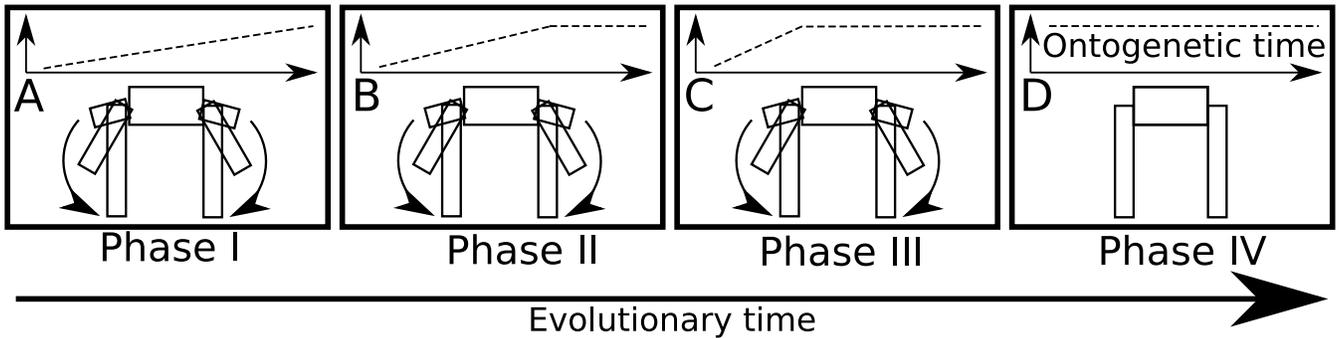


Figure 1: Schematic of morphological scaffolding. In the first phase (A), robots progress from a legless to a legged form uniformly over their lifetime. The inset figure shows the gradual rise of the robots’ center of mass as a result. In the second phase (B), robots progress from the legless to the legged form in two-thirds of the time, after which they maintain the legged form. In the third stage (C), body change occurs over the first third of the lifetime. In the fourth phase (D), robots begin, maintain and end in the legged form.

evolvability beyond that induced by using either scaffolding approach in isolation.

Morphological scaffolding is one way of investigating how change to robot body plans over evolutionary time impacts the ability to automatically evolve robot behaviors, the main goal of evolutionary robotics [13, 18]. In morphological scaffolding, body plan change is set by the experimenter; in many evolutionary robotics projects, evolution has control over changes to the controller as well as the morphology [22, 23, 17, 24, 1, 7, 16, 15, 8, 9, 11, 2, 19]. However, in these latter cases it is difficult to determine whether evolutionary changes to the body are a result of historical accident, or whether they gradually lead from morphologically-simple robots with limited ability to more complex robots with greater ability. Morphological scaffolding is an approach designed to investigate whether such a search gradient can be induced by body plan change. Such understanding should help with the design of evolutionary algorithms that more easily evolve robot morphology and control simultaneously.

The next section describes the methodology; section 3 reports results from combining the scaffolding approaches; and section 4 provides some concluding remarks.

2. METHODS

This section describes the robot, the evolutionary algorithm, and the two forms of scaffolding employed in this work.

2.1 The Robot Body

A 10 degree-of-freedom (DOF) quadrupedal robot was employed in this work (Fig. 2). All of the robots in this work are simulated, though morphological scaffolding was explored on a physical robot in [5]. The trunk of the robot is composed of a front and back segment; the segments are attached by a two 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 or right. Each of the four legs are 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 default angle of the leg relative to body.

If morphological scaffolding is used, the robot’s legs grow gradually as the robot moves, and are angled increasingly vertically. This is achieved by enforcing a change in x as the robot moves, and gradually lengthening the leg. If morphological scaffolding is not used, x is set and held at -90° , which results in vertical legs. 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 robot is equipped with proprioceptive sensors at each DOF; a vestibular and photosensor at its centroid; and a second and third photosensor at the base of its front left and right legs respectively (to allow for trilateration). Each foot is also equipped with a binary tactile sensor.

2.2 The Robot Controller

Each of the two robots is equipped with a continuous time recurrent neural network [3]. 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 $m = 10$ is the number of motor neurons in the robot, $s = 5 + 10$ is the number of sensor neurons in the robot, 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 evolutionary algorithm was found to be insensitive to different settings for the ranges, so these ranges

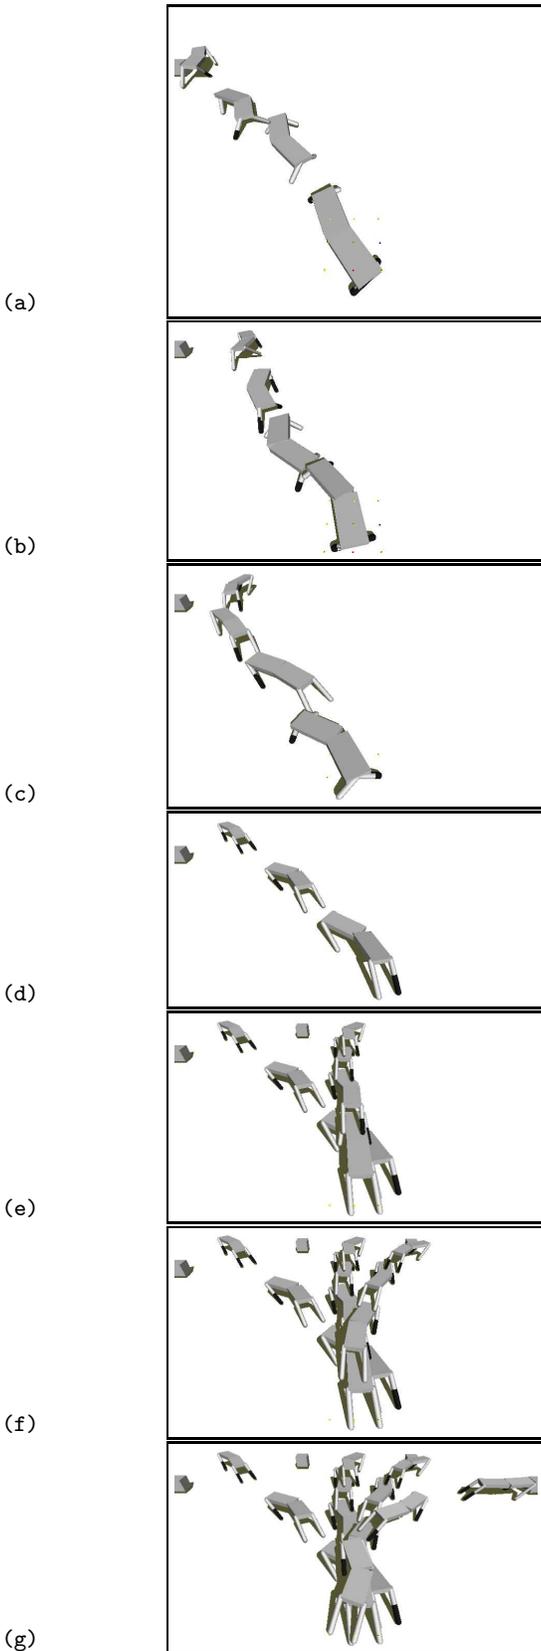


Figure 2: Results from a typical run when the robots' morphologies are scaffolded first (a-d), followed by scaffolding of their environment (e-g).

were adopted from previous reported values [3]. This yielded 270 free parameters to be optimized by evolution.

2.3 The Evolutionary Algorithm

The evolutionary algorithm used here is modeled on the Age-Layered Population System (ALPs) [14, 6]. An initial population of 400 random genomes were created, where a genome encodes values for each of the neural network parameters described above. The population was divided into 20 layers of 20 genomes each. In ALPs, genomes within a layer compete with one another; genomes between layers do not. Genomes on lower layers are younger and have lower fitness than genomes on higher layers. As evolution proceeds, genomes on lower layers periodically 'challenge' genomes on higher layers. If the lower genome is more fit than the higher genome the latter is displaced, and must in turn dislodge a higher genome or else be deleted.

In this work each genome is evaluated as follows. The encoded neural network parameters are downloaded onto the simulated robot, and the robot is allowed to perform for 1000 time steps as described in the previous section. The fitness of the robot is computed as

$$f_1 = \sum_{t=1}^{1000} \frac{p_1^{(t)}}{1000}, \quad (2)$$

$$f_2 = \sum_{t=1}^{1000} \frac{p_2^{(t)}}{1000}, \quad (3)$$

$$f_3 = \sum_{t=1}^{1000} \frac{p_3^{(t)}}{1000}, \quad (4)$$

where f_i is the i th fitness component of the controller, and $p_i^{(j)}$ is the value of the i th photoreceptor at time step j . Multiobjective optimization was employed in that only non-dominated genomes can produce offspring. A genome i is defined to be dominated if there exists another genome on its layer, j , for which

$$f_1^{(j)} > f_1^{(i)} \quad \wedge \quad f_2^{(j)} > f_2^{(i)} \quad \wedge \quad f_3^{(j)} > f_3^{(i)}.$$

Each step of the optimization process was performed as follows. A non-dominated genome from any of the 20 layers was selected at random. That genome is copied, and each network parameter in the genome was perturbed with 0.05 probability. If a network parameter is perturbed, the current value is replaced with a new value chosen with a uniform distribution from that network parameter's valid range. The new genome was then evaluated on the robot, and its fitness components were recorded.

The layer that produced the new genome is then scanned for a genome that can be dislodged by the new genome. A genome is dislodged if it is dominated by the new genome, or if it too old for its layer. Every genome is assigned an age a which is defined as $a = \lfloor n_a/400 \rfloor$, where n_a is the optimization step at which genome a 's original ancestor was created. An original ancestor is either a genome that was created during the first time step of the optimization process, or during a periodic resetting of the population's lowest layer when the genomes on that layer are deleted and replaced by new random genomes. Each of the 20 layers are also assigned a maximum age limit according to the Fibonacci sequence [1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144, 233, 377, 610, 987, 1597, 2584, 4181, 6765, 10946] (following [14]). If a genome's age

is greater than its layer’s age limit, it is defined as being too old for its layer.

If the new genome cannot dislodge a genome on its parent genome’s layer, the new genome is discarded. Otherwise, if the dislodged genome is on the highest level it is discarded. If the dislodged genome is not on the highest layer it attempts to dislodge a genome on the layer above it. It has been shown that this process allows for the continuous infusion of new genetic material into the evolving population, and the ability to dislodge solutions that have become mired on local optima [14].

Successful phototaxis is defined as a controller that achieves an average value of at least 0.5 for each of the three photoreceptors over the entire evaluation:

$$(f_1 \geq 0.5) \wedge (f_2 \geq 0.5) \wedge (f_3 \geq 0.5). \quad (5)$$

2.4 Environmental Scaffolding

When a controller is evaluated, it may be evaluated in from one to four environments. In the first environment the light source is placed to the front and 45° to the left of the robot (Fig. 2a). In the second environment the light source is placed in front of and 15° to the robot’s left (Fig. 2e). In the third and fourth environments, the light source is placed 15° and 45° to the robot’s right, respectively (Figs. 2f,g).

When environmental scaffolding is not employed, each controller is evaluated four times, one in each of the four environments (e.g. Fig. 2g). The controller’s fitness components are then averaged over all four evaluations:

$$f_1 = \sum_{k=1}^4 \left(\sum_{t=1}^{1000} \frac{p_1^{(t)}}{1000} \right) / 4, \quad (6)$$

$$f_2 = \sum_{k=1}^4 \left(\sum_{t=1}^{1000} \frac{p_2^{(t)}}{1000} \right) / 4, \quad (7)$$

$$f_3 = \sum_{k=1}^4 \left(\sum_{t=1}^{1000} \frac{p_3^{(t)}}{1000} \right) / 4. \quad (8)$$

When environmental scaffolding is employed, all controllers are initially only evaluated once in the first environment. Once a controller on one of the 20 population layer succeeds (eqn. 5), each controller on that layer is re-evaluated in the first and second environment: This layer now “hosts” two environments. Thus, over time, different layers evaluate controllers against different numbers of environments.

Thus, if a controller attempts to migrate to a layer above that hosts more environments, the controller is first evaluated in the additional environments and then tested to see whether it can unseat a controller on that layer. Alternatively, if a controller challenges a layer with fewer environments, the controller’s fitness components are recalculated across the fewer environments.

2.5 Morphological Scaffolding

Morphological scaffolding is similar to environmental scaffolding in that, as evolution proceeds, the controllers are exposed to increasingly challenging situations. In this form of scaffolding however the challenge arises from having to induce locomotion in robots that are increasingly unstable.

When morphological scaffolding is not employed, all controllers are evaluated on robots that begin evaluation in the upright, legged form, and maintain it throughout the evaluation period. When morphological scaffolding is employed,

the controllers are initially evaluated on robots that change body plans as they are evaluated. Initially, the robots start in the legless form and progress at a uniform rate to the legged form over the length of the evaluation period such that they attain the adult form just as evaluation ends (Figs. 1a, 2a).

Once a controller evolves on a layer that achieves successful phototaxis, all of the controllers on that layer are re-evaluated using a different robot: This robot transitions from the legless to the legged form over the first two-thirds of the evaluation period (Figs. 1b, 2b), and then maintains the adult form over the remaining one third of the evaluation period. If success is again achieved, controllers on that layer are re-evaluated using the third robot (Figs. 1c, 2c), and so on.

Thus, like for environmental scaffolding, one of four robot body plans is associated with each layer. Unlike environmental scaffolding however, controllers are only evaluated once, regardless of what robot is used: controllers do not need to retain the ability to produce successful phototaxis in the previous robot form. In this regime if a controller challenges the layer above and that layer hosts different robot forms, the controller is re-evaluated using that robot form before it attempts to unseat a controller on that layer.

3. RESULTS

Six experimental regimes of 100 independent evolutionary runs were conducted.

In the first regime, neither environmental nor morphological scaffolding was employed. Each controller was evaluated once in each of the four environments, using the upright form of the robot (e.g. 2g). Evolution continued until either a successful controller was evolved, or 30 CPU hours elapsed.

In the second regime, morphological scaffolding was employed, but environmental scaffolding was not. Each controller was evaluated in all four environments, but initially the controllers are evolved in the first robot form (Fig. 1a). When a successful controller is evolved, the controllers in that layer are re-evaluated in all four environments but with the second robot form (Fig. 1b). Evolution continues until a successful controller is evolved using the fourth and final robot form, or 30 CPU hours elapse.

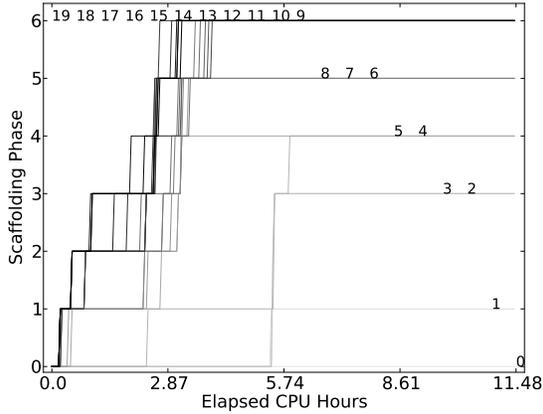
In the third regime, environmental scaffolding is employed, but morphological scaffolding is not. All controllers are only evaluated in the final, upright robot form. However, controllers are initially evaluated only in the first environment. When a successful controller is evolved, the controllers on that layer are re-evaluated in the first two environments, and so on. Evolution continues until a controller evolves that succeeds in all four environments.

In the fourth and fifth regimes, the two scaffolding regimes are both employed, but applied during evolution sequentially. In the fourth regime morphological scaffolding is applied first, followed by environmental scaffolding. This is achieved by following the process described for the second regime, but the controllers evolve through the four body plans in only the first environment. An example of four successive successful controllers evolved during a typical run from this regime are shown in Fig. 2a-d.

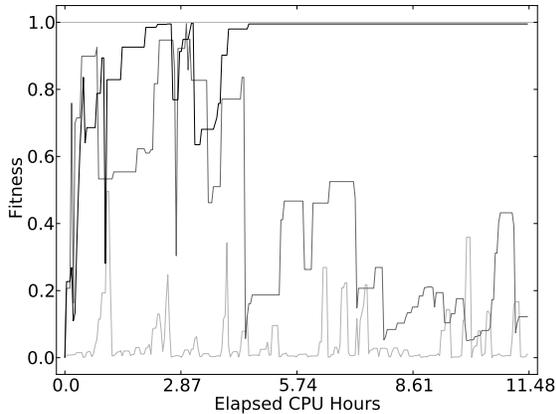
Once a controller evolves that succeeds in the first environment using the upright body plan (Fig. 2d), environmental scaffolding is applied: the controllers on that layer

Table 1: Overview of the experimental regimes.

Regime	Phases	Description
1	1	No morph or env scaffolding
2	4	Morph, but no env scaffolding
3	4	No morph, but env scaffolding
4	7	Morph, then env scaffolding
5	7	Env, then morph scaffolding
6	7	Env, morph scaffolding interleaved



(a)



(b)

Figure 3: Evolutionary dynamics observed in a typical run from the fourth experimental regime in which morphological scaffolding precedes environmental scaffolding. (a): Number of scaffolding phases supported by each of the 20 population layers. (b) Fitness of the best controller on the first (light gray), tenth (gray) and twentieth (black) layer.

are re-evaluated in the first and second environments using the upright body plan. Once a controller succeeds in this condition (Fig. 2e), the controllers on that layer are re-evaluated using the upright body plan in the first three environments, and so on. A successful controller from this same run that succeeds in the first three environments is shown in Fig. 2f, and the final controller that succeeds in all four environments—and triggers the termination of the run—is shown in Fig. 2g.

In the fifth regime, in each run environmental scaffolding is applied first, followed by morphological scaffolding. This

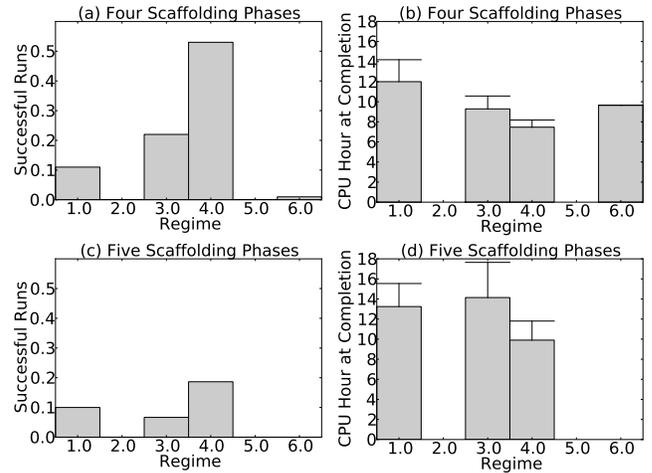


Figure 4: Final performance of the six experimental regimes. (a) Fraction of runs for each regime that evolve a successful controller before the maximum allotted time elapses. For those runs that do succeed, (b) reports the mean time required to do so. For the runs performed using five increments of morphological and environmental scaffolding, (c) reports the fraction of successful runs and (d) reports the mean time required to do so.

is accomplished by initially evolving the controllers in only the first environment using the first robot body plan. Once a successful controller is found, the controllers on that layer are re-evaluated in the first two environments using the first robot body plan, and so on. Once a controller is discovered that succeeds in all four environments, the controllers on that layer are re-evaluated in the second robot plan. This continues until a controller succeeds in all four environments using the upright body plan.

In the sixth and final regime, environmental and morphological scaffolding is interleaved. Controllers are initially evaluated using the first body plan in only environment. Once success occurs, controllers on that layer are re-evaluated using the first robot body plan in the first two environments. Once another controller succeeds on that layer, controllers are re-evaluated in the first two environments again, but with the second robot body plan. Once success occurs again, the controllers are re-evaluated in the first three environments using the second robot body plan, and so on until a controller is discovered that succeeds in all four environments using the upright body plan.

Table 1 summarizes the six experimental regimes. For the six regimes, one, four or seven scaffolding phases may be employed. For example for the fourth regime, four morphological scaffolding phases are employed in the first environment, followed by evaluations in the second, third and fourth environments.

Fig. 2 illustrates the behaviors produced by seven controllers produced by one typical run of the fourth regime. It shows the first controller to succeed in the first environment using the first environment (Fig. 2a). A short time later another controller succeeds using the second body plan in the first environment (Fig. 2b). Some time after a third and then a fourth controller succeed using the third (Fig. 2c) and fourth (Fig. 2d) body plan, respectively, a controller

succeeds in the first two environments (Fig. 2e). This controller is followed by one that succeeds in the first three environments (Fig. 2f), and then finally by one that succeeds in all environments (2g).

Fig. 3 reports the evolutionary dynamics for this same run. Fig. 3a reports the number of scaffolding phases supported by each layer as the run proceeds. It can be seen that very rapidly a controller succeeds on all layers for the first phase, except for the lowest layer. This is due to the periodic re-randomizing of the controllers on the lowest layer to introduce fresh genetic material into the population.

As the run proceeds, upper layers host increasingly more scaffolding phases. This is due to the upward migration of successful (or near successful) controllers, which triggers the addition of new phases once they arrive. By the fifth CPU hour, the tenth layer and higher require all controllers to be evaluated using the final body plan in all four environments.

It can be seen that a learning gradient has emerged. As controllers move up the first 10 layers, they must perform well in increasingly difficult situations. This corresponds to evaluation using an increasingly unstable body plan, and then evaluation in more than one task environment.

Fig. 3b reports the fitness of the best controller on the first, tenth and twentieth layer. As can be seen, in the uppermost layer the best controller is very close to success for over half the period of the run, before a controller is produced that actually succeeds in all four environments using the final robot body plan. Fitness is significantly lower, and more variable in the two lower layers. On the tenth layer this is due to the high churn of good controllers migrating upwards, replaced by less fit ones. On the first layer, this is due to periodic introduction of new, random controllers.

Fig. 4 reports the final performance of the six regimes. Fig. 4a reports the fraction of runs that succeed before the maximum time allotment, and, of those runs that do succeed, Fig. 4 reports the mean time required to do so. It is clear that the fourth regime, in which morphological scaffolding precedes environmental scaffolding, significantly outperforms the other five regimes. Additionally, even though some of the other regimes produce successful runs, those runs take significantly longer to complete than those in the fourth regime.

Fig. 5 reports the mean evolutionary progress across the six experimental regimes, and provides some insight into why the fourth regime outperforms the other regimes. Fig. 5b shows that, when the robots are always exposed to all four environments, the first robot body plan does not provide a good substrate on which to evolve successful controllers. This is as it is very rare for a population layer to experience a successful controller and thus increase the number of scaffolding phases on that layer.

Fig. 5c shows that when only the final upright, unstable body plan is used, environmental scaffolding does not provide much of a benefit: it takes a relatively long time for controllers to succeed in the first few environments. This can be explained by the fact that the upright robot provides a rough fitness landscape: it takes a long time to evolve a successful controller in the first environment using this body plan. Once a second environment is added however, the fitness landscape becomes even more rugged: even if the robot manages to maintain balance and reach the target object in the first environment, the controller may cause the robot to fall over in the second environment.

Conversely, once a controllers on a layer have progressed through the four robot body plans in the fourth experimental regime (Fig. 5d), success in the four environments occurs at a rapid pace. This is seen by the rapid rise of the curve corresponding to environmental scaffolding in Fig. 5d, compared to the slower rise of the curve corresponding to environmental scaffolding in Fig. 5c.

Fig. 5e shows that, for the fifth experimental regime, environmental scaffolding proceeds even more slowly when the first robot body plan is used, compared to the third and fourth regimes. However, there is an initial rapid rise in the number of environments the robot is exposed to: in less than one CPU hour, almost all runs are hosting controllers that can succeed in the first three environments. The addition of the fourth environment (if it occurs) takes considerably longer, and thus very few robots progress to morphological scaffolding (as evidenced by the low curve corresponding to morphological scaffolding in Fig. 5e).

One explanation for this observation is that it may be relatively easy to evolve blind locomotion for the first robot body plan (which can allow successful controllers to rapidly evolve in the first few environments), but very difficult to evolve taxis behavior for this body plan (such that all four environments are conquered). This may be because it is much easier to turn while locomoting with the final, upright body plan, but more difficult with the prone, legless form that changes into the legged form while moving. More investigation is required to determine whether this hypothesis stands.

The above explanation may also apply to the sixth regime (Fig. 5f), in which several environments must be conquered by controllers while residing in intermediate robot body plans.

One challenge in any form of scaffolding is choosing an appropriate number of learning increments, and the amount of challenge added at each increment. Here, four scaffolding increments was chosen arbitrary, and was used for both morphological and environmental scaffolding. In order to determine how sensitive the results were to the choice of the number of scaffolding increments, an additional set of runs were conducted with five morphological and five environmental scaffolding increments.

For morphological scaffolding, this meant that robots began (as before) by changing slowly from the legless to the legged form (Fig. 1a). In the second phase, the robots progressed from the legless to the legged form over the first three quarters of the evaluation period. In the third and fourth phases the robots progressed from the legless to the legged form over the first half, and then quarter of the evaluation period, respectively. In the fifth phase the robots maintained the upright legged form as before.

For environmental scaffolding this meant that there were five task environments, in which the target object was placed 45° and 20° to the front and left of the robot, respectively. It was placed directly in front of the robot in the third environment, and then 20° and 45° to the front and right of the robot in the fourth and fifth environment, respectively.

Fig. 4c,d indicate that again, the fourth regime outperformed the other five regimes, but the success rate was much lower. This suggests that there is some intermediate number of scaffolding increments that is appropriate for this task: if no increments are used, success is low, and if too many are used, the controllers are unnecessarily evaluated in too many different situations.

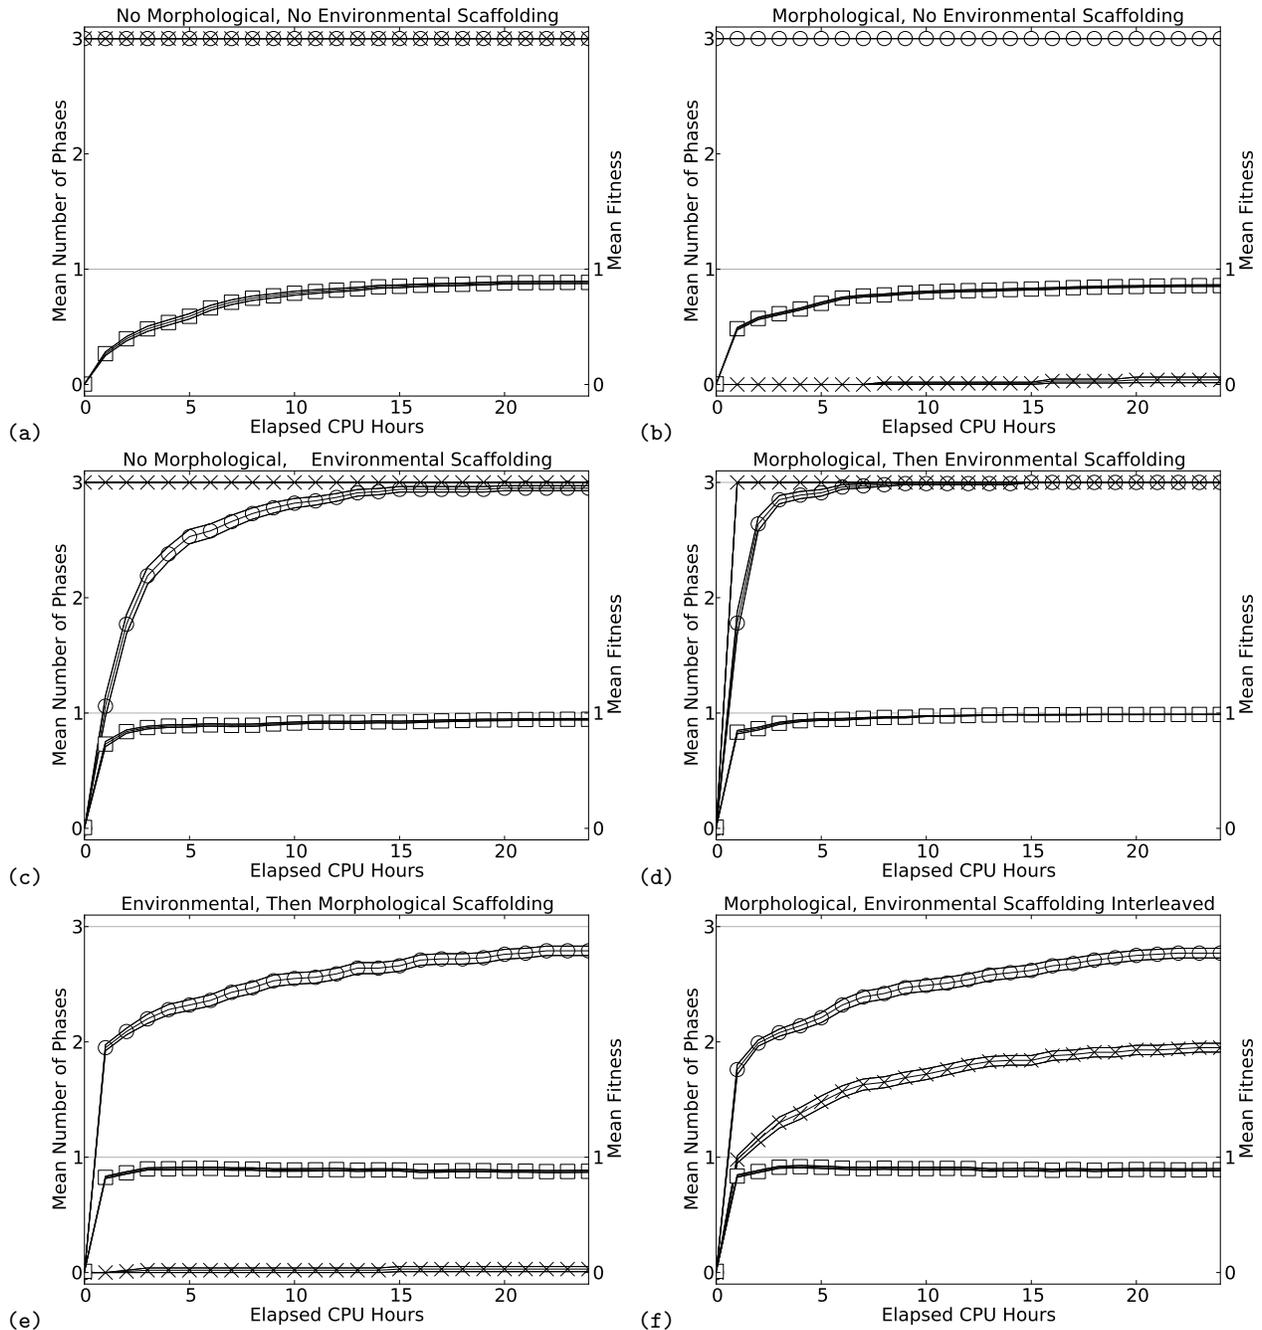


Figure 5: Mean evolutionary progress across the six experimental regimes. The curves with circles report the mean environmental scaffolding phase at the twentieth layer. The curves with crosses report the mean morphological scaffolding phase at the twentieth layer. The curves with squares report the mean fitness of the controllers on the twentieth layer. If a regime does not support one (or both) kinds of scaffolding, the last phase for that type of scaffolding is drawn. The bracketing lines of each curve report one unit of the standard error of the mean.

4. CONCLUSIONS

This paper has shown that combining different forms of scaffolding can increase the chance of evolving non-trivial robot behavior compared to (1) not using scaffolding at all; (2) using only one form of scaffolding; (3) combining the two forms of scaffolding in the wrong way; or (4) using too many scaffolding increments. It was found that if morpho-

logical scaffolding preceded environmental scaffolding, significant performance improvement was achieved. This result can be explained by reference to previous work [5] that showed that gradually evolving from stable body plans to increasingly unstable body plans smooths the fitness landscape. However, as controllers evolve through these body plans, the controllers must maintain the behaviors through

a wide range of sensor-motor relationships induced by the different body plans. This can result in increased robustness when the evolved controllers are placed in noisy environments which induce different sensor-motor relationships, but now these differences are induced by the environment rather than the robot's own body. In future work we plan to investigate the generality of this result to different robots and different tasks. In addition, we would like to investigate whether adding more forms of scaffolding further improves evolvability. For example Reil *et al.* [20] found that applying forces that restore balance to a falling bipedal robot early in evolution, and then gradually reducing those forces as evolution proceeds, resulted in better evolution of walking controllers. One might also gradually increase the amount of noise in the task environment, or gradually allow less time to complete the desired task. Finally, we wish to investigate whether the choice of which form of scaffolding to use—and how many scaffolding increments should be employed—could itself be placed under evolutionary control.

5. ACKNOWLEDGMENTS

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

6. REFERENCES

- [1] A. Adamatzky, M. Komosinski, and S. Ulatowski. Software review: Framsticks. *Kybernetes: The International Journal of Systems & Cybernetics*, 29:1344–1351, 2000.
- [2] J. Auerbach and J. Bongard. Dynamic Resolution in the Co-Evolution of Morphology and Control. In *12th International Conference on the Synthesis and Simulation of Living Systems (ALife XII)*, Odense, Denmark, 2010.
- [3] R. Beer. Parameter space structure of continuous-time recurrent neural networks. *Neural Computation*, 18:3009–3051, 2006.
- [4] J. Bongard. Behavior chaining: incremental behavioral integration for evolutionary robotics. In S. Bullock, J. Noble, R. Watson, and M. A. Bedau, editors, *Artificial Life XI*, pages 64–71. MIT Press, Cambridge, MA, 2008.
- [5] J. Bongard. Morphological change in machines accelerates the evolution of robust behavior. *Proceedings of the National Academy of Sciences*, 108(4):1234, 2011.
- [6] J. Bongard and G. Hornby. Guarding against premature convergence while accelerating evolutionary search. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 111–118. ACM, 2010.
- [7] J. Bongard and C. Paul. Investigating morphological symmetry and locomotive efficiency using virtual embodied evolution. *Proceedings of the Sixth International Conference on Simulation of Adaptive Behaviour*, pages 420–429, 2000.
- [8] J. Bongard and R. Pfeifer. Repeated structure and dissociation of genotypic and phenotypic complexity in Artificial Ontogeny. *Proceedings of The Genetic and Evolutionary Computation Conference (GECCO 2001)*, pages 829–836, 2001.
- [9] J. Bongard and R. Pfeifer. Evolving complete agents using artificial ontogeny. *Morpho-functional Machines: The New Species (Designing Embodied Intelligence)*, pages 237–258, 2003.
- [10] J. C. Bongard. Innocent until proven guilty: Reducing robot shaping from polynomial to linear time. *IEEE Transactions on Evolutionary Computation*, 2011. To appear.
- [11] N. Chaumont, R. Egli, and C. Adami. Evolving virtual creatures and catapults. *Artificial Life*, 13(2):139–157, 2007.
- [12] M. Dorigo and M. Colombetti. Robot Shaping: Developing Autonomous Agents Through Learning. *Artificial Intelligence*, 71(2):321–370, 1994.
- [13] I. Harvey, P. Husbands, D. Cliff, A. Thompson, and N. Jakobi. Evolutionary robotics: the Sussex approach. *Robotics and Autonomous Systems*, 20(2-4):205–224, 1997.
- [14] G. Hornby. Steady-state ALPS for real-valued problems. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 795–802. ACM, 2009.
- [15] G. Hornby and J. Pollack. Evolving L-systems to generate virtual creatures. *Computers & Graphics*, 25(6):1041–1048, 2001.
- [16] H. Lipson and J. B. Pollack. Automatic design and manufacture of artificial lifeforms. *Nature*, 406:974–978, 2000.
- [17] H. H. Lund, J. Hallam, and W.-P. Lee. Evolving robot morphology. *Proceedings of the IEEE Fourth International Conference on Evolutionary Computation*, 1997.
- [18] S. Nolfi and D. Floreano. *Evolutionary Robotics*. MIT Press, Boston, MA, 2000.
- [19] M. Pilat and C. Jacob. Evolution of vision capabilities in embodied virtual creatures. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 95–102. ACM, 2010.
- [20] T. Reil and P. Husbands. Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168, 2002.
- [21] L. Saksida, S. Raymond, and D. Touretzky. Shaping robot behavior using principles from instrumental conditioning. *Robotics and Autonomous Systems*, 22:231–250, 1997.
- [22] K. Sims. Evolving 3D morphology and behaviour by competition. *Artificial Life IV*, pages 28–39, 1994.
- [23] J. Ventrella. Explorations of morphology and locomotion behaviour in animated characters. *Artificial Life IV*, pages 436–441, 1994.
- [24] J. Ventrella. Designing emergence in animated artificial life worlds. In *Virtual Worlds*, pages 143–155. Springer, 1998.
- [25] L. Vygotsky. Thought and language. *Annals of Dyslexia*, 14(1):97–98, 1964.
- [26] D. Wood, J. Bruner, and G. Ross. The role of tutoring in problem solving. *J Child Psychol Psychiatry*, 17(2):89–100, 1976.