Adapting Morphology to Multiple Tasks in Evolved Virtual Creatures

Dan Lessin, Don Fussell, and Risto Miikkulainen

The University of Texas at Austin, Austin, TX 78712
dlessin@cs.utexas.edu

Abstract

The *ESP* method for evolving virtual creatures (Lessin et al., 2013) consisted of an *encapsulation* mechanism to preserve learned skills, a human-designed *syllabus* to build higher-level skills by combining lower-level skills systematically, and a *pandemonium* mechanism to resolve conflicts between encapsulated skills in a single creature's brain. Previous work with ESP showed that it is possible to evolve much more complex behavior than before, even when fundamental morphology (i.e., skeletal segments and joints) was evolved only for the first skill. This paper introduces a more general form of ESP in which full morphological development can continue beyond the first skill, allowing creatures to adapt their morphology to multiple tasks. This extension increases the variety and quality of evolved creature results significantly, while maintaining the original ESP system's ability to incrementally develop complex behaviors from a sequence of simpler learning tasks. In the future, this method should make it possible to build EVCs with complex and believable behavior.

Introduction

Since their introduction two decades ago by Sims (1994), evolved virtual creatures (or *EVCs*; Figure 1) have had a significant impact on multiple fields, including graphics, evolutionary computation, artificial life, and robotics (Shim and Kim, 2003; Lehman and Stanley, 2011; Miconi, 2008; Lipson and Pollack, 2000). But despite their many incarnations in the intervening years, their behaviors have not grown more complex. Sims' original work demonstrated light following. This level of complexity has been approximately matched multiple times since (Miconi, 2008; Pilat and Jacob, 2010), but was never clearly exceeded until the ESP system of Lessin et al. (2013). ESP was shown to construct EVC behaviors approximately twice as complex as any seen before, measured as the number of discriminable behaviors in the creature's repertoire. Moreover, in principle there is no upper limit on behavioral complexity yet established within this new framework.

The ESP system is named for its three principal components–*encapsulation*, *syllabus*, and *pandemonium*–defined as follows:

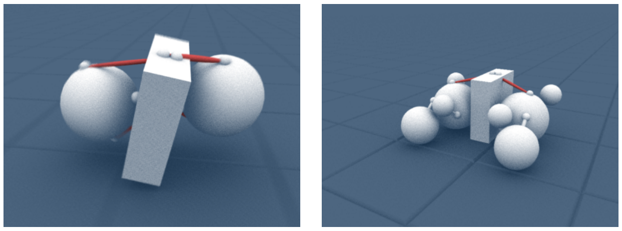

(a) Initial locomoting creature. (b) Heavy smashing arms.

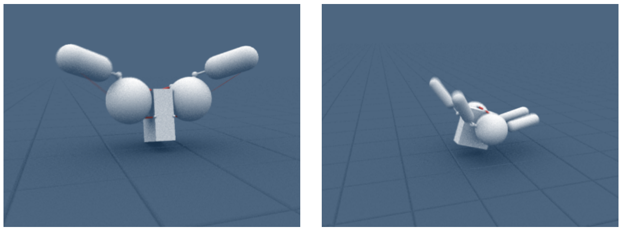

(c) Smashing flail arms. (d) Jump with anti-tip limbs.

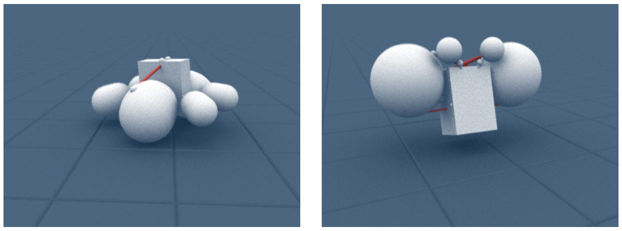

(e) Smashing tail, stabilizers. (f) Jump with heavier body.

Figure 1: Adapting EVC morphology to multiple tasks. (a) A creature adapted for locomotion. From this creature, creatures (b) through (f) were evolved using the extended ESP method described in this paper. Each of them has developed a new technique (with corresponding morphological changes) for accomplishing an additional task–in this case, delivering a strike to the ground–while still maintaining the ability to perform the initial skill (locomotion) to prescribed levels. In the original ESP system, these secondary adaptations would have been impossible.

- Once a new skill is learned, it can be *encapsulated*, which preserves that ability and makes it easily accessible to future evolution.

- A human-designed *syllabus* is used to direct the sequential acquisition of these component skills toward the larger goal of achieving new, more complex behaviors, building up hierarchically from simpler ones.

- A *pandemonium*-like mechanism (Selfridge, 1958) is employed to resolve conflicts between competing encapsulated skills in the increasingly complex brain.

The initial ESP implementation did achieve its goal of breaking the behavioral-complexity barrier for EVCs. However, it applied only to a significantly restricted case–one in which some of the most important morphological changes (those to the creature's skeletal segments and joints) were prohibited after the first skill's evolution was complete. (Note, however, that muscle drives and photoreceptors–described below–could continue to be evolved throughout all evolutionary stages, even in the original system, since they can be added without disrupting existing abilities.)

This is a serious limitation. For example, what if a creature is evolved for an initial skill such as locomotion, then is asked to adapt to a largely orthogonal skill such as reaching up to a high target? That creature may or may not have the required morphological capacity for performing the second task, as determined only by the accidents of evolution.

This paper introduces a significantly extended version of ESP, in which a retesting and reconciliation scheme replaces previous absolute limitations on morphological evolution. Morphology is thus fully evolved to suit the requirements of more than just a single skill.

In the following sections, this new ESP implementation is described. The results demonstrate a significant increase in the useful variety and quality of evolved creatures, while the ability to incrementally develop complex behaviors from a sequence of simpler learning tasks is maintained.

The Underlying EVC Implementation

Both the ESP implementation of Lessin et al. (2013)–called *original ESP* in this paper–and its improvement that is presented in later sections–called *extended ESP*–were built on a reimplementation of Sims' original work (1994) that will be referred to as *the underlying EVC implementation*. That underlying EVC system–along with some important novel extensions–is described in the balance of this section, and a representative sample of its results is shown in Figure 2.

Evolutionary Algorithm

A conventional evolutionary algorithm is used, with elitism, fitness-proportionate selection, and rank selection (Mitchell, 1998). In addition, the most challenging tasks employ shaping (Urzelai et al., 1998; Gomez and Miikkulainen, 1997). Fitness is evaluated in a physically simulated virtual environment implemented with NVIDIA PhysX.

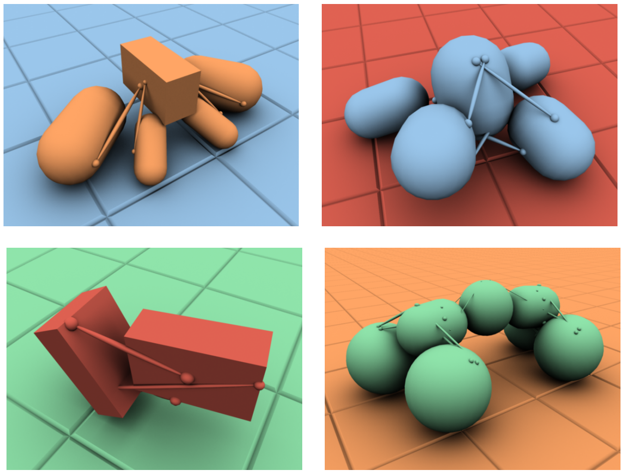

Figure 2: Typical results from the underlying EVC system. These examples were all evolved to complete a forward locomotion task—a common baseline result for EVCs.

Morphology

As in Sims' original work, creature morphology is described by a graph-based genotype, with graph nodes representing body segments, and graph edges representing joints between segments By starting at the root and traversing the graph's edges, the phenotype is expressed. Reflexive edges as well as multiple edges between the same node pair are allowed, making it possible to define recursive and repeated body substructures easily. In addition, as in Sims' work, reflection of body parts as well as body symmetry are made easily accessible to evolution, with only a single mutation required to produce each of them. In this implementation, all PhysX primitives are available for use as body segments: boxes, spheres, and capsules. Joints between segments may be of most of the types offered by PhysX, specifically: fixed, revolute, spherical, prismatic, and cylindrical. In contrast to the typical technique of separately evolving explicit joint limits, most limitations on joint movement are provided implicitly by creature structure through natural collisions between adjacent segments.

In addition to the typical segments and joints, the implementation of the underlying EVC system also evolves muscle drives and photoreceptors, as described below.

Muscles

In a break with traditional evolved virtual creature systems, which typically use forces exerted directly at joints, the underlying EVC system of this paper uses simulated muscles as actuators. Each muscle ((b) in Figure 3) is defined by two attachment points on adjacent segments, along with a maximum strength value. In simulation, the muscle is implemented as a spring, with muscle activation modifying the spring constant. In addition to acting as an effector, each muscle also produces a proprioceptive feedback signal based

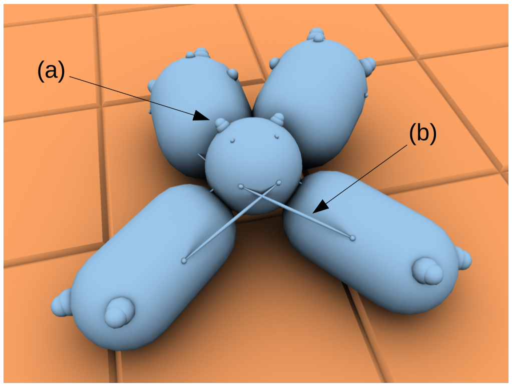

Figure 3: Photoreceptors (a) and muscles (b) bring sensing and actuation to creatures in the underlying EVC system. For both, function depends upon placement, so creature form develops meaningfully as capabilities are evolved.

on its current length. For each muscle, two nodes are added to the brain: one that accepts an input to set the muscle's activation, and another that makes the muscle's proprioceptive output signal available to the rest of the brain. Muscle drives benefit EVCs in several ways: they can be used even on creatures without joints; they only need to exist where they are useful, not at every degree of freedom of every joint; and they have the potential to embody some degree of control intelligence, with benefits for both aesthetics and the reduction of cognitive load (Lessin et al., 2014).

Photoreceptors

For tasks involving light sensing, creatures are allowed to develop simple photoreceptors ((a) in Figure 3), defined only by a direction from the center of their parent segment. This direction indicates a location on the creature's surface as well as an orientation for the receptor. The signal produced by the receptor is determined by light strength, distance, occlusion, and the difference between the direction to the light and the sensor's orientation. Multiple lights are allowed. For each photoreceptor in the body, a corresponding brain node is added which makes the receptor's output signal available to the rest of the brain.

Control

In a manner which is again very similar to that of Sims, creature control is provided by a brain composed of a set of nodes connected by wires (as in Figure 5a). Nodes receive varying numbers of input wires, and use their inputs to compute an output value (always in the range [0,1]) which may be sent to other wires. Signals originate from sensors in the body as well as certain types of internal brain nodes, travel through the network of internal nodes and wires, and ultimately control the operation of actuators (muscles) in the physically simulated body. For each step of physical simulation, control signals move one step through the brain.

In addition to the special node types for muscles and photoreceptors (described above) and for encapsulation (described in the *Encapsulation* subsection), the following node types are allowed: sinusoidal, complement, constant, scale, multiply, divide, sum, difference, derivative, threshold, switch, delay, and absolute difference.

The Original ESP System

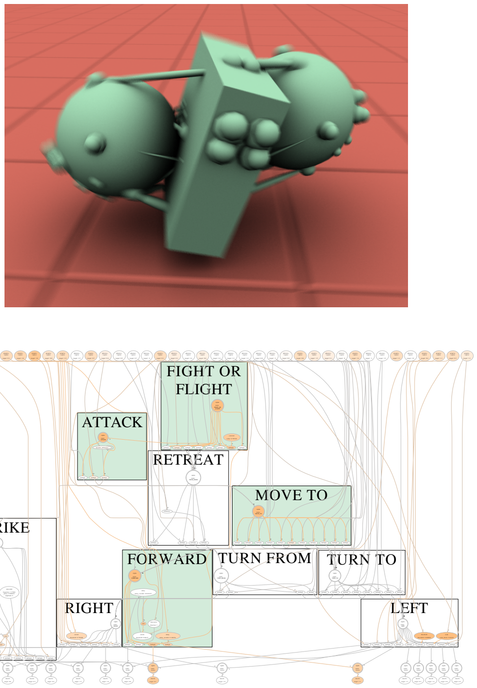

Figure 4: The body and brain of a creature evolved using the original ESP system. This creature achieved a level of behavioral complexity approximately double the state of the art. In addition to following a light, it was able to attack it or flee from it, as part of a hierarchical fight-or-flight behavior.

The original ESP system (Lessin et al., 2013) consists of three elements added to the underlying EVC system: encapsulation, syllabus, and pandemonium. In this section, each of these components is described in detail.[1]

Encapsulation

The first important element of the ESP system is a mechanism to *encapsulate* newly learned skills (Figure 5). This element accomplishes two important goals: It ensures that previously learned skills (and the body components on which they rely) are preserved, and it makes these skills easily accessible to future evolutionary development.

Figure 5a depicts a brain evolved for forward locomotion, and Figure 5b shows the result of encapsulating it. First, the nodes that compute the old skill have been preserved and

---

[1] For a video overview, see: http://youtu.be/dRLNnJlT8rY

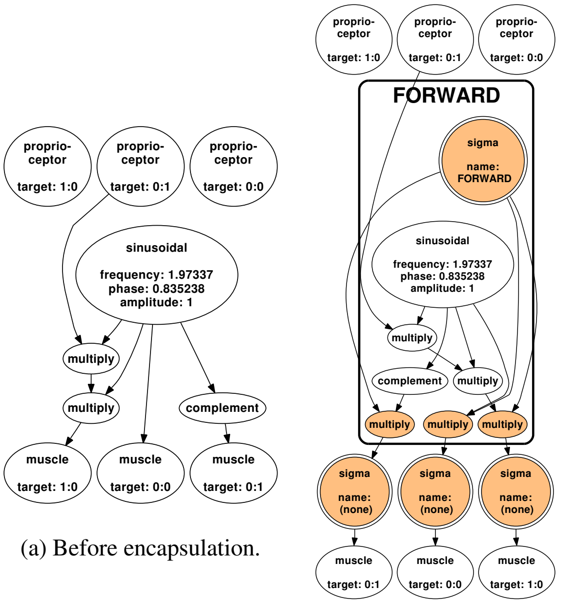

(a) Before encapsulation.

(b) After encapsulation (with new nodes shaded).

Figure 5: Encapsulation. The automated encapsulation of an evolved skill—in this case, forward locomotion—ensures that it will persist throughout future evolution, while also allowing it to be easily activated as a unit by future skills.

locked (meaning that they have been marked to disallow any changes by future evolution). Second, a new *multiply* node has been inserted into every output wire leaving the encapsulated skill. The internals of the skill will continue to function as before, always trying to perform their forward locomotion task, but now, a second signal sent to each new multiply node will modify those outgoing forward-locomotion control signals, scaling them by a number in [0,1]. Third, a single controlling node (called a *sigma node* for its function as a summation of zero or more inputs) is added, which sends its output to all of the new multiply nodes. So for each signal s_i leaving a node in the FORWARD LOCOMOTION skill (such as the *complement* node), the new signal after encapsulation (s'_i) is computed as $s'_i = \sigma s_i$ where σ is the output of the controlling sigma node.

Now, with encapsulation complete, the entire forward locomotion skill can be activated and deactivated as a unit by using the controlling sigma node just as if it were a single muscle in the creature's body. (Incidentally, note that this brain's actual muscle nodes have been hidden behind additional sigma nodes to allow future evolution to share control over them when appropriate.)

With the ability to preserve completed skills for easy future access, the next issue is how they may be acquired in sequence to achieve a larger learning objective.

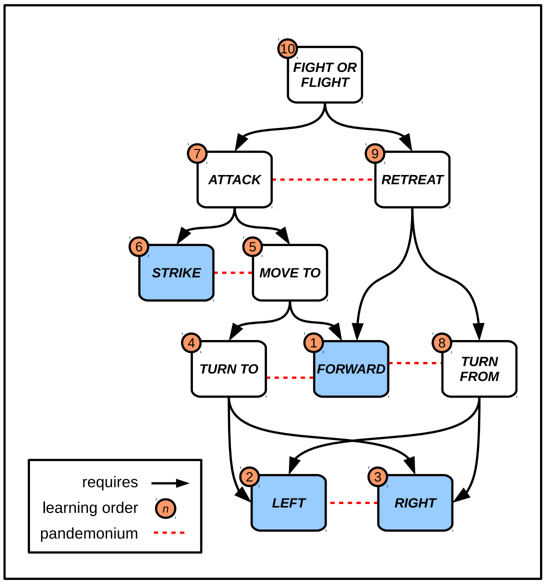

Figure 6: An example syllabus as a graph. Graph nodes represent individual subskills to be learned, directed edges indicate dependency between subskills, and the numbering indicates a proposed learning order that satisfies the dependency requirements. Pandemonium relationships are indicated by dashed red lines. The shaded nodes (called *leaf nodes*) affect only the body, rather than other nodes, and constitute the focus of the extended ESP system discussed in this paper.

Syllabus

While it is certainly possible for human students to learn a complicated topic independently, even advanced learners would be expected to develop further and with greater ease with the benefit of an expert-designed syllabus. The syllabus acts as a sequence of landmarks through the space of possible solutions, decomposing the larger learning task into a succession of more manageable steps between these waypoints.

In the ESP system, the *syllabus* consists of an ordered sequence of fitness goals used to reach the ultimate, larger goal. This collection of intermediate goals (each one defined by a fitness function) is designed by a human expert with the aim of making attainable goals more reliably learnable, and bringing previously unattained goals within reach.

For example, assume that you want to evolve a virtual creature with some of the behavioral complexity demonstrated in an internet cat video.[2] Rather than simply drifting smoothly toward a target, this creature might run to the target, then strike it, and perhaps even run away if the target is perceived as threatening. Without a syllabus, a single fitness test evaluating such a complex collection of skills might be constructed, but evolutionary progress would be unlikely.

---

[2] E.g.: "THE BEST CAT VIDEO YOU'LL EVER SEE" [sic], http://youtu.be/20mrEtabOLM

Consider, instead, how this complex behavioral goal could be broken down into an ordered sequence of smaller learning tasks. The clearly achievable goal of locomotion will be the first target. The ability to turn left and the ability to turn right are of a similarly manageable difficulty, and will be attempted next. Then, with left and right turns mastered, and the ability to develop photoreceptors, it is relatively straightforward to maintain orientation toward a light source. And with the ability to face a light and the ability to move forward, navigating to that light might be a similarly achievable goal. Proceeding in this manner, a knowledgeable human designer might produce a sequence of subskills to be learned. Each subskill would be attainable with basic EVC methods, and earlier subskills would serve to make later skills easier to learn.

This information may be conveniently summarized in a graph, encompassing subskills to be learned, dependency between subskills, learning order, and pandemonium, as seen in Figure 6. In fact, the syllabus in this figure was used by Lessin et al. (2013) to produce the creature illustrated in Figure 4. This creature achieved a level of behavioral complexity approximately double the state of the art for evolved virtual creatures at that time. More specifically, not only was it able to move to a light target, as previous creatures had done, but also strike the target upon reaching it, and flee from the target when it became dangerous, doubling the complexity of behaviors in prior work.

At this point, a complex skill can be broken into smaller subskills, and those subskills can be cumulatively acquired, but a potential problem still remains: How are competing signals from the multiple sub-brains within a single creature resolved?

Pandemonium

Consider the following example based on the syllabus graph of Figure 6. A creature evolved through this syllabus will ultimately have parts of its brain devoted to both left and right turns. But it seems unlikely that both of these abilities should ever be used at the same time. So the syllabus designer might place the left and right-turn skills in a *pandemonium* relationship with each other (Selfridge, 1958), meaning that whichever one is most active at any given moment will be allowed to send its output at full strength, and the other will have its output entirely suppressed. Under a system like this, sub-brains within the creature can compete for overall control, with little risk of sabotaging the usefulness of the entire brain's function. In Figure 6, a full set of pandemonium relationships is indicated by red dashed lines between subskill nodes.

Although this original ESP system achieved more complex behavior than before, body segments and joints could not continue to adapt beyond the first skill's completion. In the next section, a more general form of the ESP system is described that makes this possible.

The Extended ESP System: Adapting Morphology to Multiple Tasks

In this section, a new version of ESP is described with extended evolution of morphology.

Replacing Morphological Constraints with Retesting

The initial implementation of the ESP system enforced strict limits on morphological changes after the completion of the first skill: Although changes to muscles and photoreceptors were allowed, segments and joints were fixed. Due to this constraint, previously learned skills could be expected to work reliably throughout the syllabus-based construction. On the other hand, this limitation makes it difficult to develop certain abilities later. For example, a creature may succeed in developing forward locomotion and the ability to turn left, but—due to the construction of a certain joint evolved for locomotion—be unable to learn to turn right, even after many generations of evolution.

Luckily, this limitation was undertaken only to make an initial success in the original system easier to achieve. It can be removed simply by expanding and modifying the fitness evaluations applied during learning: Instead of locking segments and joints after the first skill is developed, successive skills could be allowed to change these attributes, as long as new testing shows that such changes will not invalidate earlier abilities.

However, such an increase in testing threatens to make an already computationally demanding problem significantly more difficult, especially because the system is intended to be open ended. Assuming n skills and one independent test for each skill, full retesting of all previous skills at each step of the syllabus would produce an $\mathcal{O}(n^2)$ growth in the required testing, instead of the current linear growth.

Fortunately, the retesting can be considerably reduced by focusing it where it matters. Consider again the syllabus graph shown in Figure 6. The skills that have a direct influence on the creature's body are shaded, and will be referred to as *leaf* skills. These are: FORWARD LOCOMOTION, LEFT TURN, RIGHT TURN, and STRIKE. Once these skills are successfully established, the remaining non-leaf skills can be evolved independently (in an order that meets dependency requirements), without the need for any retesting. This approach stops the $\mathcal{O}(n^2)$ growth in testing requirements significantly earlier than would otherwise be possible–in the syllabus of Figure 6, for example, after four skills instead of 10 (assuming all leaf skills are learned first).

New ESP Algorithm

This subsection describes the implementation of the new, more general form of the ESP algorithm, taking advantage of the leaf skills. The method is comprised of two stages. The first stage consists of a fixed number of generations during which the new skill's control and body evolves, as described in Algorithm 1. During this stage, existing encapsulated skills in the brain do not change, but if any morphological changes reduce these skills' fitness beyond a preset limit, the creature will be marked as unfit. In this way, the new skill is given free rein to adapt the body to its needs, provided that sufficient ability in all existing skills is retained.

The second stage runs for a fixed number of generations for each of the old skills, during which the morphology is temporarily locked–ensuring that the abilities achieved by the new primary skill are preserved–and each of the already existing skills gets a chance to reconcile itself to the new body (Algorithm 2). Since the morphology is fixed, these skills can develop completely independently–each skill can adapt to the new body, without degrading any of the other skills in the brain.

Proceeding in this manner, this extension of the ESP algorithm allows new leaf skills to seek their own adaptations to morphology as well as control, with a reasonable expectation that–as in the old system–existing skills will be maintained, allowing abilities to accumulate incrementally as in the original ESP.

Algorithm 1: Full evolution of morphology and control for new skill s'.

```
foreach generation do
    foreach individual in the population do
        mutate morphology;
        mutate control for new skill $s'$;
        foreach existing skill $s$ do
            evaluate fitness for $s$;
            if fitness for $s$ has decreased significantly then
                set individual fitness to 0;
                proceed to next individual;
            end
        end
        evaluate fitness for $s'$;
        set individual fitness to fitness for $s'$;
    end
    produce new population from existing one;
end
```

Algorithm 2: Reconciling existing skills to body changes made for new skill s'.

```
foreach existing skill $s$ do
    foreach generation do
        foreach individual in the population do
            mutate control for skill $s$;
            evaluate fitness for $s$;
            set individual fitness to fitness for $s$;
        end
        produce new population from existing one;
    end
end
```

Results

The experiments demonstrate the advantages of the continuing morphological evolution enabled by the extended ESP algorithm. In the first subsection (*Strike Results*), an experiment from the original ESP system is reproduced in the extended ESP system, with dramatically different results. In the second subsection (*High-Reach Results*), a learning challenge designed to highlight the extended system's advantages is presented, and detailed benefits are described. Note that, while extended ESP maintains original ESP's ability to construct complex hierarchical behaviors, that ability is inherited largely without modification in the new system. Therefore, the experiments in this paper are used instead to demonstrate the extended system's success in more challenging applications that were impossible in the original system. Video illustrating both of these result sections can be viewed online.[3]

Strike Results

An important part of the original ESP system's primary experimental result was to add a strike behavior to a locomoting creature (toward the larger goal of developing a complex fight-or-flight behavior). In this section, that portion of the old experiment is reproduced in the extended system, and a broad range of novel strategies and morphological changes is observed.

Strike in Original ESP Figure 1a depicts a creature evolved for locomotion in the underlying EVC system (which is common to both the original and extended versions of ESP). In original ESP, that creature consistently solved the challenge of producing a striking behavior by using its existing skeletal structure to either jump up and down or smash the ground with its limbs, without any opportunity to explore the potential for new strategies or better adaptation that might result from continuing full morphological development.

Strike in Extended ESP When the morphology is allowed to continue to evolve, however, new strategies become possible, and even old strategies may be better executed with morphological changes adapted to their specific needs. The extended ESP system develops a variety of such solutions, as can be seen in Figures 1b through 1f.

High-Reach Results

The second experiment was designed to highlight the benefits of the extended ESP system over the old implementation. Specifically, a selection of three different locomoting creatures was evolved to learn the additional skill of reaching for a high target, and the subsequent differences of results for the original and extended ESP implementations were examined in detail. The extended ESP system led to two types of improvements: 1) increased diversity of results, and 2) increased numerical fitness.

---

[3] http://youtu.be/fyVr7gdGEPE

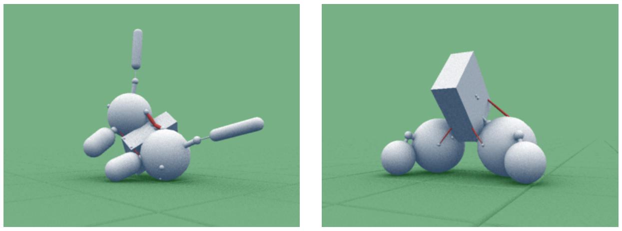

(a) Tipping, long new limbs. (b) Push-up, extended limbs.

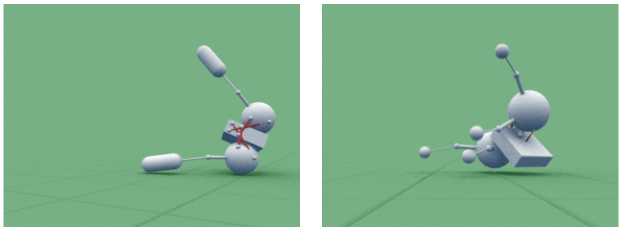

(c) Telescoping limbs. (d) Telescoping, anti-tip limbs.

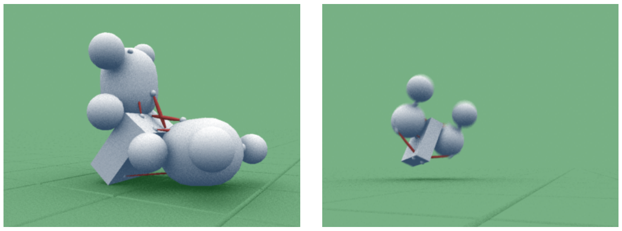

(e) Tip with enlarged limbs. (f) Jump, swing extensions up.

Figure 7: Greater variety through continued morphology evolution. The locomoting creature of Figure 1a was further evolved using the extended ESP system to adapt to a high-reach task. The results demonstrate the potential of continued morphology evolution to produce a great degree of useful variety.

Greater Variety The locomoting creature of Figure 1a was evolved toward the new high-reach goal, using both the original and the extended ESP implementation.

In the original ESP system, only two strategies were observed, within which the results were extremely uniform. Using morphology unchanged from the original locomotion result, all such creatures developed to either jump as high as possible, or reach a limb up by tipping over onto the other limb. In both cases, the results were limited by the inability of skeletal morphology to adapt to this new task.

In the extended ESP system, in contrast, a wide variety of results was observed, in which a number of novel strategies were used, often to great effect. These solutions are illustrated in Figure 7 (a) through (f).

Better Fitness Another successful solution to the locomotion task produced by the original ESP is shown in Figure 8a. This snake-like creature achieved a high reach by extending one end of its long morphology, while the rest of the body maintained balance.

Extended ESP improved upon this creature by changing its morphology for the secondary task, while its strategy remained unchanged (Figure 8b). It grew an additional body segment that enabled the higher reach, while allowing it still to perform locomotion to acceptable standards.

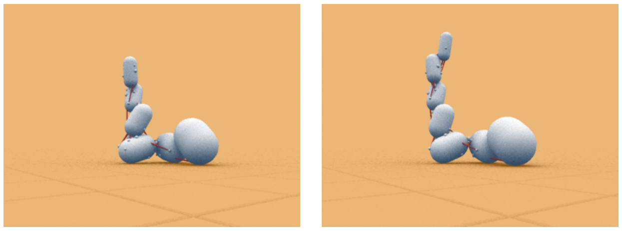

(a) Original ESP result. (b) Result in new ESP.

Figure 8: Improved fitness via continued morphology evolution. These results demonstrate how the extended ESP system (b) can produce better fitness values (i.e., a higher reach) than the original ESP system (a) by allowing the addition of new body segments.

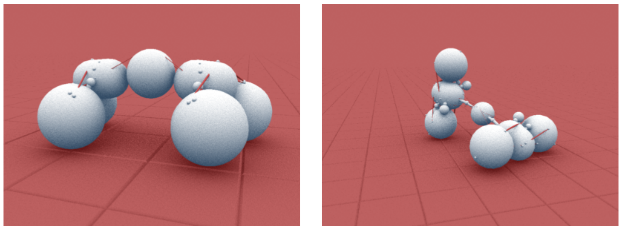

(a) Initial locomoting creature. (b) Subtle body changes.

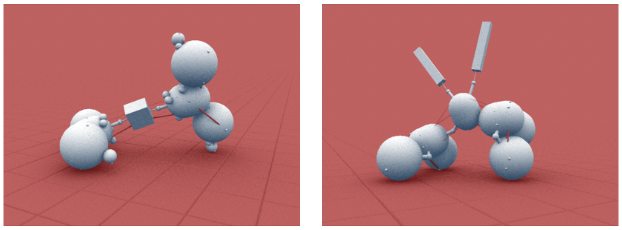

(c) More obvious body changes. (d) Dramatic changes in morphology.

Figure 9: Greater variety and improved fitness. The initial locomoting quadruped (a) is evolved for high reach in the extended ESP system (b)-(d). Through a variety of strategies, each of the extended ESP creatures scores better on this new task than any creature from the original ESP system.

Greater Variety *and* Better Fitness The relatively complex quadruped seen in Figure 9a was a third type of solution developed by the underlying EVC system for the locomotion task. In continued evolution of the high-reach task in the original ESP system, this creature's results were again extremely uniform in approach and fitness. They all reached up with a single limb, and all with approximately equal success. In the extended system, the ability to continue to adapt morphology to this new task led to a diverse set of useful results, all of which were also more fit than those produced with the original technique.

For example, Figure 9b depicts a creature that pursues the same strategy as the creature in Figure 9a, yet does so more effectively due to subtle morphological adaptations. In Figure 9c, more obvious morphological adaptations have been added to further exceed the uniform performance limit experienced by this creature in the original system, while still employing the same basic technique. In Figure 9d, even more dramatic changes to morphology provide a new way of solving the high reach: This creature employs a new pair of tall, dedicated limbs to even further exceed the performance of the original ESP system.

Discussion and Future Work

Although the extended ESP algorithm has removed the original system's explicit limitations on body changes after the first skill, development of morphology throughout the acquisition of complex skills is still not fully general and completely unlimited. First, the retesting requirements would make morphological development impractical if continued through too many steps of leaf skill addition. To mitigate this issue in the future, it may be possible to do the retesting periodically rather than universally, and run the tests in parallel. Also, the more leaf skills there are, the more likely it is that the morphological change required by one skill is harmful to the others. This limitation is more difficult to overcome, and indeed it reflects the conflicting demands that any creature faces when dealing with complex environments.

In principle, it would be desirable to continue morphology evolution throughout all skill adaptations–not just for those skills that are leaves in the syllabus graph. In this manner, it might be possible to develop morphologies that make transitions between skills easier–for example by making the creature more stable or more agile. To make such evolution possible without increasing testing too much, a rolling horizon of "leaf" skills that travels through the syllabus hierarchy might be implemented. The idea is that once a lower skill will no longer be explicitly required for any subsequent skills, it need not be retested or maintained at all. With a well-chosen sequence of skill learning, an approximately constant-sized wave of "leaves" might result, sweeping gradually through the hierarchy.

Another potential area for future study that might be possible with extended ESP is to increase the morphological complexity of EVCs. While it was not investigated in this paper, the ability of the body to embody multiple types of physical intelligence simultaneously could ultimately lead to a greater degree of physical complexity, and thereby more interesting and capable EVCs.

Conclusion

This paper described an extension of the original ESP system to continue adaptation of morphology beyond the initial skill, while still incrementally producing high-complexity behaviors. The benefits of this continued adaptation were demonstrated through experiments in which the extended ESP system generated a greater variety of solutions and solutions with higher fitness. Therefore, in the future, this method should make it possible to build EVCs with complex and believable behavior.

Acknowledgements

This research was supported by NSF grants DBI-0939454 and IIS-0915038, NIH grant R01-GM105042, and Intel's Visual Computing Program.

References

Gomez, F. and Miikkulainen, R. (1997). Incremental evolution of complex general behavior. *Adaptive Behavior*, 5:317–342.

Lehman, J. and Stanley, K. (2011). Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pages 211–218. ACM.

Lessin, D., Fussell, D., and Miikkulainen, R. (2013). Open-ended behavioral complexity for evolved virtual creatures. In *Proceeding of the Fifteenth Annual Conference on Genetic and Evolutionary Computation Conference*, GECCO '13, pages 335–342, New York, NY, USA. ACM.

Lessin, D., Fussell, D., and Miikkulainen, R. (2014). Trading control intelligence for physical intelligence: Muscle drives in evolved virtual creatures. In *Proceeding of the Sixteenth Annual Conference on Genetic and Evolutionary Computation Conference*, GECCO '14, New York, NY, USA. ACM.

Lipson, H. and Pollack, J. B. (2000). Automatic design and manufacture of robotic lifeforms. *Nature*, 406(6799):974–978.

Miconi, T. (2008). In silicon no one can hear you scream: Evolving fighting creatures. *Genetic Programming*, pages 25–36.

Mitchell, M. (1998). *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, USA.

Pilat, M. L. and Jacob, C. (2010). Evolution of vision capabilities in embodied virtual creatures. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, GECCO '10, pages 95–102, New York, NY, USA. ACM.

Selfridge, O. G. (1958). Pandemonium: a paradigm for learning in Mechanisation of Thought Processes. In *Proceedings of a Symposium Held at the National Physical Laboratory*, pages 513–526, London. HMSO.

Shim, Y. and Kim, C. (2003). Generating flying creatures using body-brain co-evolution. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 276–285. Eurographics Association.

Sims, K. (1994). Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, SIGGRAPH '94, pages 15–22, New York, NY, USA. ACM.

Urzelai, J., Floreano, D., Dorigo, M., and Colombetti, M. (1998). Incremental robot shaping. *Connection Science*, 10:341–360.