

Soft-Body Muscles for Evolved Virtual Creatures: The Next Step on a Bio-Mimetic Path to Meaningful Morphological Complexity

Dan Lessin Sebastian Risi

IT University of Copenhagen
Copenhagen, Denmark
dles@itu.dk, sebr@itu.dk

Abstract

In the past, evolved virtual creatures (EVCs) have been developed with rigid, articulated bodies, and with soft bodies, but never before with a combination of the two. In nature, however, creatures combining a rigid skeleton and non-rigid muscles are some of the most complex and successful examples of life on earth. Now, for the first time, creatures with fully evolved rigid-body skeletons and soft-body muscles can be developed in the virtual world, as well. By exploiting and re-purposing the capabilities of existing soft-body simulation systems, we can evolve complex and effective simulated muscles, able to drive a rigid-body skeleton. In this way, we can begin to bridge the gap between articulated and soft-bodied EVCs, and take the next step on a nature-inspired path to meaningful morphological complexity for evolved virtual creatures.

Introduction

Since evolved virtual creatures (EVCs) were first introduced (Sims, 1994b), the standard implementation has employed a rigid-segmented skeleton-like body, with only minor variations. There have been some recent investigations into rigid-segmented bodies with more complex segment forms (Auerbach and Bongard, 2012), and soft-bodied creatures have also produced compelling results (Rieffel, 2013; Cheney et al., 2013), but no attempt has yet been made to combine the two approaches. Evolution in the natural world has employed such a combination to produce a great variety of highly complex and successful creatures (Figure 1b). What might emerge when this method of embodiment is reproduced from life-as-we-know-it in the real world to life-as-it-could-be in the virtual world?

To begin this pursuit, appropriate simulated muscles are required. Much of the previous related work has applied simulated muscles to the more limited case of fixed-morphology creatures. These include inflated-cloth muscles on simulated robots (Glette and Hovin, 2010), muscle-inspired joint drives for simulated swimming robots (Moore and McKinley, 2014), and relatively complex spring-like muscles used to animate human-designed morphologies for entertainment purposes (Geijtenbeek et al., 2013). In other

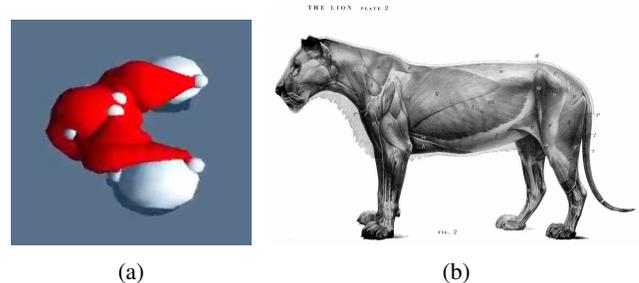


Figure 1: The inspiration for this paper’s approach (a) to morphological complexity is real world creatures (b) in which a rigid skeleton is driven by muscles. See results in motion at <http://goo.gl/rvSvFv>.

work, a simple yet fully adaptable form of evolvable musculature has been explored in which linear-spring-like muscles are evolved along with skeletal segments and control to produce true EVCs with some of the potential benefits of soft-body muscle actuation (Lessin et al., 2014b).

In this paper, a new approach is introduced in which an existing soft-body simulation system is re-purposed and exploited to produce realistic muscles with many desirable properties for the development of complex and capable EVC bodies (Figure 1a). In this way, evolved virtual creatures are produced which for the first time combine a traditional articulated skeleton with soft-body morphological elements in a new step on a nature-inspired path toward the type of morphological complexity that has proven so useful in the natural world.

Background

The primary contribution of this paper is a novel combination of traditional articulated-skeleton EVCs and soft-body simulation. This section provides a brief overview of these foundational technologies.

First, it is useful to define previous EVC systems. Classical EVCs were established in the initial work by Sims (1994b), then applied without major changes in a number of subsequent publications (Chaumont et al., 2007; Miconi,

2008; Lehman and Stanley, 2011). In these systems, the articulated rigid-body system is defined by a tree-like graph genotype, which is traversed to produce instructions for expressing the phenotype. The phenotype is typically composed of boxes or other rigid-body-simulation primitives (such as spheres and capsules), connected by a variety of joint types (such as revolute, spherical, prismatic, or cylindrical).

In the original implementation and most that followed, actuation was provided by implicit joint drives at every degree of freedom of every joint. In more recent work (Lessin et al., 2013, 2014b,a), these were replaced by linear-spring muscle activation, which is the implementation from which this paper’s method is most directly derived.

While a number of implementations have employed neural networks for control (Miconi, 2008; Lehman and Stanley, 2011), Sims’ original EVCs (Sims, 1994b,a) were controlled using a network of simple computing nodes (e.g., sinusoidal, sum, product, derivative), and that method is continued in the work of this paper. (See Figure 7b in the Results section for an example.)

It is also important to note previous work with soft bodies in EVCs. While they have never before been used in EVCs in combination with an articulated rigid skeleton, soft bodies have been applied to great effect in recent work to produce morphologically complex locomoting creatures, including some that use a combination of hard and soft body elements (though notably with the hard elements mixed in among the soft ones, not with any kind of articulated skeleton). Note that while those previous soft-body implementations employed single-purpose voxel-based soft body simulators (Cheney et al., 2013) or soft-body simulation alone (Rieffel, 2013), the work presented in this paper uses an adaptation of soft bodies within an integrated soft-and-rigid-body simulation system. This is necessary to permit the key combination of articulated-skeleton and soft-body simulation which is at the heart of this system’s novel contribution.

Approach

To take this next step in nature-inspired morphological complexity for evolved virtual creatures, previous work on EVCs with linear-spring muscles will be combined with a novel soft-body muscle implementation. Beyond the new soft-body elements, most of the important details of the system (e.g.: joint and segment types, evolutionary algorithm, and genotype encoding and expression) are as defined in those earlier implementations.

In (Lessin et al., 2013, 2014b,a), a degree of morphological complexity was added through the use of a very simple approximation to a natural muscle: a linear spring. In that work, the traditional joint-motor drives employed by most EVCs were removed, and linear springs were allowed to evolve between skeletal segments. For each joint, springs

could be added or removed by evolution, and their attachment points and strength were also evolvable. Brain activations applied to those muscles modified their underlying spring constant, increasing or decreasing the force they exerted on their attached body segments.

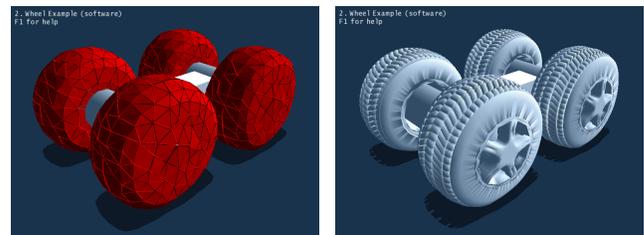
That work is the basis for the new system described here. Producing an appropriate soft-body muscle implementation is a significant challenge, but once that is achieved, it can be directly adapted to the linear-spring-muscle system just described. The system can function largely as before, simply constructing a full soft-body muscle in place of the original linear spring, using the given attachment points and strength. In the next section, the novel implementation of an appropriate soft-body-muscle system is described in detail.

Novel Soft-Body Muscle Implementation

The presented soft-body EVC musculature has a number of desired characteristics, each one obtained by employing or re-purposing existing capabilities available in an off-the-shelf physical simulation system. In this section, each of these elements of the system is described in detail.

Soft-Body Simulation

The fundamental soft-body system used is available as a part of NVIDIA PhysX, but not in most versions, and (at the time of writing) not in recent versions. (The work presented here was implemented with SDK version 2.8.1 and PhysX System Software 9.11.0621.) This system offers well-established abilities to simulate rigid bodies and joints, and can do so in combination with the soft-body simulations which are the focus of this work.

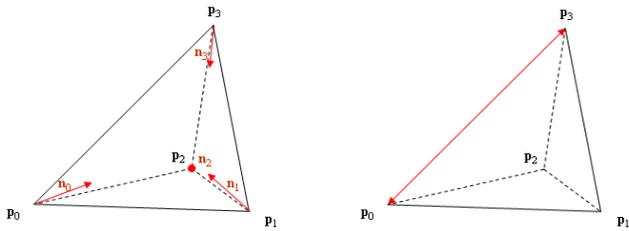


(a)

(b)

Figure 2: PhysX’s soft-body simulation is implemented using a tetrahedral mesh (a) which can be used to drive an accompanying triangular mesh for rendering (b).

In this system, soft bodies are simulated as tetrahedral meshes (Figure 2a), with vertices simulated as particles, and additional constraints applied per tetrahedron to produce some of the more complicated effects described in the following subsections. For rendering, the simulated tetrahedral mesh can be used to deform a corresponding triangle mesh of arbitrary complexity (Figure 2b).



(a) Tetrahedron vertices are moved to preserve volume, when *volume stiffness* is high. (b) Tetrahedron vertices are moved to preserve edge lengths, when *stretching stiffness* is high.

Figure 3: PhysX volume (a) and stretching (b) constraints, used to preserve muscle volume and implement muscle contraction, respectively.

Volume Preservation

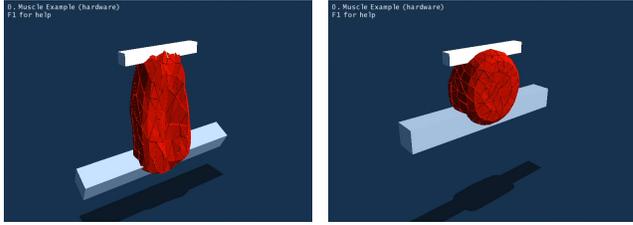
One desired characteristic of soft body muscles is the preservation of volume, which produces a familiar squash-and-stretch behavior as muscles decrease and increase in length. In PhysX, this property (which they call *volume stiffness*) is directly available as an attribute of simulated soft bodies, and is implemented as restorative movements applied to tetrahedral vertices (Figure 3a). This parameter may be varied to produce a range of soft-body styles. For this implementation, it was set to its maximum value: 1.0.

Contraction Mechanism

A vital aspect of simulated muscles is the ability to contract in a controllable manner. While PhysX provides no built-in mechanism for this, it does offer the *stretching stiffness* attribute of soft bodies, which can be repurposed to accomplish this goal. This property controls the degree to which a soft body will attempt to retain its original shape (again by changing tetrahedron-vertex positions), as illustrated in Figure 3b. By varying this attribute over time, the muscle’s tendency to return to its initial, fully contracted shape is controlled, and it can be made to relax or contract as needed. In figure 4, the ability of this contraction mechanism to lift a rigid-body object is demonstrated.

Strength that Varies with Cross Section

Variations in muscle thickness add meaningful morphological complexity to animal musculature. To have the same kind of variation occur in a meaningful way in these virtual creatures, muscle strength can be made to depend on the muscle’s cross-sectional area. A similar relationship was used in Sims’ original work (Sims, 1994b), in which joint-motor strength varied in proportion to parent-segment cross section. With the soft-body simulations used for the muscles described here, this strength increases as a natural side-effect of the scaling process. By simply modifying the initial vertex positions of the tetrahedral mesh to give it thicker por-



(a) Low stretching stiffness; muscle relaxed. (b) High stretching stiffness; muscle contracted.

Figure 4: This figure illustrates the underlying contraction mechanism employed by the soft-body muscles in this paper. As stretching stiffness is increased, the simulated muscle changes from its relaxed shape (a) to its original, contracted shape (b), and is able to do useful work, such as lifting the rigid-body segment shown here.

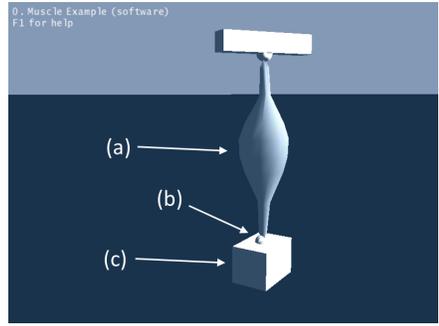


Figure 5: The tuberosity mechanism which allows muscles to realistically interact with adjacent skeletal segments: A muscle (a) is attached first to a tuberosity (b), which is in turn attached to the target skeletal segment (c).

portions, the muscle is able to pull with greater force when simulated by PhysX.

Ability to Wrap Around Skeletal Segments

The fact that muscles wrap around skeletal segments in real-world bodies is also an important contributor to the complexity of animal musculature, but in the virtual world, such behavior is not a given. When soft bodies are used in PhysX, for example, any rigid body shape that they are attached to is prohibited from producing collision interactions with them. To get around this limitation, a small additional amount of bio-mimetic complexity is added during the muscle attachment process.

Instead of attaching simulated muscles directly to the rigid-body segments that they will act upon, each end of the muscle is first attached to a small sphere, which is then attached to the target segment (Figure 5). We refer to these spheres as *tuberosities* for their similarity to the eponymous structures in animal skeletons. With this configuration, although the virtual muscles are prohibited from colliding

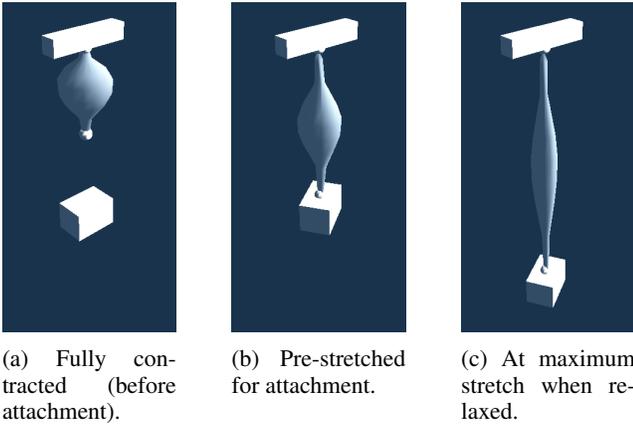


Figure 6: Muscle pre-stretch during construction. Muscles are instantiated in the fully contracted state (a), then pre-stretched to double their length for attachment to skeletal segments (b). During use, they may be contracted toward their original state, or further stretched (c) when relaxed.

with the tuberosity spheres themselves, collisions with all other skeleton segments occur as normal.

Muscle Pre-Stretch

The contraction mechanism described above requires the simulated muscles to be initially instantiated in their fully contracted state. To produce a neutral length during body construction, muscles are *pre-stretched* in the attachment process. Initially, muscles are created to span only half of the required distance between their attachment points (Figure 6a). Because of this, the attachment process serves to pre-stretch the muscle to an intermediate length (Figure 6b). From this attached pose, the muscle can be either contracted toward its initial shape, or further relaxed to its maximum length of approximately double the neutral length (a limit imposed upon soft bodies by PhysX), as shown in Figure 6c.

Experiments

Using the techniques described in the previous sections, a series of experiments were performed to evaluate this new system’s potential to produce locomoting EVCs which combine for the first time an articulated skeleton with soft-body muscles.

Experimental Setup

For these experiments, eight independent runs were started simultaneously, each on its own processor in a relatively up-to-date laptop machine. Each run was started with a different random seed, and was allowed to run for approximately two weeks, using a population size of 50. In that time (apparently due to execution speeds which were highly dependent on creature complexity), the various runs completed between 180 and 893 generations. The simulation required

for each fitness evaluation consisted of 1.5 seconds with the skeleton frozen (for muscles to recover from pre-stretch), up to 4 seconds for the body to settle (with the brain inactive), followed by the 6 seconds of simulation actually used to produce the fitness score.

Prevention of Cheating

Before presenting the results, it is important to describe an insidious form of cheating that was encountered, as well as what was done to prevent it. Many early tests fell victim to a particularly easy physics cheat apparently made possible by the muscles’ soft-body simulation. It seems to be the case that with a sufficiently strong constant contraction of muscles, the sum of small simulation inaccuracies (made worse by the number of vertices simulated in these soft-body muscles) was sufficient to produce a physically unrealistic overall force on the creature. This could be used to generate movement along the ground with no physically realistic cause. Such solutions are apparently so prevalent in the solution space, that they completely excluded any non-cheating results when fitness was based solely on locomotion.

To overcome this cheating so that valid locomotion results could be discovered, an additional requirement was added: Each muscle should significantly change its length throughout the entire fitness evaluation, making the previous continuous, strong-contraction cheating technique impossible. Specifically, each muscle was required to change its length by 25% during each 0.5 seconds of evaluation. During speed evaluation, this is strictly enforced, with zero fitness assigned to any individual not sufficiently achieving this goal. But it is also not trivial to produce a creature in which this requirement is accomplished. This challenge is met by introducing a two-stage evolutionary process, with the first stage providing a fitness gradient toward the muscle-length-change goal, and the second stage rewarding locomotion among creatures which continue to meet that initial requirement.

For stage one, fitness is calculated as the average across all muscles and all 0.5-second evaluation periods of how well the 25% length variation is achieved. So, if E is the set of evaluation periods, M is the set of the creature’s muscles, and l_{min} and l_{max} are a muscle’s minimum and maximum length during an evaluation period, stage-one fitness f_1 can be computed as

$$f_1 = \frac{1}{|E| |M|} \sum_{E, M} \max\left(1, \frac{l_{max} - l_{min}}{l_{max}} \cdot \frac{1}{0.25}\right).$$

(Note: The $\max()$ term ensures that the required fitness is developed for all muscles in all evaluation periods. Without it, the same overall f_1 fitness might be achieved with some muscles having more than the required minimum and others having less, or none.)

Once a sufficient fraction of the population (5% was used in these experiments) has achieved full fitness using f_1 , the population is refilled completely with copies of those f_1 -fit individuals, and stage two begins.

In stage two, in addition to f_1 , speed is also evaluated, and used with f_1 as described earlier. If full f_1 fitness is maintained (or nearly so), the new fitness f_2 is simply computed from speed. But if f_1 fitness falls too low, f_2 will be assigned a zero value. If s is the average speed during the evaluation, and s_{max} is the maximum target speed (20 m/s for these experiments), f_2 can be computed as

$$f_2 = \begin{cases} \frac{s}{s_{max}} & \text{if } f_1 > 0.9 \\ 0 & \text{otherwise} \end{cases}$$

Results

Of the eight experimental runs, four were successful with respect to fitness score, producing stage-two scores associated with reasonable-looking locomotion. One of those four succeeded by exploiting a weakness in the fitness definition rather than by producing true locomotion. The remaining three successful runs (runs 3, 4, and 6) are presented here as Evolved Creatures 1 through 3. To see these results in motion, please visit <http://goo.gl/rvSvFv>.

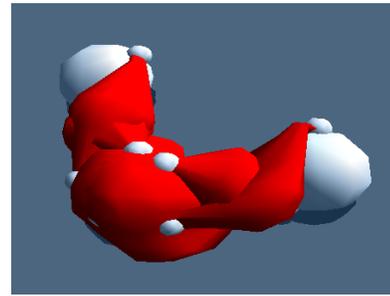
Evolved Creature 1: Crawler

In this section, the first example of a successful result (run 3, generation 172) is presented. Its body and brain are shown in Figure 7. This creature's rigid skeleton consists of a central sphere attached to spherical limbs by prismatic joints (able to telescope linearly, while maintaining orientation). Each limb is connected to the root segment by three muscles, with varying attachment points and strengths.

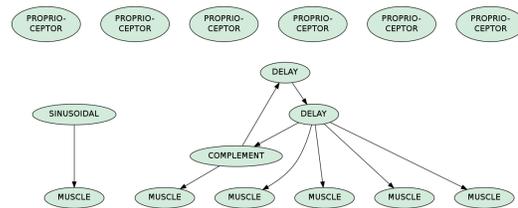
This creature's brain shows that it employs open-loop control—it ignores input from its proprioceptors (sensors indicating muscle length) when computing the output signals to the muscles. Also notable in this brain: One muscle is driven directly by a sinusoidal signal, while the rest of the brain nodes have been employed to create an improvised square-wave generator which drives the remaining muscles.

As is true for all of the results presented here, this creature's morphology makes use of bilateral symmetry, accessible to evolution as a single boolean attribute. Its method of locomotion is to move forward (toward its limbs) by quickly extending and retracting its arms in an alternating manner, producing a clawing, crawling gait (Figure 9).

The development of this creature's fitness is shown in Figure 8. Generations in stage one (learning to keep muscle lengths changing) are shown on the left of the graph in red, with the generations spent in stage two (learning locomotion while maintaining muscle movement) shown on the right of the graph in green. In this creature, the ability to change all muscle lengths and complete the first stage was achieved



(a) The body of Evolved Creature 1, consisting of a root sphere, connected to two sphere limbs by prismatic joints. Each limb is connected to the root with three muscles. From this viewing angle, this creature's locomotion direction would be up and to the right.



(b) The brain of Evolved Creature 1.

Figure 7: The body and brain of the first example of a successful result from this system. This creature was produced at generation 172 of evolutionary run 3.

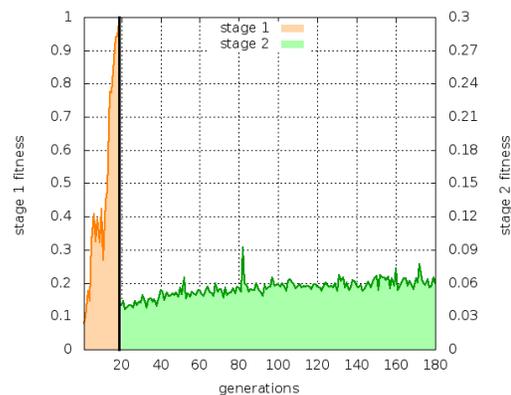


Figure 8: Fitness graphed over time for the evolutionary run producing Evolved Creature 1. The graph shows the development of stage-one fitness on the left in orange (in which muscle-length changes are rewarded) and stage-two fitness on the right in green (in which locomotion is rewarded, and muscle-length changes are required).

very quickly, followed by slow, steady optimization of control with no obvious morphological changes.

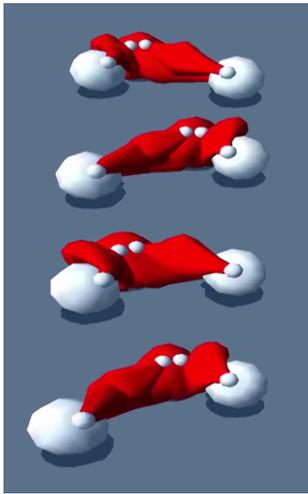


Figure 9: The locomotion method of Evolved Creature 1. Viewing the images from top to bottom, the creature is seen to shift its weight from side to side, extending alternating limbs forward (toward the viewer) to produce a crawling gait in that direction.

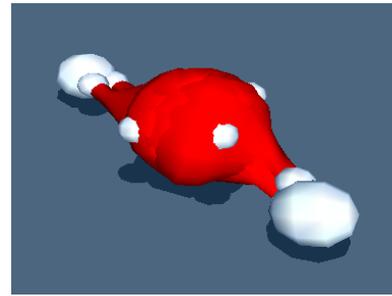
Evolved Creature 2: Inchworm

The next successful result presented (run 4, generation 518) is discussed in this section. Its body and brain are illustrated in Figure 10. Its skeleton consists of a large central sphere, joined to two smaller spheres by spherical (ball-and-socket) joints. The central segment is connected to each of its limbs by two muscles. In contrast to the previous creature, this creature's brain is harder to interpret. It is notably different, however, in that it appears to employ some closed-loop control, having connections from its proprioceptive sensors into the rest of the control network.

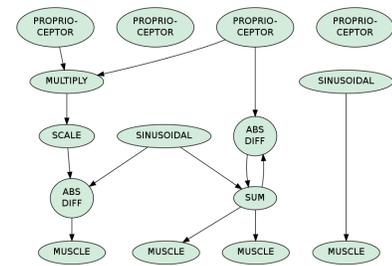
By aggressively swinging its central segment toward one of the limbs, then more gently returning it towards the other, this creature locomotes along the direction of its body length (Figure 12), producing the fastest creature by far among the results presented in this paper. As before, both stages of this creature's fitness development are illustrated in Figure 11. In contrast to the first creature, this creature spent many more generations developing full stage-one fitness (rewarding muscle-length change), and its stage-two fitness makes significant jumps after appearing to level off on two separate occasions.

Evolved Creature 3: Shuffler

In this section, the final successful result (run 6, generation 348) is presented, with body and brain illustrated in Figure 13. This creature's skeleton is composed of a central capsule segment, with capsule limbs extending out and to the back. The limbs are connected with cylindrical joints, which allow both telescoping along the joint's axis and rolling around it, as well. Each limb is connected to the



(a) The body of Evolved Creature 2, consisting of a central sphere, with two smaller spheres attached as limbs by ball-and-socket joints. Each limb is connected to the root segment with two muscles. From this viewing angle, the creature's locomotion direction is out and down, to the right, along the length of its body.



(b) The brain of Evolved Creature 2.

Figure 10: The body and brain of the second example of a successful result from this system. This creature was produced at generation 518 of evolutionary run 4.

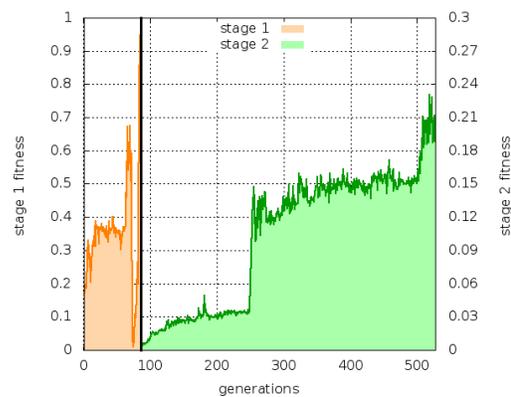


Figure 11: Fitness graphed over time for the evolutionary run producing Evolved Creature 2.

central segment by two muscles. By shifting its weight from side to side, forward (away from the limbs) shuffling locomotion is produced. The two stages of fitness development are shown in Figure 14. In contrast to the two previous crea-

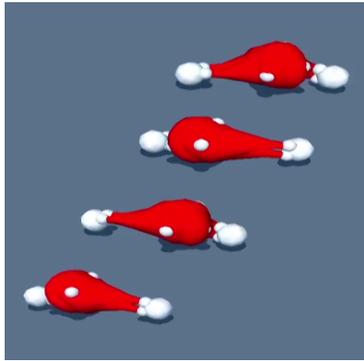
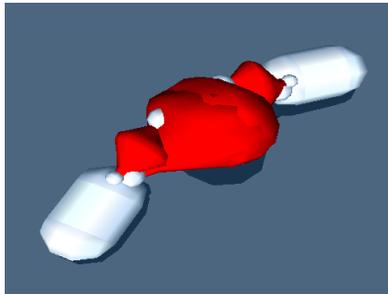
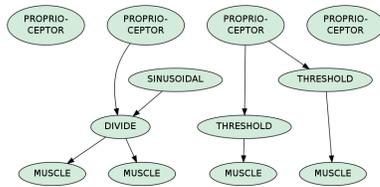


Figure 12: The locomotive technique of Evolved Creature 2. Considering the images from top to bottom, the creature first draws its weight back (to the right in this view), then aggressively throws it forward (left in this view) producing an inchworm-like gait along the direction of the length of its body.



(a) The body of Evolved Creature 3, with a central capsule and two capsule limbs, joined by cylindrical joints. Each limb is connected to the root by two muscles. From this viewing angle, the creature's direction of locomotion would be up and back, to the left.



(b) The brain of Evolved Creature 3.

Figure 13: The body and brain of the third example of a successful result from this system. This creature was produced at generation 348 of evolutionary run 6.

tures shown, this creature spent over half of its simulation time developing stage-one fitness.

Results Summary

The experiments presented here have demonstrated that this paper's novel soft-body muscle system is indeed sufficient

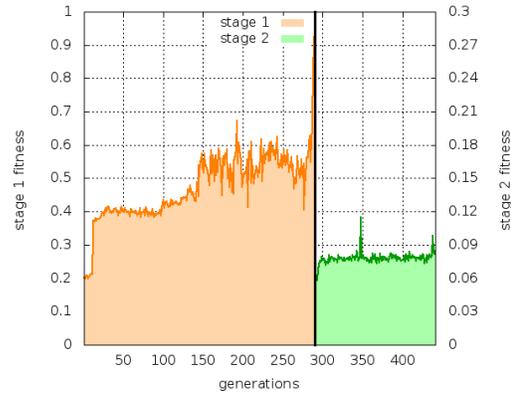


Figure 14: Fitness graphed over time for the evolutionary run producing Evolved Creature 3.

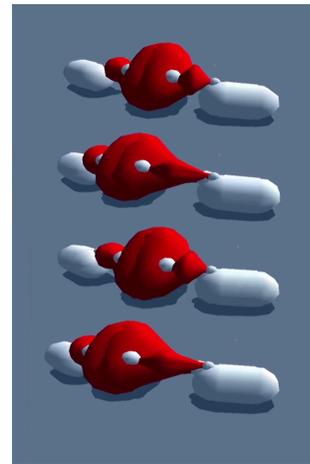


Figure 15: The locomotive technique of Evolved Creature 3. From top to bottom, the images show how the creature shifts its weight from side to side, resulting in a forward-sliding, shuffling gait. From this viewing angle, the locomotion direction is out, down, and slightly to the left.

to allow, for the first time, the evolution of virtual creatures which combine a rigid-body skeleton with complex non-rigid muscles. These evolved results demonstrated multiple locomotion solutions, employing a variety of segment types, joint types, control strategies, and gaits.

Discussion and Future Work

One important issue to discuss is the fact that the anti-cheating mechanism employed (requiring muscle-length changes) means that some otherwise valid solutions may never be evolved. This was an acceptable trade-off for these initial experiments, but in the future, a more accurate physical simulator or different method of preventing cheating might allow this restriction to be removed.

One observation about the results presented here is that,

despite the demonstrated variety in body structure, brain structure, and gaits, all of these successful creatures share a similar body plan, with a limited number of segments in somewhat similar configurations. Future work might determine the cause for this or get beyond it, perhaps through the use of a diversity-promoting mechanism (Lehman and Stanley, 2011).

In the larger picture, the work presented here is part of a potentially rewarding path to the kind of rich morphological complexity seen in creatures in the natural world. Now that bodies can be evolved with functional muscles which contribute in a meaningful way to morphology, one clear next step would be the addition of simulated skin. (In fact, this was even proposed in Sims' original work). The same kinds of physical simulation systems that provide soft-body simulation can also simulate cloth, which might be an ideal way to achieve this next layer of bio-mimetic realism in virtual creatures.

Another compelling topic on the same path might be the evolution of bone shapes, as enabled by the ability to effectively evolve three-dimensional forms (Clune and Lipson, 2011). Not only would this allow richer morphological interactions with muscles and skin, but it might even permit the development of emergent joints. Instead of using implicit, externally enforced relationships to define the relative movement of connected segments, evolvable bone shapes might allow these relationships to emerge naturally from the interactions between bones, muscles, and skin. Additionally, evolvable bone shapes might permit the development of creatures with exoskeletons, which may well emerge naturally through this process as long as they are not explicitly disallowed.

Conclusion

This paper has presented a technique for combining rigid articulated skeletons with soft-body muscles in evolved virtual creatures for the first time. This was made possible by a novel combination and re-purposing of a number of existing physical simulation components from an off-the-shelf simulation system, as detailed above. Initial experiments with evolving morphology and control for locomotion were shown, along with a new method for counteracting evolution's destructive exploitation of this simulator's inaccuracies. These experiments had a relatively high rate of success in producing locomoting creatures, and a number of useful results were presented. This novel bio-mimetic synthesis of two highly successful EVC techniques represents a new step toward matching the morphological complexity of some of the most successful creatures in the natural world.

References

Auerbach, J. E. and Bongard, J. C. (2012). On the relationship between environmental and morphological complexity in evolved robots. In *Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation*

Conference, GECCO '12, pages 521–528, New York, NY, USA. ACM.

Chaumont, N., Egli, R., and Adami, C. (2007). Evolving virtual creatures and catapults. *Artificial Life*, 13(2):139–157.

Cheney, N., MacCurdy, R., Clune, J., and Lipson, H. (2013). Unshackling evolution: Evolving soft robots with multiple materials and a powerful generative encoding. In *Proceeding of the Fifteenth Annual Conference on Genetic and Evolutionary Computation Conference*, GECCO '13, pages 167–174, New York, NY, USA. ACM.

Clune, J. and Lipson, H. (2011). Evolving 3d objects with a generative encoding inspired by developmental biology. *ACM SIGEVOLUTION*, 5(4):2–12.

Geijtenbeek, T., van de Panne, M., and van der Stappen, A. F. (2013). Flexible muscle-based locomotion for bipedal creatures. *ACM Trans. Graph.*, 32(6):206:1–206:11.

Glette, K. and Hovin, M. (2010). Evolution of artificial muscle-based robotic locomotion in PhysX. In *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pages 1114–1119. IEEE.

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. (2014a). Adapting morphology to multiple tasks in evolved virtual creatures. In *Proceedings of The Fourteenth International Conference on the Synthesis and Simulation of Living Systems (ALIFE 14) 2014*.

Lessin, D., Fussell, D., and Miikkulainen, R. (2014b). Trading control intelligence for physical intelligence: Muscle drives in evolved virtual creatures. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2014*.

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

Moore, J. M. and McKinley, P. K. (2014). Evolving joint-level control with digital muscles. In *Proceedings of the 2014 Conference on Genetic and Evolutionary Computation*, GECCO '14, pages 209–216, New York, NY, USA. ACM.

Rieffel, J. (2013). Heterochronic scaling of developmental durations in evolved soft robots. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*, GECCO '13, pages 743–750, New York, NY, USA. ACM.

Sims, K. (1994a). Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372.

Sims, K. (1994b). 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.