

Trading Control Intelligence for Physical Intelligence: Muscle Drives in Evolved Virtual Creatures

Dan Lessin
Dept. of Computer Science
The Univ. of Texas at Austin
Austin, Texas, USA
dlessin@cs.utexas.edu

Don Fussell
Dept. of Computer Science
The Univ. of Texas at Austin
Austin, Texas, USA
fussell@cs.utexas.edu

Risto Miikkulainen
Dept. of Computer Science
The Univ. of Texas at Austin
Austin, Texas, USA
risto@cs.utexas.edu

ABSTRACT

Traditional evolved virtual creatures [12] are actuated using unevolved, uniform, invisible drives at joints between rigid segments. In contrast, this paper shows how such conventional actuators can be replaced by evolvable muscle drives that are a part of the creature's physical structure. Such a muscle-drive system replaces control intelligence with meaningful morphological complexity. For instance, the experiments in this paper show that control intelligence sufficient for locomotion or jumping can be moved almost entirely from the brain into the musculature of evolved virtual creatures.

This design is important for two reasons: First, the control intelligence is made visible in the purposeful development of muscle density, orientation, attachment points, and size. Second, the complexity that needs to be evolved for the brain to control the actuators is reduced, and in some cases can be essentially eliminated, thus freeing brain power for higher-level functions. Such designs may thus make it possible to create more complex behavior than would otherwise be achievable.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning—*connectionism and neural nets*; I.6.8 [Simulation and Modeling]: Types of Simulation—*animation*; I.2.1 [Artificial Intelligence]: Applications and Expert Systems—*games*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*animation*

Keywords

evolved virtual creatures; artificial life; muscle drives; physics-based character animation; physical intelligence; morphological complexity

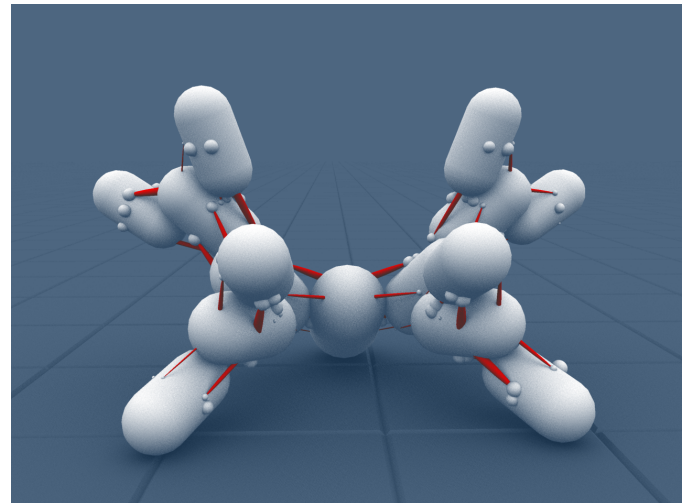


Figure 1: A creature evolved for jumping (Section 4.2.5) using the method described in this paper, demonstrating the morphological complexity that results from replacing implicit joint-motor drives with an evolvable musculature. As with all other examples in this work, the physical intelligence embodied by these muscle drives enables this creature to perform a useful task essentially without control intelligence (Section 3.2). Video of this and all other results can be viewed at <http://youtu.be/csZ9JZcuBfE>.

1. INTRODUCTION

Morphological complexity is an important goal for evolved virtual creatures (or *EVCs*; Figure 1) [2]. How can it be increased to approach the morphological complexity of creatures evolved in the real world? Traditional segmented EVCs [12, 3, 10, 8] achieve some measure of complexity through the placement, dimensions, and types of their rigid segments and joints. More recently, creatures with morphology based on implicit definitions such as CPPNs and gene regulatory networks [2, 6, 4] have demonstrated a different—and arguably greater morphological complexity, albeit based on indirect developmental mechanisms. In contrast, this paper demonstrates that it is possible to increase the complexity of the rigid-bodied model directly by employing a more advanced approach to actuation.

In a conventional EVC, actuation is provided by implicit joint motors. Such motors are completely uniform, i.e., present at every free axis of every joint, and fixed over time.

They are also typically unseen, perhaps because their uniformity would provide little reward for making them visible.

However, our recent work [9] has demonstrated that EVCs can also be successfully actuated by a simple form of simulated muscle—a variable-strength linear spring attached to two segments across a joint. Although, in that implementation, the muscles were controlled by a complex brain, one particularly interesting property of such drives is that they do not actually require this control complexity. As will be shown in this paper, these muscles are able to embody and replace a significant portion of the control intelligence that would normally be provided by the creature’s brain. In fact, creatures that are almost entirely without control intelligence can still develop sufficient *physical* intelligence (in the form of their evolved musculature) to perform rudimentary, yet useful tasks, such as jumping and simple locomotion.

One particularly beneficial result of this shifting of intelligence from brain to body is the fact that, where the control intelligence was externally invisible, the physical intelligence that replaces it is visible in the morphological complexity of the muscles. And although the muscle-drive model described in this work is in some ways simple, it nevertheless communicates meaningful complexity through its evolved characteristics: the density or lack of muscles at a joint, their size (with rendered thickness indicating strength), orientation (indicating direction of force application), and their attachment points.

In addition, this effective obviation of control logic could ultimately prove useful as an evolutionary robotics application if mechanisms similar to the simple simulated muscles used here can be implemented in the real world, especially where control intelligence is at a premium. Robots that need to be particularly small, for example, might benefit from replacing a relatively complex controller with a properly evolved actuator musculature.

Even traditional EVC applications could benefit from this easing of the demands on brainpower. By removing the cognitive load that can now be borne by the muscle drives, this implementation frees the brain to devote equivalent computational power to achieving more complex behavioral goals.

In the balance of this paper, this new step on the biologically inspired path to meaningful morphological complexity and reduced control loads is described, and successful results are presented and evaluated.

2. BASIC EVC SYSTEM

The basic EVC system described in this paper is in large part built on a traditional EVC framework [12, 3, 7, 10]. This section briefly sets out the components of this system, which—while not the primary focus of this paper—are nevertheless fundamental.

2.1 Evolutionary Algorithm

The specifics of the evolutionary algorithm are largely conventional, making use of elitism, fitness-proportionate selection, and rank selection [11]. In addition, some degree of shaping [13] is employed, as described in Sections 4.2 and 4.3. Fitness is evaluated in a physically simulated virtual environment implemented with NVIDIA PhysX.

New individuals are created through crossover followed by mutation. The directed graph genotypes used (see Section 2.2) are sufficiently tree-like that crossover can be implemented by choosing a random point in each of the parent

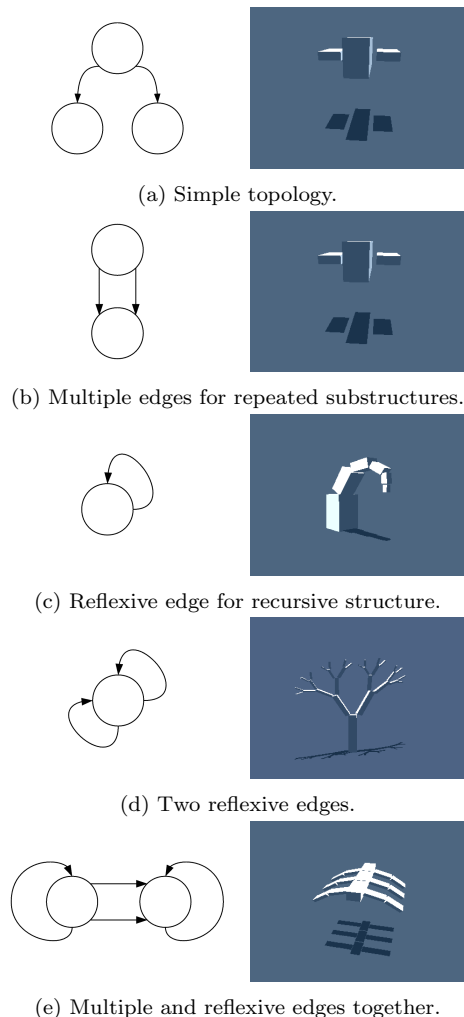


Figure 2: Hand-designed genotype/phenotype pairs (after those of Sims [12]) demonstrate the encoding power of EVC systems.

graphs and swapping sub-trees at those points. This produces two child graphs, one of which is randomly discarded. Mutation mechanisms and probabilities vary by attribute type.

2.2 Morphology

As with most EVCs, 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, as illustrated in Figure 2. In addition, reflection of body parts as well as body symmetry are made readily available to evolution.

In the implementation of EVCs in this paper, all PhysX primitives are made 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. (Note that,

consistent with this paper’s focus on muscle drives, the joints themselves are omitted from all renderings.) 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.

3. EVOLVABLE MUSCULATURE

The evolvable musculature described in this section makes it possible to transfer control intelligence to physical intelligence, which is the focus of this paper. First, the implementation of the muscle drives themselves is described, then the extremely minimal control that it enables is specified.

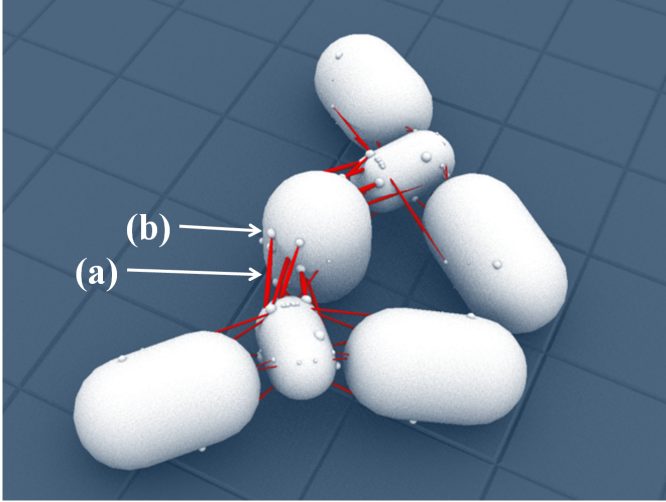


Figure 3: Evolvable musculature, with example muscle body (a) and attachment point (b) indicated. The density of muscles at a joint, their thickness (indicating strength and activation), orientation, and attachment points all contribute meaningfully to the creature’s morphological complexity.


3.1 Muscle Drives

The muscle drives are implemented as simple linear springs. Each muscle (Figure 3(a)) is completely described by its attachment points and maximum strength. An attachment point (Figure 3(b)) may be placed anywhere on a rigid body segment, and each pair of attachment points must exist across a joint connecting two such segments. Muscles may be added and removed by evolution, and their attachment points and maximum strength are evolvable. During simulation, a muscle’s activation (in $[0,1]$) determines what portion of its maximum strength that muscle will apply.


The muscle is implemented using a standard PhysX joint called a *distance joint*, modifying its attributes so that it acts as a simple linear spring. A PhysX distance joint allows the specification of a maximum distance between two attachment points, and this maximum is enforced by spring-like behavior when exceeded. By setting the distance joint’s maximum distance to zero, only the spring-like enforcements are applied. The spring constant is adjusted during simulation to reflect the tension that results from combining the muscle’s activation with its maximum force value. Note that

this implementation—with numerous joints of varying types between a single pair of rigid body segments—is probably not typical for PhysX, and initial results with normal settings resulted in simulations that were not sufficiently stable. All experiments presented in this paper rely on a much smaller simulation step— $1/240$ th of a second—as well as other configuration settings, all of which can significantly affect simulation efficiency.

From the three evolvable properties of each muscle (two attachment points and the maximum strength), as well as the fact that muscles may be added or removed at any joint, a great degree of visually obvious meaningful morphological complexity emerges. This design also has the potential to embody sufficient physical intelligence to perform basic behaviors with only a trivial degree of control intelligence required, as will be described next.



(a) Jump activation (single square pulse).



(b) Locomotion activation (repeating square wave).

Figure 4: With the muscle drives’ capacity for physical intelligence, simple but useful behaviors can be performed effectively without control intelligence. Here, the fixed global muscle activations that replace the typical EVC’s relatively complex brain for all experiments in this work are illustrated.

3.2 Minimal Control

As a demonstration of approximately how much physical intelligence the evolvable musculature can embody, the examples in this paper all function almost entirely without control intelligence (Figure 4).

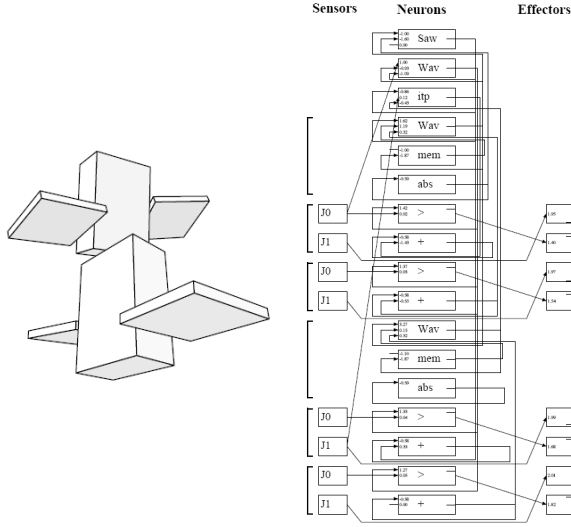


Figure 5: For comparison, this figure illustrates a body and brain from one of Sims’ conventional EVCs—this one evolved for locomotion [12]. Note the different balance of control and physical intelligence, with greater complexity hidden in the brain, while the morphology is less elaborate.

In a conventional evolved virtual creature, control intelligence is implemented as a neural network [10] or a directed graph of simple computing nodes [12], as shown in Figure 5. In the creatures in this paper, far less is required. For these new creatures, the typical brain is replaced by a single activation function, which is applied to all muscles simultaneously. This activation function was arbitrarily chosen and fixed before each experiment began, being neither evolved nor hand-tuned for suitability. These functions (a half-second unit-amplitude square pulse for jumping, and a 1-Hz unit-amplitude square wave for locomotion) are illustrated in Figure 4.

4. EXPERIMENTS

In this section, the results of two experiments are presented, in which creatures evolve body and musculature for the tasks of jumping and locomotion. In each case, the potential for physical intelligence in the muscle drives effectively obviates control intelligence and also demonstrates the muscle drives’ potential to exhibit meaningful morphological complexity. All of the results described here can be seen in motion in the accompanying video¹.

4.1 Experimental Setup

For all of the results in this paper (Sections 4.2 and 4.3), a typical population size was on the order of 100. Between 221 and 500 generations were used to obtain all of the jump results shown, and between 1000 and 2000 generations were used for the locomotion results. In both the locomotion and jumping tasks, 10 copies of the experiment would be executed in parallel, each with its own random seed. The champions of these 10 runs typically form a diverse collection of successful results, although certain variations of mor-

phological themes tend to recur. Illustrative examples are presented in the subsections below.

The variety demonstrated in these results suggests fertile ground for the future development of new, more challenging tasks, as discussed in Section 5.

4.2 Jump Results

The following five subsections illustrate the various solutions found for a simple jumping task. For this skill, fitness is defined using a number of intermediate shaping steps, resulting in a sequence of fitness goals like the one below. Each stage is complete when a sufficient fraction of the population—on the order of 5%—has achieved full fitness. At that point, that percentage having full fitness is replicated to fill a new population, and evolution continues in the next stage.

A useful concept in defining these goals is the axis-aligned bounding box (AABB)—particularly its top and bottom, which describe the creature’s highest and lowest extents. Both static (i.e., at rest) and highest (as measured throughout a single fitness evaluation) AABB measures are employed.

1. **static AABB top**
(to encourage a static size requiring multiple segments)
2. **static AABB top + highest AABB top**
(to encourage muscles and upward action, while maintaining static size)
3. **highest AABB top + highest AABB bottom**
(to encourage ground clearance, while maintaining upward action)

For all jump evolution experiments, control consists solely of the single fixed global activation signal depicted in Figure 4(a), with all other required intelligence residing entirely within the body, including the evolvable musculature.

In each result illustration (Figures 6-10), the left and right sides show the creature before and during its jump, respectively.

4.2.1 Jump Result 1




Figure 6: Two-armed swing (repeatable).

The creature in Figure 6 adapts its morphology to the given minimal control signal by developing heavy arms that are swung up by appropriately placed muscles. The upward momentum of these limbs is then sufficient to make the creature airborne.

Foreshadowing a common technique observed in the locomotion results, this creature happens to end its jump in the same configuration from which it began, demonstrating a potential for repeated action.

¹<http://youtu.be/csZ9JZcuBfE>

4.2.2 Jump Result 2

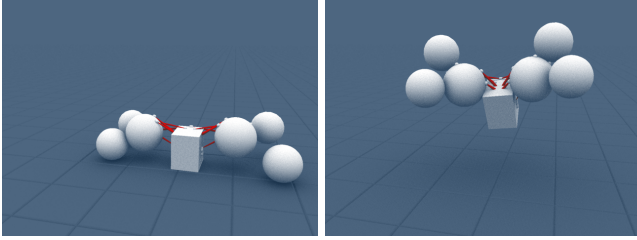


Figure 7: Two-armed swing (non-repeatable).

The creature in Figure 7 applies the same basic limb-swinging strategy to a different morphology, resulting in a strong jump that does not happen to end in the same pose from which it began.

4.2.3 Jump Result 3




Figure 8: One-armed swing.

The strategy in Figure 8 is similar to that of the previous two, but it works here with a single limb instead of a symmetrical pair of limbs. As with Jump Result 1, this creature's consistent begin and end poses foreshadow the successful technique seen in the locomotion results—in this case matching almost exactly the morphology and behavior of Locomotion Result 2.

4.2.4 Jump Result 4




Figure 9: Four-legged push.

The creature in Figure 9 employs the far less common (for this experiment) technique of pushing off the ground rather than swinging limbs up. This bias may result from the particular method of fitness shaping used for this skill, in which an initial upward extension of the creature's axis-aligned bounding box is rewarded as an intermediate goal on the way to a true jumping behavior.

4.2.5 Jump Result 5

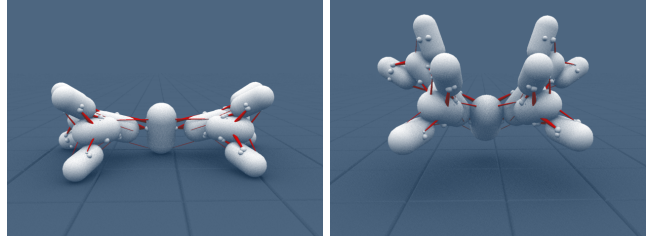


Figure 10: Complex-arm swing.

In Figure 10—the most morphologically elaborate of the jump results—a particularly complex collection of segments, joints, and muscles is applied to the work of swinging heavy arms up to induce a successful leap. (See Figure 1 for a more detailed illustration.)

4.3 Locomotion Results

In this section, results from a locomotion experiment are presented, in which the single fixed square-pulse control signal of Figure 4(a) was replaced with the repeating fixed square-wave signal of Figure 4(b), and the ultimate fitness function was changed from jump height to distance traveled in a given amount of time. A sequence of fitness goals like the one below—with the first three the same as in the jump results—is used for shaping. Note that in the fourth step, only a modest amount of horizontal travel is required for success. Once that initial degree of locomotion has been established, and the final stage begins, the previous requirement for jumping fitness is removed, and evolution is allowed to focus solely on optimizing horizontal travel towards an effectively unlimited distance goal.

1. **static AABB top**
2. **static AABB top + highest AABB top**
3. **highest AABB top + highest AABB bottom**
4. **highest AABB bottom + (modest) horizontal distance traveled**
(to encourage locomotion while ground clearance (jumping) is maintained)
5. **horizontal distance traveled (unlimited)**
(to optimize locomotion alone once it has begun to develop.)

In each of the following eight examples, the left image is a closeup of the creature with muscles relaxed (as during the trough of the activation square wave), and the right image depicts the creature with muscles activated, during locomotion, with approximate direction of movement indicated by the arrow.

4.3.1 Locomotion Result 1




Figure 11: Double front-armed swing hop.

In Figure 11, the square-wave activation of muscles is used to swing the front limbs up, accumulating momentum which produces forward translation during repeated jumps.

4.3.2 Locomotion Result 2




Figure 12: Single front-armed swing hop.

With morphology and action very similar to that of Jump Result 3, the creature in Figure 12 also employs a repeating forward-translating jump for simple but highly effective locomotion.

4.3.3 Locomotion Result 3




Figure 13: Front-armed swing step.

The mode of locomotion of the creature in Figure 13 is surprisingly complex and subtle, given the abrupt simplicity of the global activation signal. In a two-stage sequence of actions, this creature swings front legs up, which causes the middle box segments first to tip forward, then step ahead, pulling the back limbs along with them.

4.3.4 Locomotion Result 4




Figure 14: Delta wheelbarrow.

The creature in Figure 14 employs a dense concentration of muscles at its central joints to produce upward and forward momentum, which results in a wheelbarrowing forward slide.

4.3.5 Locomotion Result 5




Figure 15: Front-hinged swing drag.

In Figure 15, muscles sharply raise forward segments that are hinged so as to provide a lifting and forward-moving impulse, which drags the stabilizing rear legs along the ground.

4.3.6 Locomotion Result 6




Figure 16: Square wheelbarrow.

In Figure 16, a different morphology employs the same basic technique as Locomotion Result 4 to again produce a sliding wheelbarrow-like forward movement.

4.3.7 Locomotion Result 7




Figure 17: Complex swing step.

In Figure 17, the most morphologically complex of the locomotion results, one cluster of segments forms a stable base, while another such cluster is swung up to produce an elegant raise-tip-and-step sequence of actions, resulting in forward motion.

4.3.8 Locomotion Result 8




Figure 18: High hop.

In one of the simplest yet most effective locomotion results (Figure 18), the creature uses clusters of strong muscles to swing up heavy limbs, lifting its comparatively small root segment in a high-jumping locomotive technique.

5. DISCUSSION AND FUTURE WORK

It is important to note that the claimed benefits of the evolved muscle drives—removing a measure of the burden from control intelligence and embodying that intelligence as functional morphological complexity—are not expected to be limited to this particular form of adaptable drive. Any sufficiently inhomogeneous evolvable drive system should be able to accomplish the same goal. For example, if traditional EVC joint-motor drives had evolvable strengths, a similar transfer of intelligence from brain to body should be possible. The increase in morphological complexity in that case might be smaller (perhaps variable motor sizes displayed at a joint, rather than the varied number, orientation and attachment points exhibited by muscle drives), but still useful.

Another important point is that the work presented here is intended to establish that this evolvable musculature can embody *some* useful degree of control complexity, but does not yet include a quantification of that amount. This topic is worthy of a more systematic examination in the future.

Although the main goal of this work is to trade control intelligence for physical intelligence, the evolved muscle drives also make it possible to meaningfully increase morphological complexity. This goal can be advanced in several ways in the future.

One obvious next step would be to replace the current system’s simulated linear springs with a more complex model

employing simulated soft bodies or pressurized cloth for the bulk of the muscles. Previous work with simulated muscles [5] already demonstrated that such an approach is feasible, and has done so specifically within PhysX. Allowing muscles to help define the distribution of the body’s mass, as they do in real creatures, would significantly advance the process of biologically inspired purposeful complexification of morphology. Also, the meaningful change of such muscles’ shape during simulation—indicating the degree of their extension and activation—would add an additional layer of realistic detail.

Another biologically inspired refinement of the rigid-segment EVC model would be to simulate skin, as anticipated by Sims 20 years ago [12]. With powerful cloth simulation widely available (including in PhysX), this extension has become a conceivable next step on the path of life-like morphological complexification. In particular, combining simulated-cloth skin with massed muscles (as described above) might produce a particularly rich simulation, with a skin stretching and sliding over muscles as they extend and contract.

Looking even further into the future, if evolved creatures can embody the right kinds of morphological complexity, perhaps externally imposed joint mechanisms could be replaced by more realistic and more expressive joints whose properties arise directly from their morphology. By allowing the shape of the rigid-body segments to evolve [1], and permitting the inclusion of other necessary anatomical elements such as tendons and ligaments, it may be possible for rich and useful joint properties to emerge naturally, adding yet another layer of purpose-driven morphological complexity to evolved virtual creatures.

Similarly, although on a somewhat different path, the ability to evolve sufficiently detailed exoskeleton segments, along with the necessary muscles and connecting elements, could permit the development of exoskeleton-based virtual creatures. In this style of morphology—where again, body function follows from its form—meaningful complexity should emerge.

Moving forward into more demanding tasks in the future, the diversity of results displayed in Sections 4.2 and 4.3 is encouraging. The variety of available solutions for these simple behaviors promises greater opportunity for continued success as new constraints are applied.

6. CONCLUSION

This paper has described a version of evolved virtual creatures in which traditional joint-motor drives are replaced by a simple yet powerful evolvable musculature. The results presented here have demonstrated that this new substrate can support a significant degree of physical intelligence, sufficient to almost entirely replace the control intelligence that would normally be used for basic but useful tasks such as jumping and locomotion. The process of shifting this intelligence into a new complexity in the body makes it visible, enabling progress toward the goal of meaningful morphological complexity. And the fact that this can make (for these basic tasks) the typical EVC brain essentially superfluous, gives some indication of how much of a control burden these muscle drives can embody, and the degree to which they can liberate the brain’s computational resources for other, more complex work. In addition, it demonstrates that, in some real-world applications, where these simple tasks are useful,

but brains are difficult to support, a sufficiently evolvable drive system like this one may enable an entirely new class of solutions to emerge.

7. ACKNOWLEDGEMENTS

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

8. REFERENCES

- [1] J. E. Auerbach and J. C. Bongard. Evolving CPPNs to grow three-dimensional physical structures. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation*, GECCO '10, pages 627–634, New York, NY, USA, 2010. ACM.
- [2] J. E. Auerbach and J. C. Bongard. 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, 2012. ACM.
- [3] N. Chaumont, R. Egli, and C. Adami. Evolving virtual creatures and catapults. *Artificial Life*, 13(2):139–157, 2007.
- [4] N. Cheney, R. MacCurdy, J. Clune, and H. Lipson. 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, 2013. ACM.
- [5] K. Glette and M. Hovin. 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, 2010.
- [6] M. Joachimczak and B. Wróbel. Co-evolution of morphology and control of soft-bodied multicellular animats. In *Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation Conference*, GECCO '12, pages 561–568, New York, NY, USA, 2012. ACM.
- [7] N. Lassabe, H. Luga, and Y. Duthen. A new step for artificial creatures. In *Artificial Life, 2007. ALIFE'07. IEEE Symposium on*, pages 243–250. IEEE, 2007.
- [8] J. Lehman and K. Stanley. 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, 2011.
- [9] D. Lessin, D. Fussell, and R. Miikkulainen. 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, 2013. ACM.
- [10] T. Miconi. In silicon no one can hear you scream: Evolving fighting creatures. *Genetic Programming*, pages 25–36, 2008.
- [11] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, USA, 1998.
- [12] K. Sims. 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, 1994. ACM.
- [13] B. Skinner. The behavior of organisms: An experimental analysis. *The Century psychology series*, 1938.