Open-Ended Behavioral Complexity for 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

In the 19 years since Karl Sims' landmark publication on evolving virtual creatures [22], much of the future work he proposed has been implemented, having a significant impact on multiple fields including graphics, evolutionary computation, and artificial life. There has, however been one notable exception to this progress. Despite the potential benefits, there has been no clear increase in the *behavioral complexity* of evolved virtual creatures (EVCs) beyond the light following demonstrated in Sims' original work.

This paper presents an open-ended method to move beyond this limit, making use of high-level human input in the form of a syllabus of intermediate learning tasks—along with mechanisms for preservation, reuse, and combination of previously learned tasks. This method (named *ESP* for its three components: *encapsulation*, *syllabus*, and *pandemonium*) is employed to evolve a virtual creature with behavioral complexity that clearly exceeds previously achieved levels. ESP thus demonstrates that EVCs may indeed have the potential to one day rival the behavioral complexity—and therefore the entertainment value—of their non-virtual counterparts.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning—*connectionism and neural nets*; I.6.8 [**Simulation and Modeling**]: Types of Simulation—*animation*

General Terms

Algorithms

Keywords

evolved virtual creatures, artificial life, physics-based character animation, task decomposition, content creation

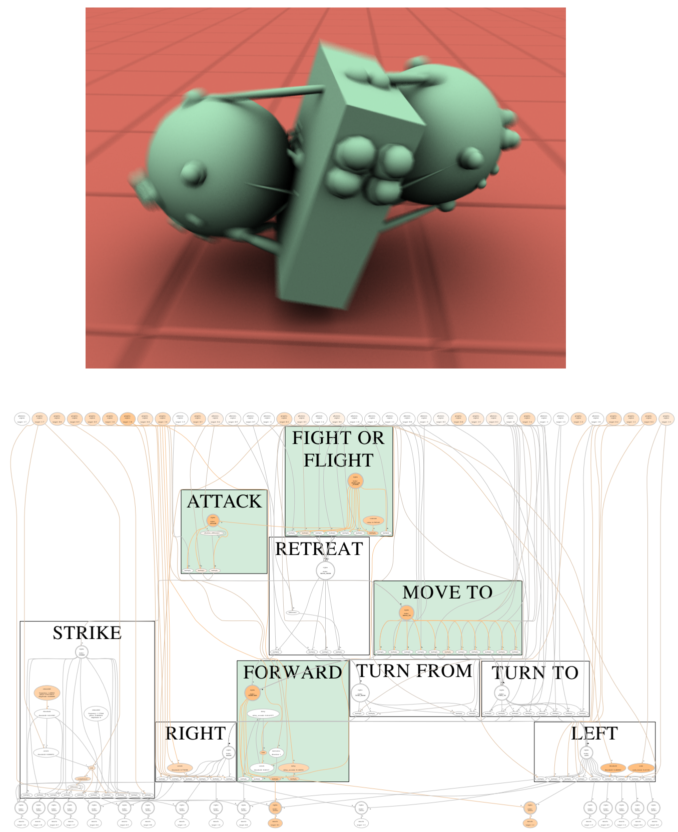

Figure 1: The body and brain of a creature evolved using the ESP method. At this point, the creature has achieved a level of behavioral complexity (as defined in the introduction) which is approximately double the state of the art for evolved virtual creatures.

1. INTRODUCTION

Defining *behavioral complexity* as the number of discriminable behaviors in a creature's repertoire, many of Karl Sims' creatures [22] could be said to have minimal complexity, employing repertoires which contain only a single behavior. His examples of locomotion on land and in water, as well as jumping, fall into this category, as does much of the work that others have since completed. For example, locomotion in air for EVCs has been demonstrated [20], and a specialized form of ground-locomoting EVCs have been produced which can be converted into functional real-world robots [13]. Soft-bodied virtual creatures have been evolved [9, 7], and many other variations at this level of complexity have been presented [3, 8, 1, 11, 10, 12]. The highest level of behavioral complexity demonstrated by Sims—creatures with the ability to follow a target or a path by switching between perhaps up to four discriminable behaviors—has since been matched multiple times [17, 21, 14], but never clearly exceeded. (Miconi's work [14] is a particularly interesting case, as he is the first to produce a form of real combat between EVCs, but with respect to behavioral complexity as defined above, his creatures do not differ significantly from those of Sims, as their combat can be essentially viewed as target following with damage assignment layered on top—the target following leads to collisions, and these collisions produce a score interpreted as damage, but no additional behavioral complexity is required or produced.)

This lack of progress is despite the apparent usefulness of more complex behaviors. Numerous examples of valued creature content from the real world—nature documentaries, animal and human combat, even internet cat videos[1]—have in common a level of behavioral complexity that is clearly greater than what has been demonstrated in EVCs to date. Perhaps if we can bring greater behavioral complexity to EVCs, they can begin to approach the entertainment value of their non-virtual counterparts.

In fact, there is suggestive evidence in support of this proposition. Cognitive science and psychology describe a striking effect in which the right kinds of relatively complex behaviors—even by the simplest of geometric figures—lead to the perception of intentionality and desires (*perceptual animacy*) [18], as described in the classic work by Heider and Simmel [6]. For a particularly clear non-academic example of this same effect, consider the academy-award-winning animated short "The Dot and the Line" (Chuck Jones, 1965). In much of this film, the only elements added to a simple dot and line to transform them into the protagonists of a compelling love story are their movements–their behavioral complexity.

Motivated by this potential, this paper describes a method designed to significantly increase behavioral complexity in evolved virtual creatures: *ESP*. The primary elements of this method—*encapsulation*, *syllabus*, and *pandemonium*—are defined as follows:

1. A human-designed *syllabus* breaks the development of a complex behavior into a sequence of smaller learning tasks.

2. Once each of these subskills is learned, it is *encapsulated* to preserve it throughout future evolution, and also to allow future skills to more easily incorporate its function.

3. Finally, a mechanism inspired by Selfridge's *pandemonium* [19] is used to resolve disputes between competing skills or drives within the increasingly complex brain.

At this point, it is important to note that, while there are particular challenges in applying it to EVCs, this kind of task decomposition has been known in multiple related fields for many years. Selfridge's pandemonium, Minsky's society of mind [15], and Brooks' subsumption architecture [2] are prominent examples from artificial intelligence and robotics. And in reinforcement learning and evolutionary computation, work such as layered learning and hierarchical task decomposition [24, 25, 4] explores similar concepts. In EVCs, however, no previous system has demonstrated the use of such an approach to increase behavioral complexity beyond existing limits.

In the remainder of this paper, the details of the ESP system are presented, and this method is employed to approximately double the state of the art in behavioral complexity for evolved virtual creatures (Figure 1).

[1]e.g., "THE BEST CAT VIDEO YOU'LL EVER SEE" [sic], http://www.youtube.com/watch?v=20mrEtabOLM

2. BASIC EVC SYSTEM

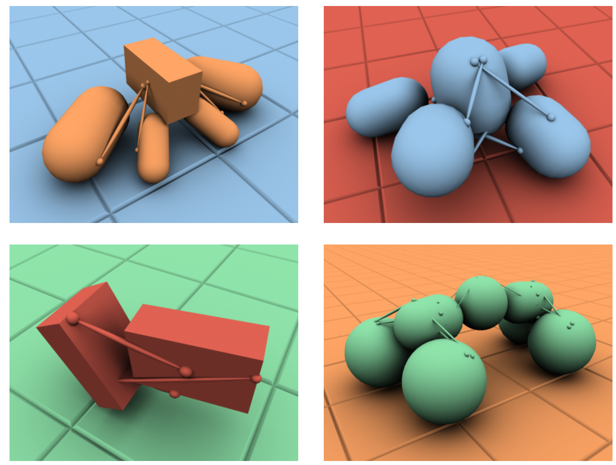

Figure 2: Typical results from the basic EVC system. These examples were all evolved to complete a forward locomotion task—a common baseline result for EVCs.

The basic EVC system described here is largely derived from the work of Karl Sims [22]. This section briefly sets out the components of this system, which—while not the primary focus of this paper—are nevertheless fundamental to its comprehension. A representative sample of results is shown in Figure 2.

2.1 Evolutionary Algorithm

The specifics of the evolutionary algorithm are conventional, making use of elitism, fitness-proportionate selection, and rank selection [16]. In addition, the most challenging tasks employ some degree of shaping [23]. Fitness is evaluated in a physically simulated virtual environment implemented with NVIDIA PhysX.

2.2 Morphology

As in Sims' original work [22], 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 easily define recursive and repeated body substructures, as illustrated in Figure 3. In addition, as in Sims' work, reflection of body parts as well as body symmetry are made easily accessible to evolution. In this implementation, all PhysX primitives are made available for use as body segments: boxes, spheres, and capsules.

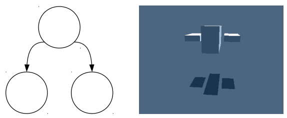

(a) Simple topology.

(b) Multiple edges for repeated substructures.

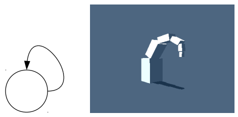

(c) Reflexive edge for recursive structure.

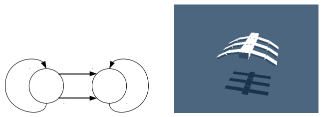

(d) Multiple and reflexive edges together.

Figure 3: Hand-designed genotype/phenotype pairs (as in [22]) demonstrate the encoding power inherited from Sims' original EVC system.

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 in this system are provided implicitly by creature structure through natural collisions between adjacent segments.

2.3 Control

Again in a manner 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 6a). 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 special node types for muscles and photoreceptors (described below) and one special type used in encapsulation (see Section 3.2), the following node types are allowed: sinusoidal, complement, constant, scale, multiply, divide, sum, difference, derivative, threshold, switch, delay, and absolute difference.

2.4 Photoreceptors

For tasks involving light sensing, creatures are allowed to develop simple photoreceptors ((a) in Figure 4), 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.

2.5 Muscles

In a break with traditional EVC systems, which typically use forces exerted directly at joints, this system uses simulated muscles as actuators. Each muscle ((b) in Figure 4) 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 on its current length. For each muscle, one node is added to the brain which accepts an input to set the muscle's activation, and another node is added which makes the muscle's proprioceptive output signal available to the rest of the brain. Muscle drives bring the following potential benefits to EVCs: flexibility (they can be used even on creatures without joints), efficiency (effectors need only exist where useful, not at every degree of freedom of every joint), and beauty (by tapping into the human affinity for elegant, functional body structure).

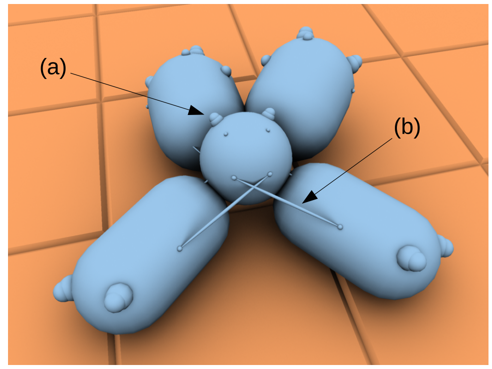

Figure 4: Photoreceptors (a) and muscles (b) bring sensing and actuation to creatures in the basic EVC system. For both, function depends upon placement, so creature form develops meaningfully as capabilities are evolved.

3. ESP

The ESP method consists of three elements added to the basic EVC system: a syllabus, encapsulation, and pandemonium. In this section, each of these components is described in detail.

3.1 Syllabus

While it is certainly possible for human students to learn a complicated topic independently, their development is typically faster and surer 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. 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 all of these skills might be constructed, but evolutionary progress would be unlikely.

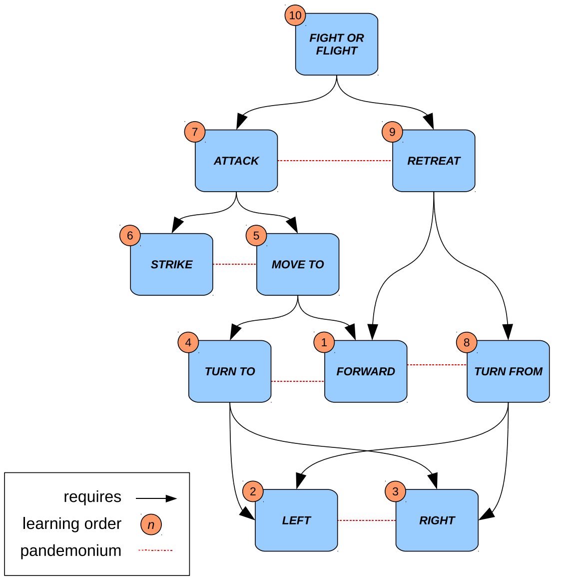

Figure 5: An example syllabus as a graph. In this depiction, graph nodes represent individual subskills to be learned, directed edges indicate dependency between subskills, and the numbering indicates a proposed learning order which satisfies the dependency requirements. Pandemonium relationships are indicated by dashed red lines.

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 would seem 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 the following sequence of subskills to be learned, in which each subskill is probably attainable with basic EVC methods, and in which earlier subskills serve to make later skills easier to learn:

1. FORWARD LOCOMOTION
2. LEFT TURN
3. RIGHT TURN
4. TURN TO LIGHT (using LEFT TURN and RIGHT TURN)
5. MOVE TO LIGHT (using TURN TO LIGHT and FORWARD LOCOMOTION)
6. STRIKE
7. ATTACK LIGHT (using MOVE TO LIGHT and STRIKE)
8. TURN FROM LIGHT (using LEFT TURN and RIGHT TURN)
9. RETREAT FROM LIGHT (using TURN FROM LIGHT and FORWARD LOCOMOTION)
10. FIGHT OR FLIGHT (switching between ATTACK LIGHT and RETREAT FROM LIGHT based on external circumstances)

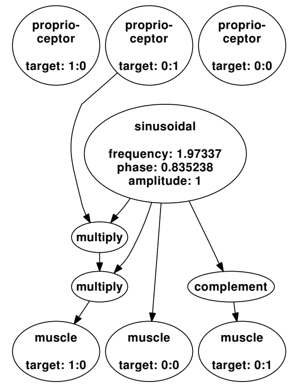

(a) Before encapsulation.

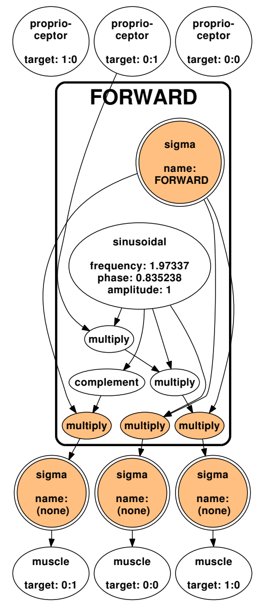

(b) After encapsulation (with new nodes shaded).

Figure 6: 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.

This information may be conveniently summarized in a graph, encompassing subskills to be learned, dependency between subskills, learning order, and pandemonium (Section 3.3), as seen in Figure 5.

At this point, using high-level human knowledge, a previously impractical learning task has been broken into a sequence of potentially attainable subgoals. But how can a single evolving creature learn new skills while retaining and making use of the ones it already has?

3.2 Encapsulation

The second important element of the ESP system is a mechanism to *encapsulate* previously learned skills. This accomplishes two important goals: It ensures that previously learned skills (and the body components they rely on) are preserved, and it makes these skills easily accessible to future evolutionary development. Both of these goals are achieved through the automated encapsulation process illustrated in Figure 6.

Figure 6a depicts a brain evolved for forward locomotion, and Figure 6b shows the result of encapsulation. Note the following aspects of this new brain. The nodes that compute the old skill have been preserved and locked (meaning that they have been marked so as to disallow any changes by future evolution). Also, 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]. And finally, 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.) As progress through the syllabus continues and the next skill after FORWARD LOCOMOTION is evolved, its newly added nodes will be the only ones in the brain that are not already locked, and will therefore be easily identifiable when it is their turn to be encapsulated.

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

3.3 Pandemonium

Consider the following example based on the syllabus graph of Figure 5. 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, 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. In Figure 5, a full set of pandemonium relationships is indicated by red dashed lines between subskill nodes.

With this final component of the ESP system described, it is now possible to consider a full example, in which previously achieved levels of behavioral complexity are first matched, then exceeded.

4. RESULTS

The primary result of this paper is an application of the ESP method, using the syllabus of Figure 5, to evolve a virtual creature through a sequence of ten learning tasks, the first five of which approximately match the previously demonstrated behavioral-complexity limit for EVCs, and the second five of which approximately double it. (These results are best viewed in the accompanying video[2].)

4.1 FORWARD LOCOMOTION

A FORWARD LOCOMOTION result from the basic EVC system has been chosen, and its control abilities have been encapsulated, as shown in Figure 7. This creature was evolved through traditional EVC techniques, including the use of shaping, with the ultimate fitness being defined by the interleaving of an efficiency score into a discretized score for speed. Specifically, if s is the creature's speed, s_{max} is the maximum speed, σ is the discretization step, and ϵ is a measure of the creature's efficiency (in [0, 1]), the combined fitness is defined as

$$\frac{\sigma(\lfloor \frac{s}{\sigma} \rfloor + \epsilon)}{s_{max}}.$$

This is intended to ensure that speed is the primary factor in fitness, but increased efficiency (while maintaining approximate speed) is also rewarded.

At this point, the creature has developed the rigid body segments, muscles, and control system it needs for successful locomotion, and these elements will be preserved as evolution continues.

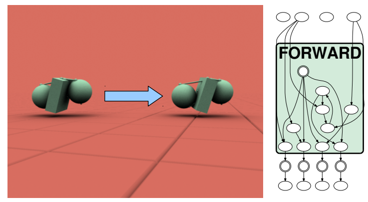

Figure 7: FORWARD LOCOMOTION encapsulated.

4.2 LEFT TURN

With the LOCOMOTION skill preserved, a new run of evolution begins, this time with the fitness function rewarding the ability to rotate counterclockwise while largely maintaining position. The addition of new muscles is allowed during this process. The resulting completed skill is shown (after encapsulation) in Figure 8.

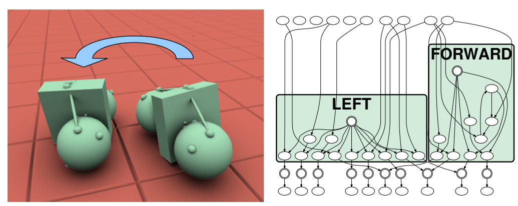

Figure 8: LEFT TURN added.

4.3 RIGHT TURN

With the first two skills preserved, a clockwise turn is evolved in the same way as the counterclockwise turn, and the result is encapsulated (Figure 9). At this point, the creature has all of the low-level skills that it will need to reach any point on the ground, with the majority of future skills relying ultimately on reapplications of FORWARD LOCOMOTION, LEFT TURN, and RIGHT TURN.

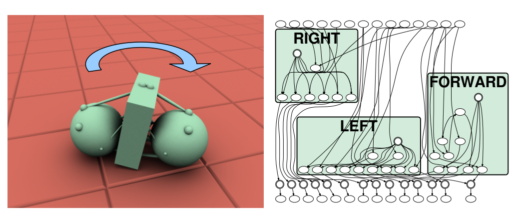

Figure 9: RIGHT TURN added.

[2] http://youtu.be/dRLNnJlT8rY

4.4 TURN TO LIGHT

At this point, the creature is allowed to develop photoreceptors, while being tested on its ability to orient itself to a target (which is perceived as a point light source) using the previously encapsulated LEFT TURN and RIGHT TURN skills. The fitness evaluation is an average over four runs, each with a fixed light source at a different heading from the creature. Figure 10 shows the completed and encapsulated result, which is able to consistently aim its locomotion direction at a user-controlled target.

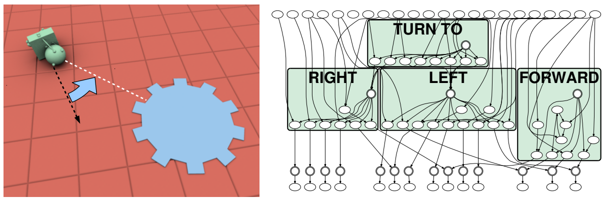

Figure 10: TURN TO LIGHT has been added, which keeps the locomotion direction (black dashed arrow) oriented toward a target.

4.5 MOVE TO LIGHT

Now, with TURN TO LIGHT and FORWARD LOCOMOTION available, and with the evolution of new photoreceptors allowed, the creature is evaluated on its ability to navigate to a light source. As with TURN TO LIGHT, fitness is averaged over multiple runs (in this case five), again with a fixed light source at a different relative angle each time. The result (Figure 11) is a creature whose behavioral complexity approximately matches the current state of the art.

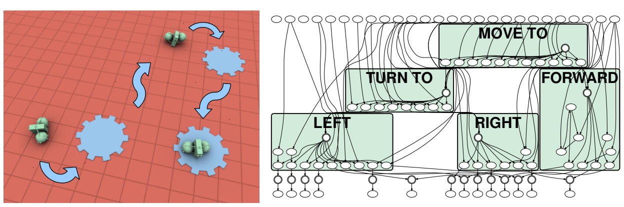

Figure 11: MOVE TO LIGHT has been added, allowing the creature to follow a target along a curving path, catching the target when it finally stops.

4.6 STRIKE

In anticipation of the upcoming ATTACK task (see Figure 5), the creature must first learn to deliver a strike to the ground underneath it. For this creature, that goal is accomplished with a vertical jump, as seen in Figure 12. To facilitate the evolution of this new low-level skill, the development of new muscles is allowed.

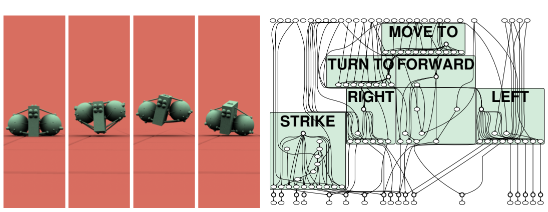

Figure 12: This creature's STRIKE solution employs a vertical jump.

4.7 ATTACK

Having learned MOVE TO LIGHT and STRIKE, it is now possible to produce an ability slightly more complex than simply moving to a target. By first moving to the target, *then* striking, this creature (Figure 13) takes another small step toward the behavioral complexity of compelling creature content from the real world. For this task, fitness is an average across four directions of distance from the target when the first sufficiently strong ground impact occurs (with a penalty for producing such an impact when the scene contains no light).

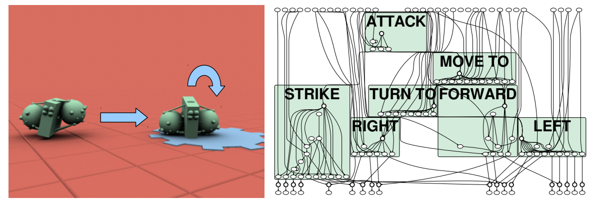

Figure 13: In the newly added ATTACK, the creature navigates to the target, then strikes it.

4.8 TURN FROM LIGHT

Now, in preparation for the upcoming RETREAT skill (see Figure 5), the creature must learn to turn away from a light source (as shown in Figure 14). Although obviously similar to TURN TO LIGHT, this task also required a fitness term to discourage an initial wrong-direction turn, so as to achieve reasonable results for targets near the creature's front. Also, significantly more evaluation directions (thirteen) were used (particularly near the front) to similarly motivate appropriate reactions in these cases.

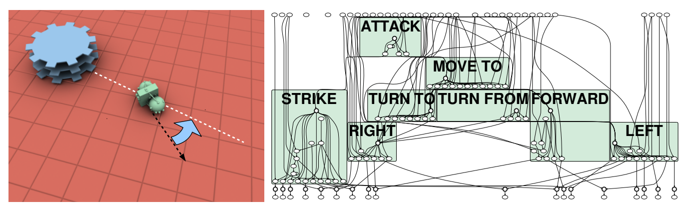

Figure 14: TURN FROM LIGHT has been added, which keeps the locomotion direction (black dashed arrow) oriented away from the target.

4.9 RETREAT

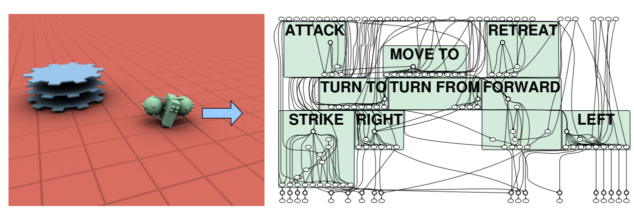

Figure 15: RETREAT added.

At this point, using TURN FROM LIGHT and FORWARD LOCOMOTION, the creature learns to maximize its average distance from a light target. As with TURN FROM LIGHT, penalties for initial wrong-direction moves and multiple tests with targets near the front are combined to discourage inappropriate initial reactions. With this skill complete (Figure 15), the necessary components are in place for the final top-level skill of the syllabus.

4.10 FIGHT OR FLIGHT

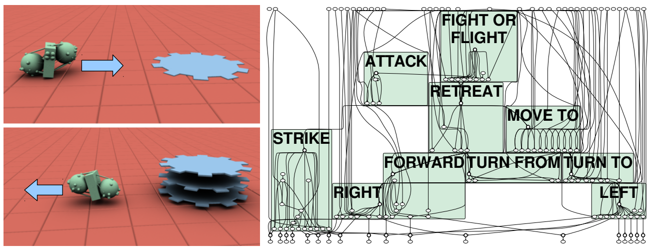

Figure 16: FIGHT OR FLIGHT has been added, completing the progression through the syllabus.

The task of this final, highest skill is to choose between ATTACK and RETREAT based on the perceived environment. For this evaluation, the creature is confronted with a vulnerable target (a single disc on the ground), which the creature should attack, and a dangerous target (a spinning vertical stack of three such discs), which will destroy the creature if touched.

The fitness score is again the result of averaging over initial light directions, but in this case there is some additional complexity. At each direction, one evaluation is made with a vulnerable target, and one with a dangerous target. While the proper reaction in a single case (vulnerable vs dangerous) should be rewarded, the real challenge is to motivate a discrimination between the two, so that the right action can be taken in *both* cases. To accomplish this, a small fraction of the final score is based on the average maximum of the two component scores (to motivate *any* development, especially initially), and a much larger fraction of the final score is based on the average minimum of the two component scores (to reward the ultimate goal of finding the proper reaction in *both* cases). The weighting is chosen so that a single perfect result for a minimum component will be worth more than perfect scores in *all* of the maximum components. So if f^+ is the average *maximum* score across all n test directions, and f^- is the average *minimum* score across all n test directions, then the final overall fitness is computed as

$$\frac{f^+ + 2n \cdot f^-}{2n+1}.$$

Without these additional motivations, solutions emerged which chose a single (higher-scoring) hard-coded reaction to be used for each light position—regardless of target type—without making the leap to the increased scores available if discernment between the two types of target could be developed.

Figure 16 shows a successful result for this task, marking the completion of the full syllabus and the acquisition of its highest, most complex skill. This result demonstrates that the ESP system can enable evolved virtual creatures to achieve a level of behavioral complexity which is a clear advance on the state of the art.

5. DISCUSSION

This section examines the roles of human input, evolution, and physical simulation in this system, as well as the future potential of ESP for EVCs.

The human input utilized by this method in the form of the syllabus is at a usefully abstract level—on a par with the kind of input employed by human learners. This syllabus, along with the opportunity for human selection among high scorers at the end of each subskill stage, offers great potential value as a mechanism for exerting relatively high-level creative control over creature development.

In addition, numerous benefits accrue from the fact that this system's results are evolved and that this evolution takes place in physical simulation. Thanks to evolution, the creatures this system produces are unceasingly novel, developing new solutions for morphology, muscle and eye placement, and mechanism and style of movement each time the process is restarted. And the fact that these solutions are evolved to operate in a physically simulated environment adds a particular level of realism, demonstrating results that are convincingly physically plausible, and even include some of the subtle imperfections of action that bring so much character to creatures in the real world. Note, also, that creating controllers for bodies like these by hand would be impractical, but that this difficulty is in this case handled entirely by the evolutionary algorithm.

Finally, one of the most important aspects of the ESP system is that it is designed to be open ended. While a significant increase in behavioral complexity has been demonstrated, there are no obvious barriers to continued reapplication of this technique to achieve results of still greater complexity in the future.

6. FUTURE WORK

Relaxing Encapsulation In the work described above, once encapsulation is applied, it is absolute. But relaxing this restriction could bring a number of benefits. Once the top-level skill is acquired, it might be suitable to remove all encapsulation restraints and continue evolution, so as to achieve smoother, better-integrated solutions. Note also that encapsulation as applied above fixes segments and joints (although not muscles and photoreceptors) after the first skill, which may place unacceptable limitations on some applications (e.g., *combat*, as described below). In cases like this, continued testing of any previously learned abilities that need to be maintained might take the place of such strict constraints.

Muscles and Skin The fact that this EVC system is built around functional muscles presents a great opportunity. In the current system, muscles are implemented as simple linear springs, but a more advanced model might employ physically simulated soft bodies or cloth (as in [5]) for the bulk of the muscles themselves, along with a variation of cloth simulation for skin (as proposed in [22]). This could bring a new level of realism and beauty to such an EVC system, as well as the potential for more complex and realistic functional damage.

Combat While Miconi has already produced one limited form of combat for EVCs [14], there is a great deal more that can be done in this area. The ESP method, in combination with the future-work topics described above (and the ability to vary body-part materials, the importance of which was recognized by Miconi), could potentially produce a far richer and more compelling form of combat for evolved virtual creatures than what has been seen to date.

Fauna on Demand Finally, a more refined and automated version of the ESP system could make it possible to populate virtual worlds with continually novel creature content (especially with the help of techniques such as those seen in [12]). As virtual boundaries are pushed back, human users could (subject to limitations of computing power) continually encounter never-before-seen creatures, all developed from a single high-level human-designed syllabus.

7. CONCLUSION

The ESP system described in this paper has allowed evolved virtual creatures to achieve a new level of behavioral complexity (as defined in the introduction) which is approximately double the state of the art. This advance demonstrates that the behavioral complexity of evolved virtual creatures has not yet been exhausted, and in fact suggests that it may continue to increase so as to one day match the behavioral complexity of creatures from the real world—with all of the potential for content creation that this might bring.

8. ACKNOWLEDGEMENTS

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

9. REFERENCES

[1] J. C. Bongard and R. Pfeifer. Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. pages 829–836. Morgan Kaufmann, 2001.

[2] R. Brooks. A robust layered control system for a mobile robot. *Robotics and Automation, IEEE Journal of*, 2(1):14–23, 1986.

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

[4] J. Doucette, P. Lichodzijewski, and M. Heywood. Hierarchical task decomposition through symbiosis in reinforcement learning. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference*, pages 97–104. ACM, 2012.

[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] F. Heider and M. Simmel. An experimental study of apparent behavior. *The American Journal of Psychology*, pages 243–259, 1944.

[7] J. Hiller and H. Lipson. Automatic design and manufacture of soft robots. *Robotics, IEEE Transactions on*, 28(2):457–466, 2012.

[8] G. Hornby, J. Pollack, et al. Body-brain co-evolution using l-systems as a generative encoding. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 868–875, 2001.

[9] 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*, pages 561–568. ACM, 2012.

[10] M. Komosinski. The Framsticks system: versatile simulator of 3D agents and their evolution. *Kybernetes: The International Journal of Systems & Cybernetics*, 32(1/2):156–173, 2003.

[11] 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.

[12] 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.

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

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

[15] M. Minsky. *Society of mind*. Simon & Schuster, 1988.

[16] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, USA, 1998.

[17] M. L. Pilat and C. Jacob. 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, 2010. ACM.

[18] B. Scholl and P. Tremoulet. Perceptual causality and animacy. *Trends in cognitive sciences*, 4(8):299–309, 2000.

[19] O. G. Selfridge. 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, Nov. 1958. HMSO.

[20] Y. Shim and C. Kim. 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, 2003.

[21] Y. Shim, S. Kim, and C. Kim. Evolving flying creatures with path-following behavior. In *ALife IX: Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems*, pages 125–132, 2004.

[22] 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.

[23] B. Skinner. The behavior of organisms: An experimental analysis. *The Century psychology series*, 1938.

[24] P. Stone and M. Veloso. Layered learning. In *Machine Learning: ECML 2000: 11th European Conference on Machine Learning Barcelona, Catalonia, Spain May, 31-June 2, 2000 Proceedings*, volume 1810, page 369. Springer, 2000.

[25] S. Whiteson, N. Kohl, R. Miikkulainen, and P. Stone. Evolving keepaway soccer players through task decomposition. *Machine Learning*, 59(1):5–30, 2005.