Evolved Virtual Creatures as Content: Increasing Behavioral and Morphological Complexity

Daniel Gregory Lessin

Report TR-15-01
January 2015

dlessin@cs.utexas.edu
http://nn.cs.utexas.edu/?danlessin

The University of Texas at Austin
Austin, TX 78712

The Dissertation Committee for Daniel Gregory Lessin
certifies that this is the approved version of the following dissertation:

Evolved Virtual Creatures as Content: Increasing Behavioral and Morphological Complexity

Committee:

Don Fussell, Co-Supervisor

Risto Miikkulainen, Co-Supervisor

Dana Ballard

Vladimir Lifschitz

Joe Marks

Evolved Virtual Creatures as Content: Increasing Behavioral and Morphological Complexity

by

Daniel Gregory Lessin, A.B.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

THE UNIVERSITY OF TEXAS AT AUSTIN

May 2015

Dedicated to Nicole. All I had to do was write this dissertation; she had the significantly tougher job of putting up with me while I did it.

Acknowledgments

To begin, I want to thank my supervisors, Don Fussell and Risto Miikkulainen. If I am ever tempted to be too proud of what I think I've accomplished on my own, it will only be an indication of how well they have done their job.

I would like to thank my first academic advisor, Joe Marks. Although he has surely been a mentor for over 20 years now, he has more importantly been an inspiration, not only in his work, but in his respect and caring for those around him. I may not be able to repay him for his kindness, but I know I can try to do the same for others in the future.

I would also like to thank Vladimir Lifschitz. Academically, he is the most powerful example I know of the scientific ideal of humility before the truth. As a teacher and advisor, he is equally committed to his students' success and happiness, and is universally liked and respected for it.

To my daughters Katie and Gwen: Hey, you're in my dissertation! I'm sorry I've been so busy for so long. :(Maybe now I'll have a little more time to do all the things you guys want to do. :) By the way, I left some candy for you in the right-most drawer in the kitchen. Share nice!

Also, I want to thank my mother, Dona. Beyond all of her efforts and sacrifice to give me a good life, she dedicated her professional career to helping those most in need. I hope she always remembers the good she has done for so many, and feels them surrounding her with the love and gratitude she so richly deserves.

Finally, I want to thank my father, Len. I could say I wished he were here to see this, but thanks to him, I know that the kindness, love, and generosity that we show to the world every day are far more important than personal accomplishments. To attempt to carry on this lesson from him will be a greater honor than any degree I could ever earn.

Evolved Virtual Creatures as Content: Increasing Behavioral and Morphological Complexity

Publication No. __________________

Daniel Gregory Lessin, Ph.D.
The University of Texas at Austin, 2015

Supervisors: Don Fussell
Risto Miikkulainen

Throughout history, creature-based content has been a highly valued source of entertainment. Recently, *evolved virtual creatures* (or *EVCs*; Sims 1994) were proposed as a potential new source of creature content. In EVCs, the creature's morphology and the control network driving its behavior are evolved together to accomplish naturalistic tasks. Despite their immediate appeal, however, EVCs still lag far behind their natural counterparts: Neither their morphology nor their behavior is sufficiently complex. This dissertation presents three contributions to address this problem. First, the *ESP* system, which combines a human-designed syllabus with encapsulation and conflict-resolution mechanisms, is used to approximately double the state of the art in behavioral complexity for EVCs. Second, an extension to ESP is presented that allows full morphological adaptation to continue beyond the initial skill. It produces both a greater variety of solutions and solutions with higher fitness. Third, a muscle-drive system is demonstrated to embody a significant degree of physical intelligence. It increases morphological complexity and reduces demands on the control network, thus freeing resources for more complex behaviors. Together, these contributions bring evolved virtual creatures, in both action and form, a step closer to matching the entertainment value of creatures from the real world.

Table of Contents

List of Figures

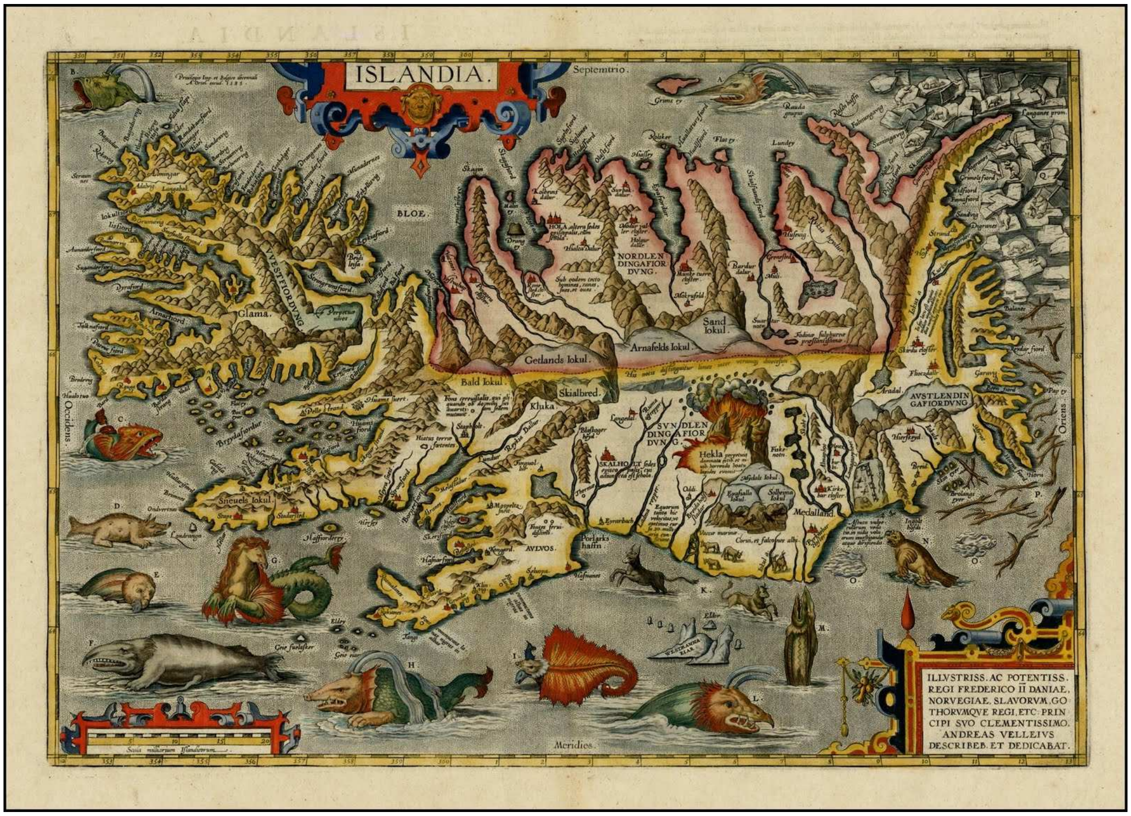

Figure 1.1: Islandia map by Abraham Ortelius, circa 1600. We once lived in a world with mysterious realms worthy of the phrase Here Be Dragons*. Today it is becoming possible to create such worlds virtually, and populate them with automatically generated creature content.*

Chapter 1

Introduction

Imagine living in a time when, just by traveling over the horizon, you could discover not only new geography, but new life as well. Ancient maps filled uncharted areas with illustrations of imagined monsters and inscriptions like *Here Be Dragons* to express this state of exhilarating ignorance about the world (Figure 1.1). And just when it seemed the age of discovery was over, virtual worlds have provided a new avenue for exploration, offering novel terrain that is created faster than we can experience it. However, while it is possible to construct and program creatures to inhabit these worlds, something vital is missing: creature content that is compelling and truly novel, like the kind seen in nature. Taking a step in this direction by making evolved virtual creature content complex and compelling is the focus of this dissertation.

1.1 Motivation

Creature content has been highly valued throughout human history (Figure 1.2). From the cruel spectacle of animal (and human) combat, to educational and inspirational nature documentaries, to the emergent slapstick of pet videos on television and the internet with millions of views[1], real-world creatures have long been a prominent source of compelling and valued entertainment. In recent decades, advances in computer power and sophistication have made it possible to consider not only creature content evolved and observed in nature, but also the entirely new parallel category of creature content evolved and simulated in virtual environments. This technology opens the door from evolved-creature-content-as-we-know-it to evolved-creature-content-as-it-could-be.

This new font of creature content offers numerous potential advantages over its non-virtual predecessor. By applying the creative power of evolution, this new field has the potential to produce endlessly novel virtual creature content; instead of hoping to someday discover alien life evolved on other planets or in some as-yet-inaccessible region of our own, we can restart randomly seeded virtual evolution whenever we wish. And further, virtual evolution may be modified to suit specific content requirements. Abstract goals such as morphological or behavioral novelty can be rewarded [26], more specific constraints such as topological complexity or simplicity can be enforced, and indirect pressures may be applied, such as environmental modifications. Note that many of these techniques, although they involve human input, still result in genuine creativity: Only the desired outcome is given, with the specifics of the solution emerging purely from evolutionary processes.

In addition, these creatures can be used (either as finished products or along with mechanisms for continued evolution) in non-real environments, from computer-generated movie scenes to the larger worlds of video games and open virtual environments. With sufficient progress in performance to allow near-real-time creature evolution, the exploration of a virtual world could become like real-world exploration. Pushing back boundaries in it results in the genuine discovery of new forms of life.

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

(a) Combat.

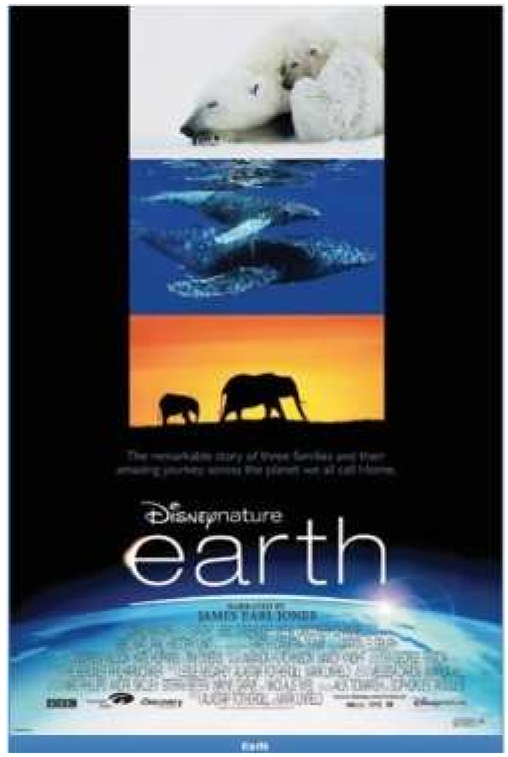

(b) Nature documentaries.

(c) Pet videos on television.

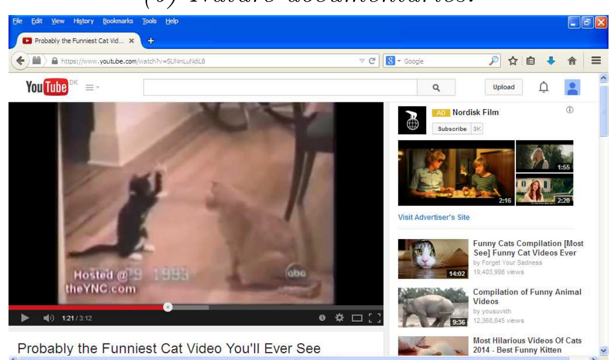

(d) Internet cat videos.

Figure 1.2: Compelling creature content. Throughout history, many forms of highly valued entertainment have been derived from animal and human life in the natural world. The goal of this dissertation is to create similarly entertaining content for virtual worlds.

1.2 Challenge and Opportunity

In 1994, this endeavor got off to a strong start with the publication of Karl Sims' landmark paper on evolving virtual creatures (*EVCs*) [46]. In it, physical simulation and artificial evolution were combined to develop morphology and control simultaneously for novel and compelling evolved virtual creatures. Their abilities included locomotion on land and in water, jumping, and, most impressively, phototaxis (the ability to move toward a user-controlled light source). Since that time, significant extensions to Sims' work have been demonstrated in multiple fields, including computer graphics, artificial life, evolutionary computation, and even robotics. But there has, to date, been one notable exception to this progress. Despite the potential benefits to creature content, there has been no clear increase in the behavioral complexity of EVCs beyond the light following demonstrated in Sims' original work.

Defining *behavioral complexity* as the number of discriminable behaviors in a creature's repertoire, many of Karl Sims' creatures could be said to have minimal complexity, employing repertoires that 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 was demonstrated [42], and a specialized form of ground-locomoting EVC was produced, which can be converted into functional real-world robots [29]. Soft-bodied virtual creatures were evolved for locomotion [21, 18], and many other variations at this level of complexity were presented [7, 20, 5, 25, 23, 26].

The highest level of behavioral complexity demonstrated by Sims—creatures with the ability to follow a target or a path by switching between up to four discriminable behaviors—has since been matched multiple times [37, 43, 30]. However it has never been clearly exceeded, even though more complex behaviors would be useful. There are numerous examples of creature content in the real world that are valued precisely because they are complex—much more than what has been demonstrated in EVCs to date. If we could give the behavior of EVCs similar complexity, they might begin to approach the entertainment value of their non-virtual counterparts.

In fact, there is suggestive evidence in support of this proposition. There is a striking effect in cognitive science and psychology 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*) [38]. This principle is well described in the classic work by Heider and Simmel [17], in which viewers watching simple geometric forms performing complex motions (i.e., behaviors) readily ascribe internal motivations and emotions to them. 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 this film, the simple dot and line are transformed into the protagonists of a compelling love story simply through animation, i.e., their complex behavior.

A second way to make EVCs more interesting and entertaining is to make their physical appearance, i.e. morphology, more complex. While morphology can be interesting in its own right, it can be especially powerful when it matches behavior—that is, when the EVC utilizes behavior that naturally emerges from its morphology. Unlike the almost complete stagnation in behavioral complexity since Sims, there has been a slow but steady (and recently increasing) interest in more complex bodies for evolved virtual creatures.

For example, Bongard and Pfeifer's [5] bodies built of multiple spheres provide one potential avenue for complexifying EVC bodies. They are based on using a larger number of smaller primitives to produce a more finely grained physical description. Similarly, soft-bodied EVCs with voxel-based body descriptions [8] increase morphological resolution. This approach makes it possible to fine tune creature form in ways that arrangements of a few predefined primitives cannot.

Another approach is to apply evolutionary techniques to the same goal with traditional block-based EVC morphologies. For example, Lehman and Stanley [26] motivate morphological diversity explicitly through novelty search and niching. Taking a slightly different approach, with a strong focus on the value of physical intelligence [36] and a novel metric for morphological complexity, Bongard and Auerbach pursued this goal by defining creatures using Compositional Pattern-Producing Neural Networks (CPPNs) [2], and by attempting to use complex environments to promote complex bodies [4, 3].

Given that behavioral and morphological complexity is important, why has more of them not been seen in EVCs to date? It seems that behavioral complexity has been limited by the monolithic developmental process of typical EVC evolution. Using that approach, it is difficult to make the kinds of leaps in brain architecture that complex behaviors require. On the other hand, morphology evolution has been limited, missing the kind of complex underlying mechanics found in natural bodies. However, if the large-scale process could be directed through human intuition, it might be possible for evolution to solve the details; if more of body mechanics, such as muscles and joints, can be incorporated into the morphology, more complex bodies should emerge. Increasing behavioral and morphological complexity in this manner is the challenge and opportunity taken on in this dissertation.

1.3 Approach

Motivated by the potential value of behavioral complexity for content creation, this dissertation describes a method, ESP, for constructing significantly more complex EVC behaviors than have been seen before. 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 creature 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. A mechanism inspired by Selfridge's *pandemonium* [41] is used to resolve disputes between competing skills or drives within the increasingly complex creature.

These behaviors are evolved in a brain consisting of a network of nodes connected by wires. Each node computes a simple function, making it easier to construct high-level behaviors and to implement encapsulation. Network topology and all node properties emerge entirely through evolutionary processes.

To create useful morphological complexity, this dissertation proposes an approach that extends the traditional EVC body with biologically motivated elements. The traditional Sims-like segments are envisioned as approximations to bones, and muscle-like springs are added to them for actuation. In addition to a mechanically intricate increase in morphological complexity, this bone and muscle structure results in physical intelligence. By encoding part of the behavior in the creature's body, physical intelligence has the additional benefit of making it easier to evolve more complex behaviors, since the demands on the brain are reduced.

1.4 Overview

This dissertation is organized as follows:

Chapter 2 describes the basis upon which the ESP approach is built: the foundational concepts and related work on evolutionary algorithms and evolved content.

Chapter 3 presents the basic EVC system, which reproduces Sims' original work with some changes, including those that support more complex bodies and behaviors. As a basis for evolving more complex behaviors in Chapters 4 and 5, it will be demonstrated on the task of evolving a variety of creatures for locomotion.

In chapter 4, the ESP method is employed to approximately double the state of the art in behavioral complexity for evolved virtual creatures (Figure 1.3a). This chapter focuses almost entirely on behavior; the morphology evolved in Chapter 3 is kept largely fixed throughout most of this process.

In Chapter 5 (Figure 1.3b), an extension to the ESP system is described which allows full morphological development to continue throughout a larger portion of the syllabus-based evolutionary process. This technique can produce creatures with greater variety and greater fitness when applied to learning goals with multiple and differing morphological requirements. This ability is demonstrated by evolving a diverse set of solutions to the strike and high-reach tasks using the creatures from Chapter 3 as a starting point.

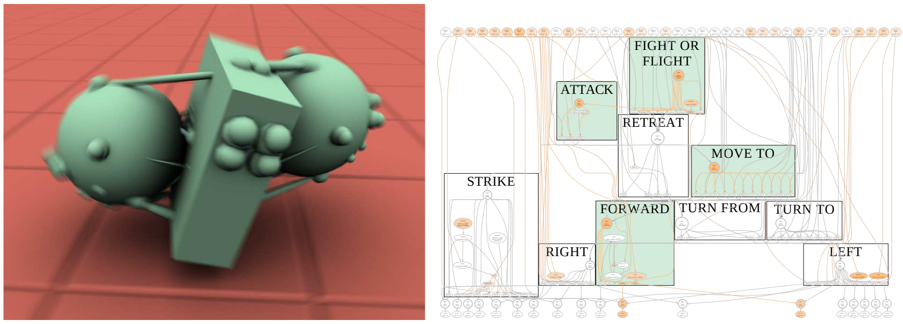

(a) ESP. This system produces a significant increase in behavioral complexity over traditional EVCs.

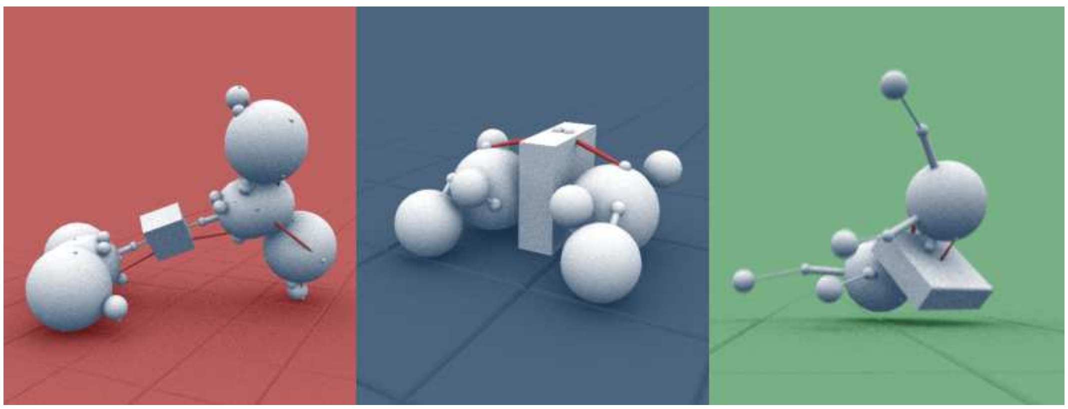

(b) Extended ESP. This augmented version of ESP allows greater body adaptation to multiple tasks than what Original ESP allows.

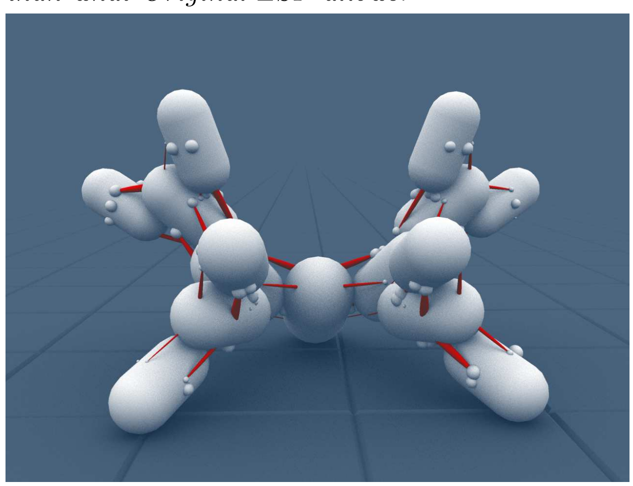

(c) Muscle Drives. This novel drive system simultaneously increases morphological complexity and reduces computational demands on the brain.

Figure 1.3: This dissertation's primary contributions, which improve both the behavioral and morphological complexity of evolved virtual creatures.

In Chapter 6, the muscle drive system employed throughout the dissertation is evaluated in detail (Figure 1.3c). It is shown to produce two important benefits for evolved virtual creatures. First, it manifests a significant measure of intelligence that would otherwise be hidden in the brain, adding to the creature's morphological complexity. Second, transferring intelligence from the brain to the body reduces the computational load that must be borne by the brain, making it easier to evolve more complex control.

These three new methods—ESP, extended ESP, and muscle drives for EVCs—comprise the primary contributions of this dissertation. The degree to which they are successful in creating valued content, and can be extended in the future to provide creature content on demand will be discussed in the last two chapters.

1.5 Conclusion

This dissertation presents significant new steps in the pursuit of increasing complexity in evolved creature content. This goal is pursued as a rewarding avenue for improving the entertainment value of evolved virtual creatures. With continued progress down this path to more compelling EVCs, we may someday witness the dawn of a new age of naturalistic exploration, in which the borders of our world are once again worthy of the phrase *Here Be Dragons.*

Chapter 2

Background

The results of this dissertation are built on the foundation of evolutionary algorithms, evolution of content, and in particular, evolved virtual creatures (EVCs). These are discussed in this chapter, both as foundational material and for their relation to the dissertation's contributions.

2.1 Evolutionary Algorithms

At the heart of this dissertation's contributions is the evolutionary algorithm. Inspired by the workings of evolution in nature, this algorithm is at once extremely powerful and extremely versatile.

In order to find valuable results, an evolutionary algorithm starts with a population of candidate solutions—locations within the space of all possible solutions. These candidate solutions are *genotypes*—individual genetic encodings, analogous to the set DNA of a single individual in nature. Also defined for the algorithm is a method for mapping from genotypes to the *phenotypes* which they encode (analogous to individual creatures in nature). The process of converting a genotype into a phenotype is called *expression* (Figure 2.1).

The evolutionary algorithm itself functions as depicted in Figure 2.2. Individual genotypes may be encoded in a variety of ways, from a simple binary string to a grammar-like rule list to a complex directed graph. Starting with the population of genotypes (possible solutions), each one is expressed as the corresponding phenotype. The expression process also varies, sometimes mapping directly from the genotype, and sometimes involving organic growth-like processes, in which a complex final result emerges in hard-to-predict ways from a relatively simple genotypic encoding. The expressed phenotype is evaluated using a *fitness function*, which may be accomplished by many means. Phenotypes may be evaluated in simple mathematical terms such as the number of ones in a binary string or the physical dimensions of a virtual creature's body, or they may be put through a battery of tests in physical simulation, or they might even be scored using human input. In fact, this flexibility of genotypic encoding, phenotypic expression, and fitness evaluation is a particular advantage of evolutionary algorithms: As long as a fitness function can be computed from it, few other constraints on a genotype are required.

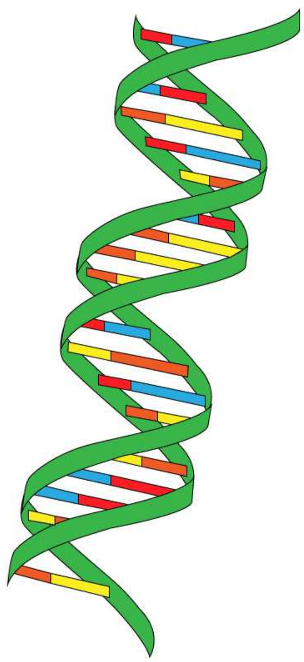

(a) *Genotype: complete genetic encoding for a single individual.*

(b) *Phenotype: the individual produced by expressing a genotype. (José Mujica, 40th President of Uruguay*

Figure 2.1: A genotype (a) is expressed *to produce a phenotype (b).*

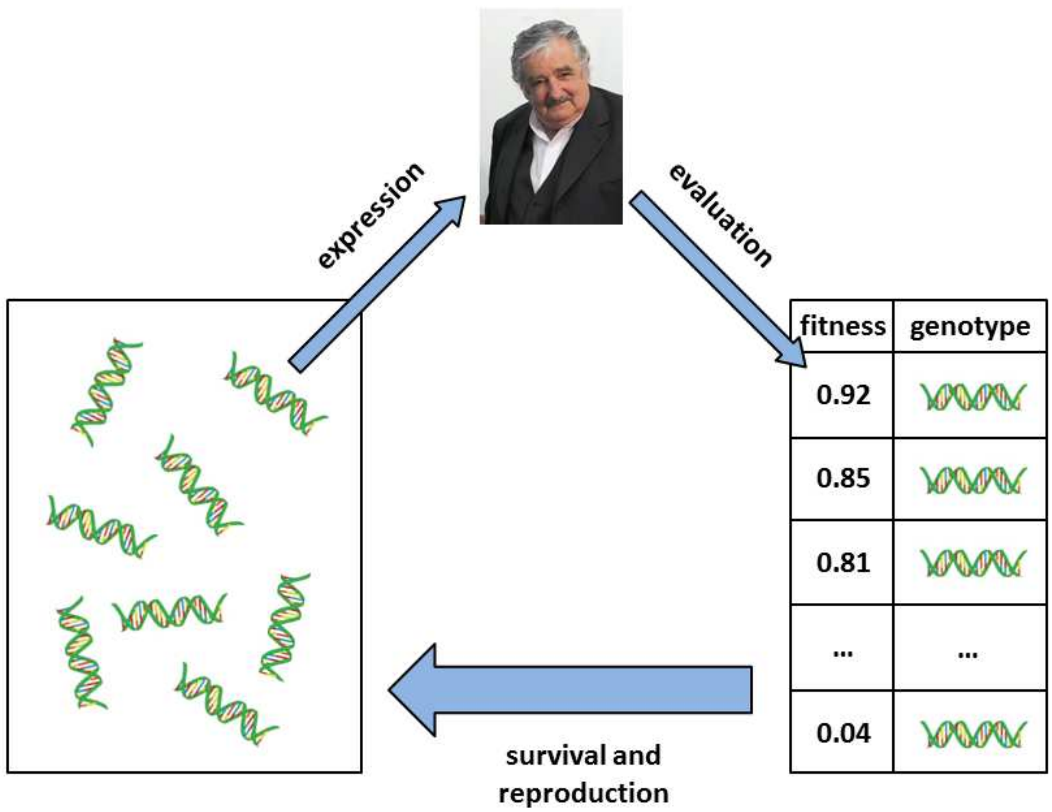

Figure 2.2: Evolutionary algorithm. A population of genotypes (left) are individually expressed to produce a phenotype (center), which can then be evaluated by a fitness function. This produces a fitness value for each genotype (right), which determines the survival and reproduction of genotypes to produce the next generation. As the cycle is repeated, fitness in the population tends to improve.

Once the population of genotypes has had its fitness evaluated, the next generation can be created. As in nature, fitter individuals are more likely to survive and reproduce, thereby making up a larger proportion of the next generation's population. Typical ways to select parents include fitness-proportionate selection (likelihood proportional to fitness value) and rank selection (likelihood proportional to fitness rank). Using crossover, selected parents produce new individuals by splicing together their genotypes. For example, two string-like genotypes can be combined by choosing a random point in the encoding, and copying everything up to that point in one and everything after that point in the other. Analogously, graph-based parent genotypes can combine randomly selected subgraphs to produce a child genotype. As an alternative to crossover, individuals may be produced by mutation of a single selected genotype. Mutation can range from simple local modifications to a genotype, such as flipping a bit in a binary-string encoding, to relatively complex topological changes in a graph-based encoding. Highly fit genotypes may also be copied directly from one generation to the next without modification, which is referred to as *elitism* [32]. These processes continue until an entire new population has been created to replace the old one.

From that point, a new round of fitness evaluations can occur, and the process repeats. Over time, the prevalence of fit individuals in the population tends to increase. With a sufficient number of generations, impressive demonstrations of creativity and problem solving are often observed, as seen in results from fields as diverse as satellite antenna design [19], vibration-reducing truss design [33], wind farm layout [16], protein structure prediction [52], aircraft wing design [34], and many more. The work in this dissertation builds on this creativity as well.

2.2 Evolved Content

One particularly relevant application of evolutionary algorithms is the production of content—making things that are interesting, entertaining, or compelling to human viewers.

By making the fitness function dependent on user perception, visually appealing images can be created. An early example of this (1986) was Richard Dawkins' *Blind Watchmaker* algorithm [10], which was created to demonstrate basic ideas about evolution. In this system, genotypes were defined by nine numerical values, expressed as phenotypes that were simple, yet organic-looking tree-like structures, which Dawkins referred to as *biomorphs* (Figure 2.3). Through interactive evolution based on user preference of phenotype images, more visually pleasing biomorphs were evolved–an early example of evolved content, which is the ultimate focus of this dissertation.

In 1991 [44], Sims created evolved content with interactively evolved growth rules for generating plant designs (Figure 2.4). While Dawkins' biomorphs were simple two-dimensional images composed of straight lines, these plants were three-dimensional and

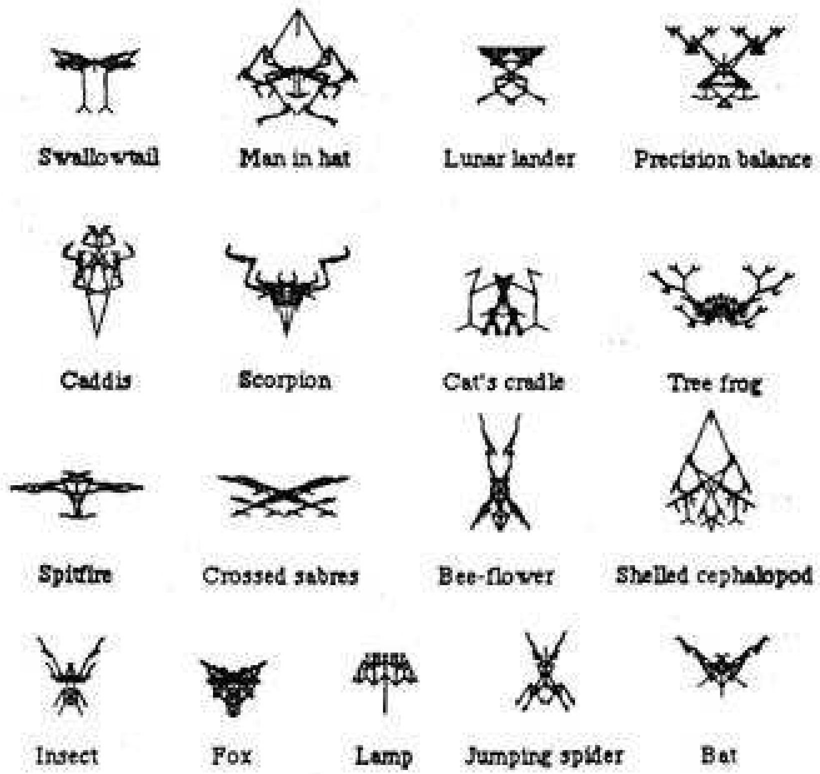

Figure 2.3: Biomorphs. An early example of interactive evolution of visual imagery from Richard Dawkins' Blind Watchmaker *algorithm [10]. Although created to demonstrate principles of evolution, and despite the simplicity of their construction, these results have clear aesthetic appeal.*

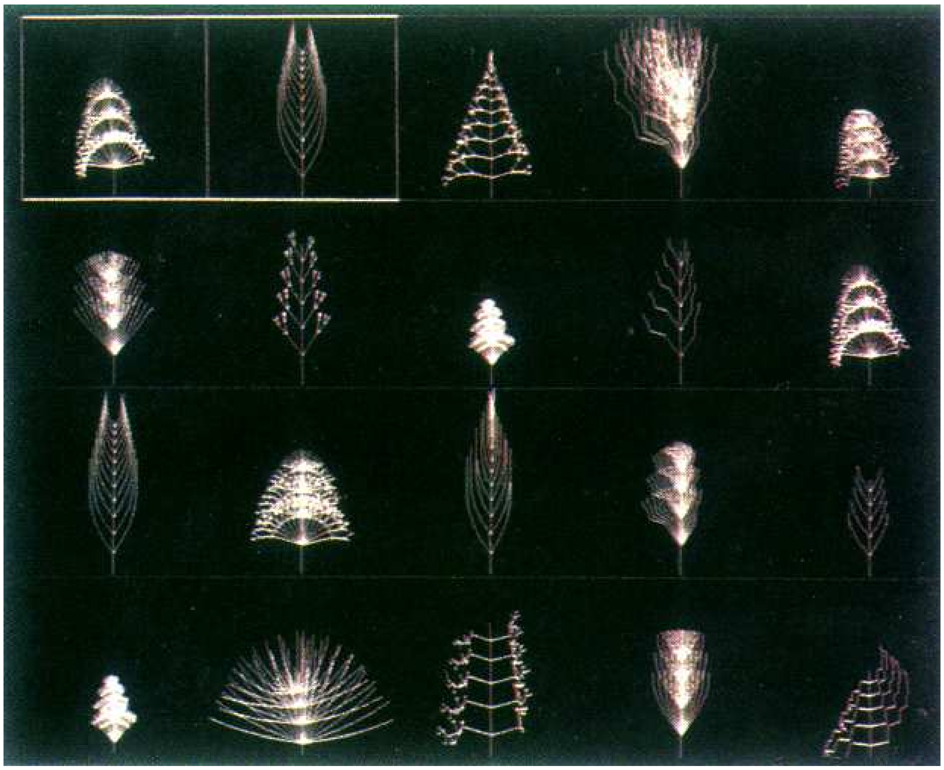

Figure 2.4: Plant forms evolved based on their visual appearance from a system designed by Sims [44]. In a system similar to Dawkins' Blind Watchmaker, genotypes (in this case, growth rules) are selected for evolution using fitness determined by a human user viewing the plant-like phenotypes. The resulting content was of sufficient quality to be included in a short film by Sims: Panspermia *(1990).*

Figure 2.5: Images evolved based on their visual appeal from a system by Sims [44]. In this system, the genotypes are hierarchical lisp expressions, but the basic interactive-evolution selection method is essentially unchanged from the one used in Figure 2.4. This further example of evolved content is notable for its persistent visual appeal.

more complex, with 21 evolved parameters instead of Dawkins' nine. This advance demonstrated more complex and aesthetically appealing content creation through evolution, with the resulting phenotypes having sufficient visual appeal to play a central role in Sims' short film *Panspermia* (1990).

In addition, in the same publication, Sims demonstrated an early yet powerful implementation of interactively evolved images (Figure 2.5), with genotypes encoded as hierarchical lisp expressions. With a human user as the selector, these images remain arguably unsurpassed in their visual appeal in evolutionary image generation.

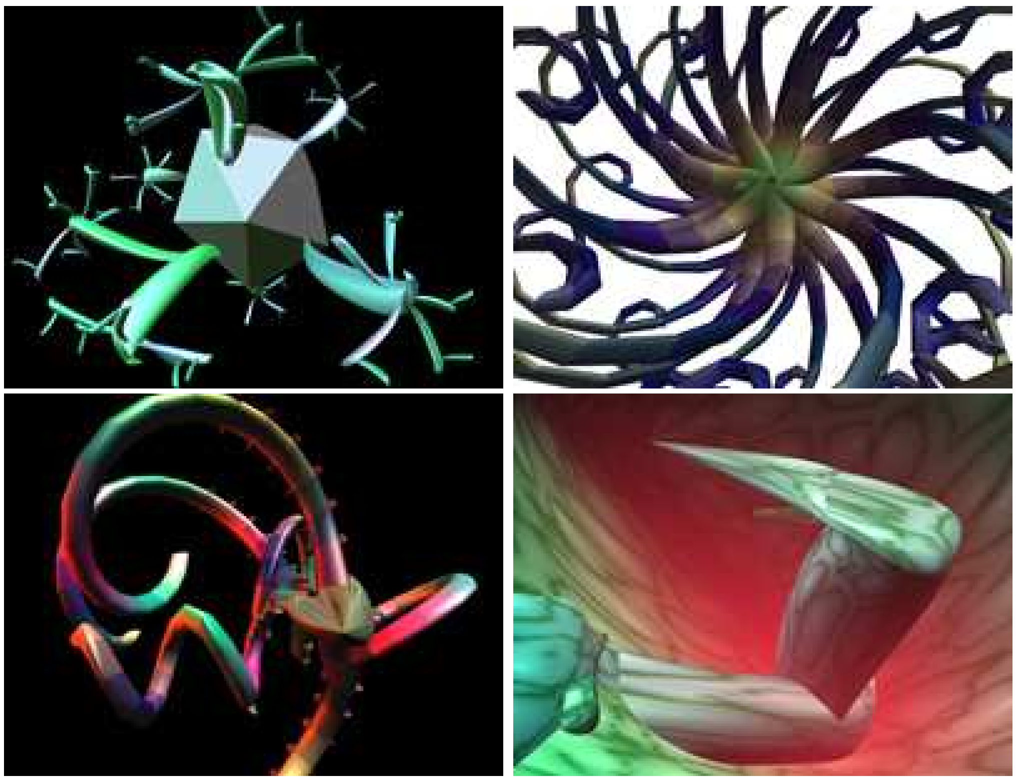

Figure 2.6: Evolved animated content from Sims' 1997 Galápagos exhibit (Intercommunication Center, Tokyo). This work employed an early form of crowdsourcing implemented using in-person interaction with sensors at the museum installation.

In 1997, Sims presented the *Galápagos* museum installation (Intercommunication Center, Tokyo)—an animated three-dimensional version of the same concept, even employing an early form of crowdsourcing to obtain fitness evaluations. Viewers stood on sensors in front of displayed creatures that they preferred, leading to increased fitness for those individuals, and increased odds of survival in the future (Figure 2.6).

In 2008, *Picbreeder* [40] made a result similar to Sims' two-dimensional evolved images available for worldwide user input by way of the internet, producing visually appealing

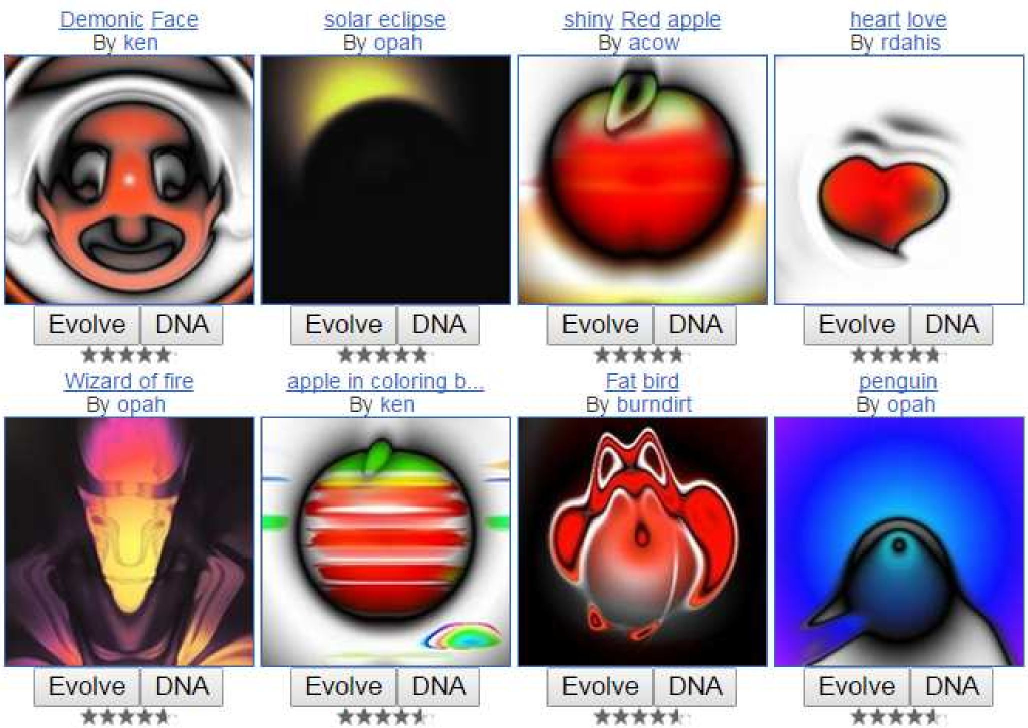

Figure 2.7: Images produced using worldwide crowdsourcing in Picbreeder [40], a prominent example of evolved content. Similar in concept to the evolved images of Sims, these results encode their genotypes as Compositional Pattern-Producing Networks [48].

results such as those seen in Figure 2.7. In this example, genotypes are encoded as Compositional Pattern-Producing Networks (CPPNs) [48]—similar in concept to the hierarchical lisp expressions used in Sims' work. Due to its internet-based open-ended nature, this is one of the more prominent examples of evolution of content.

Figure 2.8: Results from Draves' cooperative evolution of animated imagery Electric Sheep *[12]. With highly effective interaction and world-class aesthetic results, this system is an example of valued evolved content.*

The interactive evolution of imagery achieved perhaps its highest expression with the development of Draves' *Electric Sheep* [12], a web-based crowdsourcing system for the development of animated visuals based on his *fractal flame* algorithm, with powerful aesthetic appeal (Figure 2.8). With a particularly well-conceived interaction method, as well as world-class aesthetics, this system has achieved significant recognition for its visual beauty beyond the academic world.

Most recently, *Endless Forms* [9] employed an extended version of the encoding used by Picbreeder to produce interactively evolved three-dimensional shapes suitable for 3-D

Figure 2.9: 3-D printed shapes evolved using Clune and Lipson's Endless Forms *[9]. This extension of Picbreeder's web-based interactive concept into three dimensions demonstrates a potential application for the generation of content through physical evolution.*

printing (Figure 2.9), also using online user input. This extension of interactive evolution provides an example of how physical content can be produced through artificial evolution.

In these examples, the ability of evolutionary algorithms to produce content is made particularly plain. It is this font of creativity that is harnessed for evolved virtual creatures, i.e. the core application at the heart of this dissertation.

2.3 Evolved Virtual Creatures

For the purposes of this dissertation, evolved virtual creatures (EVCs) are defined as digital organisms having co-evolved bodies and brains, evaluated in physical simulation. The foundations of this dissertation—the definition of EVCs and the established research with them—is discussed in this section.

(a) A creature evolved for locomotion in water.

(b) A creature evolved for locomotion on land.

(c) A creature evolved to follow a light source.

Figure 2.10: Sims' evolved virtual creatures [46]. Sims evolved creatures for locomotion in water and on land, to jump, and to follow a light source (phototaxis).

2.3.1 Sims' EVCs

The first and most influential example of evolved virtual creatures are due to Sims ([46]; Figure 2.10). The advances of this dissertation are also built on this foundation, as described in detail in Chapter 3.

The genotypes for Sims' creatures were directed graphs, able to encode complex body structures, as shown in Figure 2.11a. The bodies of these creatures were composed of boxes, connected by joints with varying degrees of freedom and evolvable limits to their revolution. Actuation was provided by implicit joint motors, able to apply force at every degree of freedom of every joint.

Corresponding to the body, Figure 2.11b depicts the brain, with nodes computing simple functions and signals carried between nodes by evolved connections. In Sims' im-

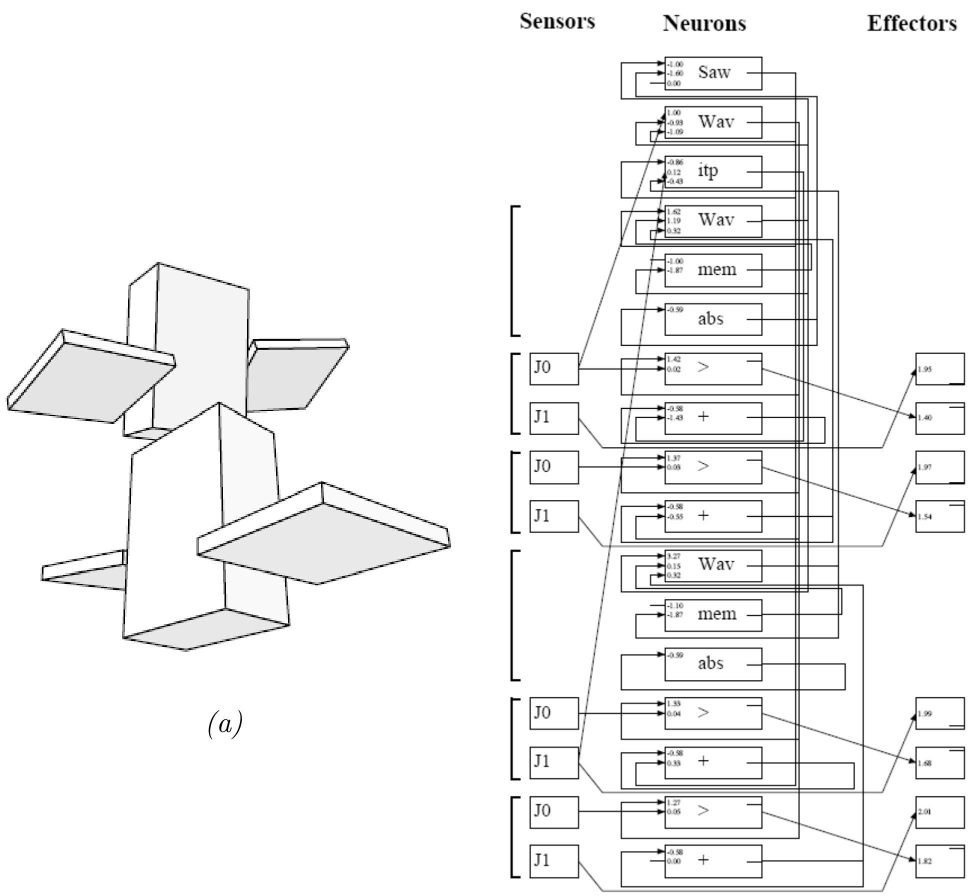

(a)

(b)

Figure 2.11: This figure illustrates a body and brain from one of Sims' conventional EVCs—this one evolved for locomotion [46].

plementation, brain elements can be embedded within body segments, where they can take advantage of the same kinds of repetition and recursion as the creature's morphology.

Evolution in Sims' system proceeds as described in Section 2.1, making use of fitness-proportionate selection, crossover, mutation, and elitism. Interestingly, while Sims' computation was performed in a massively parallel fashion on a Connection Machine CM-5 (described at the time as a supercomputer), similar work can now be performed on a typical desktop machine.

Using this system, Sims demonstrated impressive results in multiple tasks. A variety of creatures were evolved for locomotion, both on land and in water (Section 2.3.2), and creatures with the ability to jump off the ground were produced. Most impressively, Sims demonstrated creatures evolved for phototaxis (light seeking) behavior (Section 2.3.3), a level of behavioral complexity not exceeded until the work presented in this dissertation [27] (Chapter 4). Many of these behaviors have since then become benchmarks in EVCs. They have been replicated many times on different platforms, and other similar behaviors have been added to this repertoire, as will be described next.

2.3.2 Locomotion

The standard benchmark task for an EVC system is locomotion. Sims presented locomotion on land and water (Figures 2.10a and 2.10 b), and this result has been repeated for many different purposes by many different researchers. An illustration of the breadth of these results is seen in Figure 2.12

Lipson and Pollack evolved creatures for locomotion in a system that allowed the results to be 3-D printed and activated in the real world (Figure 2.12a) [29]. Creatures composed of rigid segments and linear actuators were evolved for locomotion in physical simulation. The body parts (including joints) could then be 3-D printed, and only the fitting of actuators required special attention during assembly. Notably, these creatures were required to maintain static stability at all times (i.e., have their center of gravity always supported by the body). In this manner, a consistent transition to the real world was guaranteed, where dynamics might differ from simulation, but geometry should not.

Shim and Kim evolved virtual creatures for another type of locomotion—flight (Figure 2.12b) [42]. Lehman and Stanley used locomoting EVCs as the subject for an investigation of novelty promotion (Figure 2.12c) [26]. Cheney et al. (Figure 2.12d) demonstrated their new encoding for soft-bodied EVCs by applying them to the locomotion benchmark [8].

As a preliminary step, locomotion results from this dissertation's EVC system are presented in Section 3.6. In Chapter 4, this fundamental ability will be used as a starting point for the incremental acquisition of more complex behaviors.

(a) *Automatic Design and Manufacture of Robotic Lifeforms [29].*

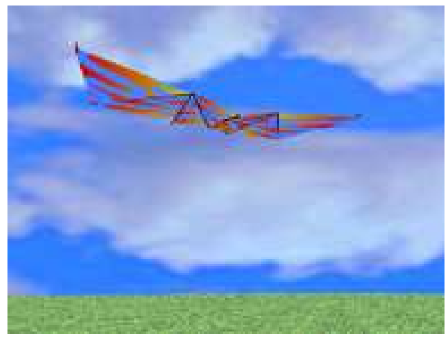

(b) *Generating Flying Creatures Using Body-Brain Co-Evolution [42].*

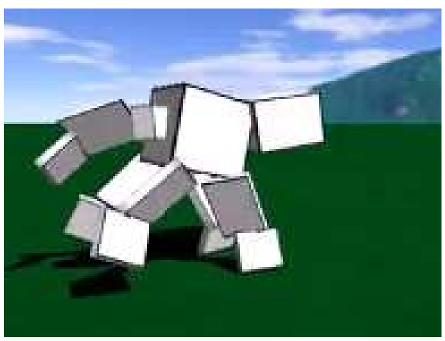

(c) *Evolving a Diversity of Creatures through Novelty Search and Local Competition [26].*

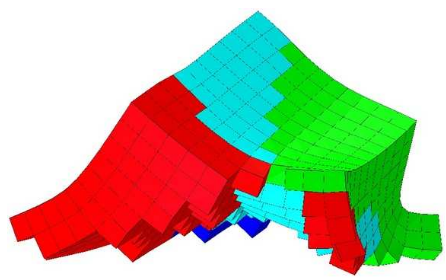

(d) *Unshackling Evolution: Evolving Soft Robots with Multiple Materials and a Powerful Generative Encoding [8].*

Figure 2.12: A selection of work indicating the breadth of research using locomotion in EVCs, including physical robots, soft robots, flying robots, and diversity promotion.

2.3.3 Phototaxis

Phototaxis (the ability to move to a light source) was the most complex behavior demonstrated by Sims (Figure 2.10c). By testing the ability to move toward a light target placed at multiple positions, creatures were developed with a generalized ability to perform phototaxis. This remained the most complex EVC behavior for almost two decades until the work of this dissertation (Chapter 4) approximately doubled that complexity (measured as the number of discriminable behaviors, as defined in Section 1.2).

Pilat and Jacob reproduced the behavioral complexity of Sims' phototaxis approximately in their 2010 work [37], although their implementation differed in some respects. While Sims' photoreceptors were embedded in each body segment and produced signals relative to the segment's orientation, Pilat and Jacob used a single sensor for the entire creature, and that sensor's signals were preprocessed to give one output for heading to the light and another for the light's elevation angle. Also, Pilat and Jacob's creatures had simpler morphology, allowing only single-degree-of-freedom hinge joints between segments. Unlike the control networks of Sims, with nodes computing a variety of predefined functions, Pilat and Jacob's creatures used a more conventional artificial neural network (ANN).

Shim and Kim also achieved a similar result in 2004 with flying creatures able to follow paths [43]. This dissertation's EVC system demonstrates phototaxis as an intermediate step on the path to more complex behaviors (Section 4.3.5).

2.3.4 Combat

Another important EVC result, especially with respect to entertainment value, is combat. Sims demonstrated a stylized form of combat in his block-control competition [45]. In it, creatures compete in one-on-one contests for control of a target block. Successful results employ a variety of entertaining strategies (Figure 2.13), leading to a new form of particularly compelling EVC content.

In 2008, Miconi [30] implemented a more direct form of EVC competition (Figure 2.14), again with very entertaining results. In Miconi's relatively realistic implementation of combat, creatures are evolved for their ability to damage each other with physical impacts. In the process, many of Sims' accomplishments were replicated, including both locomotion and phototaxis. As Miconi anticipates, a fully convincing implementation of collision damage would require bodies composed of different materials (otherwise, every collision may well inflict equal damage on both the attacker and the victim). However, even with a simplified damage model, Miconi produces a broad variety of compelling creature-combat results. Note that this dissertation's proposed future work on morphological complexity (Section 7.3) might well enable exactly the kind of body-material trade-offs needed to take Miconi's EVC combat to the next level of realism.

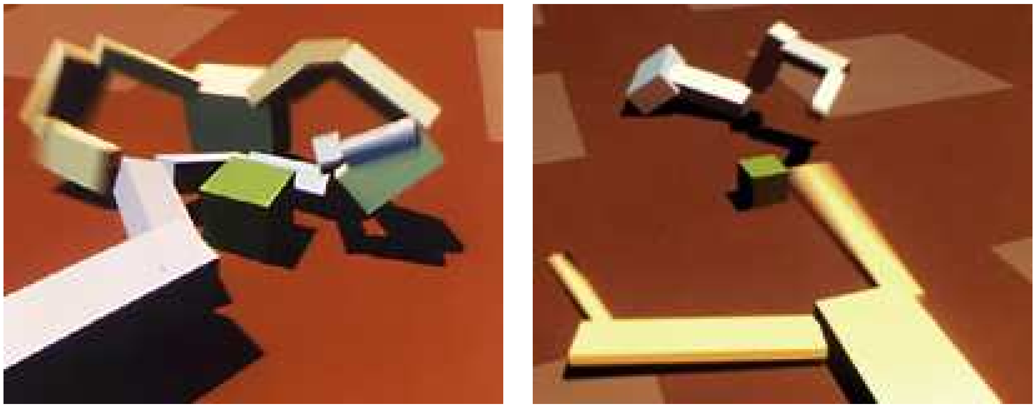

Figure 2.13: Evolved creatures compete one-on-one for control of a target block in this work by Sims [45].

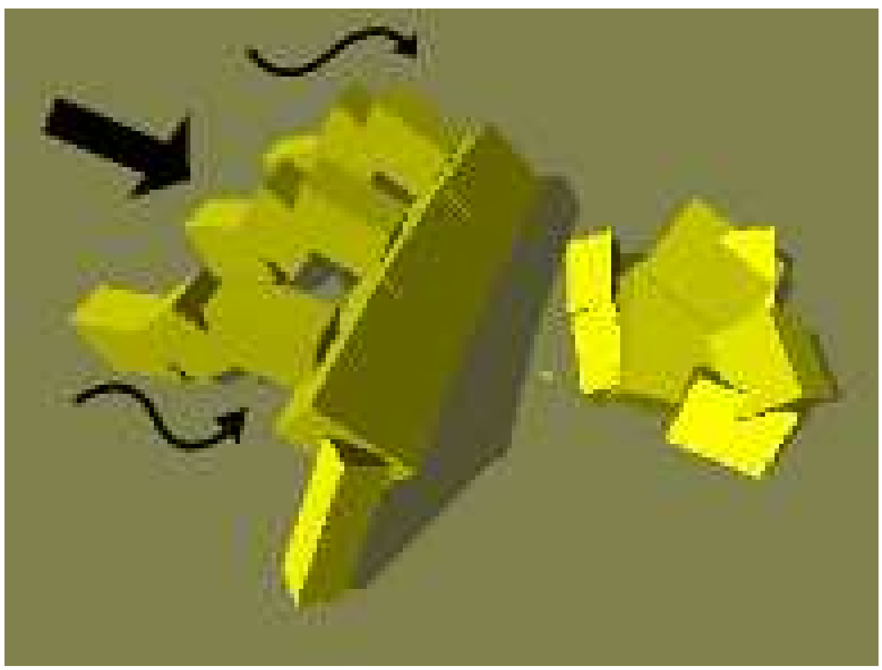

Figure 2.14: EVC combat as implemented by Miconi [30]. Creatures are evolved based on their ability to damage each other in one-on-one competition. In this image, a larger creature (left) attacks a smaller one with a steamroller-like technique. Combat is a demanding and natural goal for EVCs, requiring a combination of several behaviors. It is also part of the motivation for the work in this dissertation.

In Chapter 4, this dissertation's *ESP* system demonstrates the ability to develop skills relevant to complex and entertaining combat, such as seeking and striking a target (Section 4.3.7), fleeing from a dangerous target (Section 4.3.9), and deciding between the two actions based on perception of its environment (Section 4.3.10).

In Chapter 5, the *Extended ESP* system demonstrates combat-relevant capabilities which allow EVC morphology to adapt to multiple skills while still developing the complex behaviors made possible by ESP.

2.4 Conclusion

This chapter described the creative engine of the evolutionary algorithm, its application to the evolution of content, and in particular its use to produce evolved virtual creatures. This dissertation is built on these foundations. In the next chapter, these general techniques are specialized for this dissertation, with the presentation of the basic EVC system upon which all of this dissertation's contributions are built.

Chapter 3

The Basic EVC System

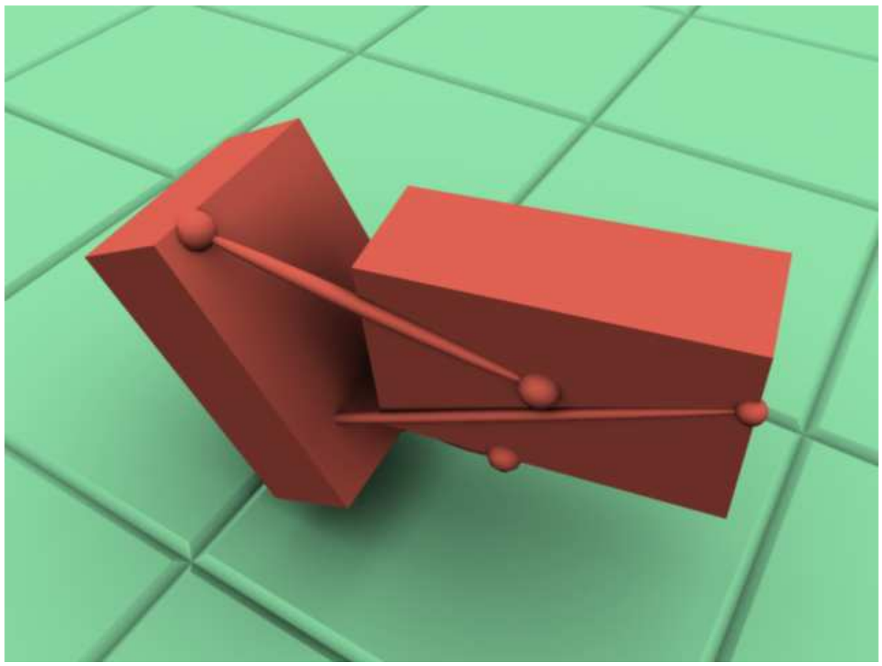

Figure 3.1: A typical result from the Basic EVC System. This creature was evolved for locomotion as described in Section 3.6, and was produced at generation 2000 of Run 1 (Figure 3.9). This phenotype is the expression of the genotype shown in Figures 3.4 and 3.5.

The Basic EVC System described in this chapter is the foundation underlying this dissertation's three primary contributions. In Chapter 4, this system is extended by a mechanism called *ESP*, which makes it possible to increase behavioral complexity dramatically. In Chapter 5, the Basic EVC System is further extended to permit not just behavior but also morphology to be adapted to multiple skills. In Chapter 6, the Basic EVC System's muscle drives (Section 3.5) are demonstrated to embody physical intelligence, which may then be removed from the brain, leaving it to develop higher functions. The Basic EVC System largely replicates the work of Sims [46] and other traditional EVC systems [7, 24, 30], although it does include some novel elements to support the technologies developed in Chapters 4-6.

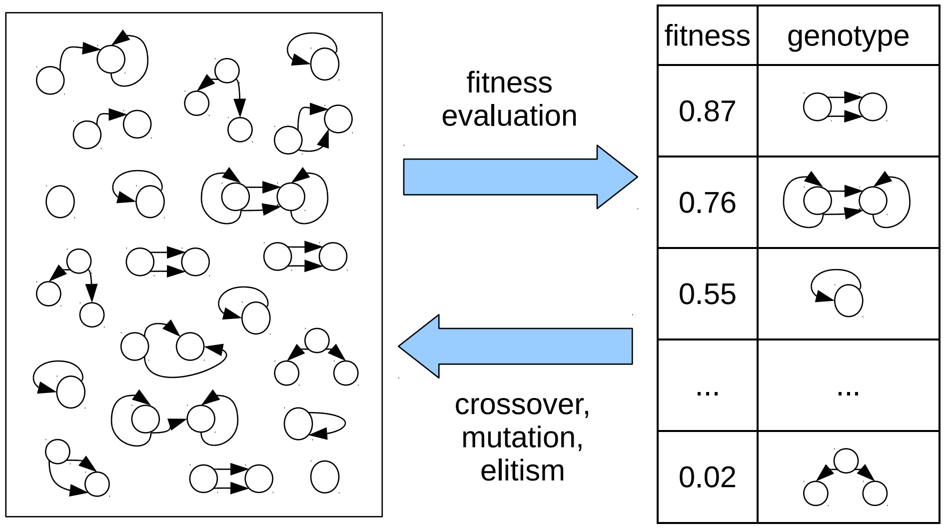

Figure 3.2: An evolutionary algorithm as used in the Basic EVC System. A population of creature genotypes (left) is evaluated for fitness, so that each one can be assigned a fitness score (right). Based on these scores, a new population is created using mechanisms such as crossover, mutation, and elitism. Over time, as this process is repeated, fitness in the population tends to improve. After a significant number of generations, one or more individuals are chosen as the winners, usually based on fitness, but possibly using other criteria as well, such as how visually appealing they are.

3.1 Evolutionary Algorithm

In the Basic EVC System, the engine of development is an evolutionary algorithm (Section 2.1). It begins with a population of genetic representations (genotypes) for virtual creatures. These genotypes are expressed to produce the physical representations of the creatures (phenotypes), and these phenotypes are evaluated in a physically simulated virtual environment implemented with NVIDIA PhysX. This evaluation (via a user-defined fitness function) produces a score for each creature's genotype, and based on these scores, genotypes are selected for survival, breeding, and mutation (Figure 3.2) to produce the next generation. Over time, the fitness of the population tends to improve. After a sufficient number of generations, the creature with the best fitness may be considered the winner, or a user may select a winner from among individuals with relatively high fitness values.

The evolutionary algorithm is conventional, making use of elitism (the intact preservation of the population's best), fitness-proportionate selection (choosing the next generation's parents based directly on their fitness), and rank selection (choosing the next generation's parents based on their place in a fitness-based ranking) [32]. In addition, the most challenging tasks employ some degree of shaping [35], a process described by B.F. Skinner in which fitness is based on increasingly close approximations to the ultimate behavioral goal [47]. A specific example is given in Section 3.6.

3.2 Encoding Morphology

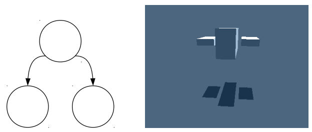

(a) Simple topology. Starting with the genotype's root node, the phenotype's root segment (the center block) is defined. Each graph edge leaving the root describes a joint, which connects to its own child node and corresponding segment (the two arms).

(b) Multiple edges for repeated substructures. In this example, a simpler genotype produces the same phenotype as in (a). Both joint edges connect to a single child node/segment, rather than to two separate child nodes, as was done in (a). As each edge is traversed, a new copy of the child node is created.

(c) Reflexive edge for recursive structure. In this example, an edge from one node to itself describes a repeating joint-node sequence. Each transition through the edge applies that joint's scale and orientation transformations to all structures below it. Recursion stops when the parent-segment's evolved recursion limit is reached.

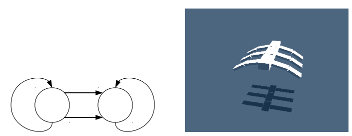

(d) Multiple and reflexive edges together. In this example recreated from Sims' original, a 15-segment bug-like phenotype is described using only two nodes and four edges in the genotype. This result is made possible by the combination of the techniques shown in (b) and (c).

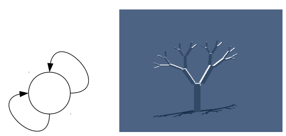

(e) Two reflexive edges. In a further application of the recursive technique shown in (b) (again recreating an original by Sims), this genotype of only one node and two edges encodes a phenotype of significant complexity and aesthetic interest.

Figure 3.3: Hand-designed genotype/phenotype pairs (as in [46]) demonstrate the encoding power inherited from Sims' original EVC system. With relatively simple genotype graph topologies such as these, complex and useful morphological phenotypes can be defined.

As in traditional EVC systems, 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 easy to define recursive and repeated body substructures, as illustrated in Figure 3.3. In addition, as in Sims' work, reflection of body parts as well as body symmetry are made easily accessible to evolution as single attributes susceptible to activation and deactivation by single mutations.

In the Basic EVC implementation, all PhysX primitives are made available for use as body segments: the boxes of traditional EVC systems, as well as spheres and capsules. Joints between segments may be of most of the types offered by PhysX, specifically: fixed, revolute, spherical, prismatic, and cylindrical. In contrast to the typical technique of evolving explicit joint limits separately, most limitations on joint movement in the Basic EVC System are provided implicitly by creature structure through natural collisions between adjacent segments. In this way, evolution of a single number—the width of a joint—implicitly defines joint limits in a straightforward and natural way.

For convenience, the genotype is stored in the text-based *graphviz* [13] format. This format is well suited to represent the genotype's graph-like nature, and may be easily converted into graph visualizations using utilities in the graphviz package such as *dot*. In this encoding, additional genotype attributes beyond graph topology are stored as comments within the graphviz file. For example, the body of the creature phenotype shown in Figure 3.1 (evolved for locomotion, as described in Section 3.6) was produced from the genotype encoded in the graphviz file shown in Figure 3.4. Figure 3.5 shows the same genotype graphviz file converted into a graph visualization by dot. With this technique, an already convenient text-based storage format brings the added benefit of a powerful data visualization system with minimal added cost.

3.3 Encoding Control

Again in a manner similar to that of traditional EVCs, creature control is provided by a brain composed of a set of nodes connected by wires, in a network identical to that of the brain genotype (Figure 3.7). 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 (the *sigma* node; see Section 3.2), the following 13 node types are allowed:

sinusoidal: Generates sine wave based on evolved amplitude, frequency, and phase.

complement: Outputs 1 - input.

constant: Outputs an evolved constant value.

scale: Multiplies input by an evolved constant.

```
1     digraph body_genotype {
2       s_0 [
3         label = "
4         primitive_type: 0
5         dimensions: 0.10  0.16  0.05
6         symmetry_flag: 0
7         symmetry_direction: 0
8         max_recursion: 1
9         "]
10      s_0_j_0 [ shape = record,
11        label = "
12        joint_type: 1
13        parent_attach_point: -0.19  -2.72
14        child_orientation: 0.42  0.60 0.43 0.53
15        child_scale: 0.98 1.01 1.03
16        mirror_flag: 1
17        mirror_direction: 1
18        terminal_only: 0
19        joint_width: 0.50
20        | {
21          parent_attach_theta_phi: -1.30 -1.74
22          child_attach_theta_phi: -2.03 2.78
23          max_spring_coefficient: 82.60 |
24          parent_attach_theta_phi: 1.24 -1.97
25          child_attach_theta_phi: -2.96 0.39
26          max_spring_coefficient: 86.68 |
27          parent_attach_theta_phi: 2.06 2.27
28          child_attach_theta_phi: -1.91 -1.31
29          max_spring_coefficient: 51.43
30        }"]
31      s_0->s_0_j_0
32      s_0_j_0->s_0
33    }
```

Figure 3.4: An example morphology genotype encoded in the graphviz *format [13]. (Some formatting data removed for clarity.) The graphviz format naturally encodes the genotype's connectivity, with additional attributes stored as comments. Lines 2-9 describe a segment node, and lines 10-30 describe joint data, including muscle data at lines 20-30. This encoding is depicted as a graph in Figure 3.5, and its expression as a phenotype is shown in Figure 3.1.*

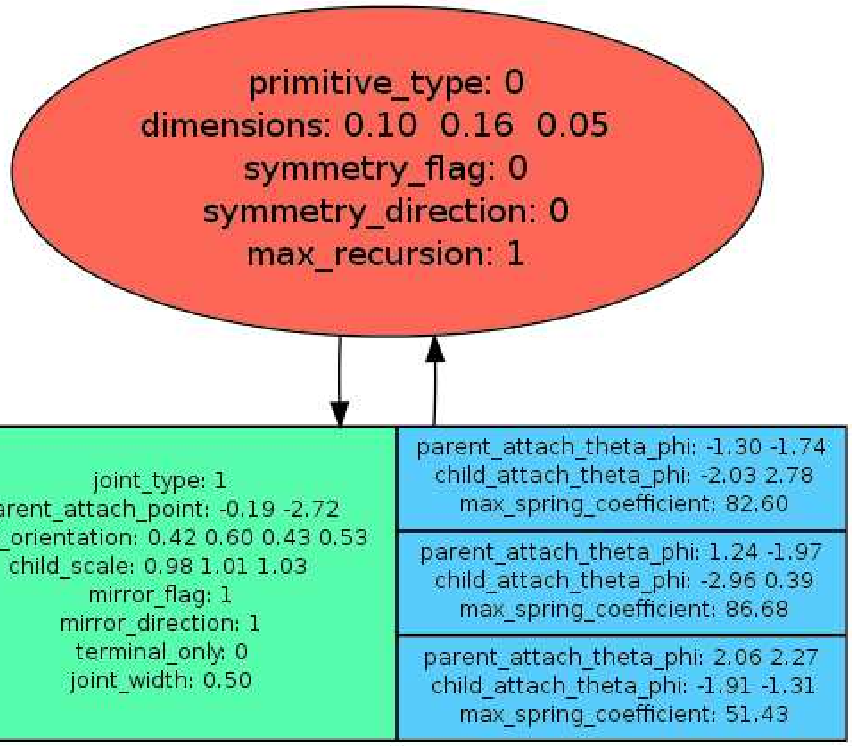

Figure 3.5: The morphology genotype encoding of Figure 3.4 rendered as a graph. Oval-shaped nodes (red) encode body segments, with square nodes recording attributes of joints (muscle attributes in blue, other joint attributes in green). Segment and joint attributes stored as comments in the encoding are visible as text within the nodes. Note that in this format, each joint-encoding edge produces two edges in the rendered graph—one from the parent segment to the joint record, and one from the joint record to the child segment.

```
1  digraph brain_genotype {
2    s_0x9079018 [label = "muscle \n target: 0:1"]
3    s_0x1c347198 [label = "proprioceptor \n target: 0:0"]
4    s_0x21eafe10 [label = "muscle \n target: 0:0"]
5    s_0x11da5b48 [label = "proprioceptor \n target: 0:1"]
6    s_0x8d0d528 [label = "complement"]
7    s_0x3c91dc10 [label = "switch \n threshold: 0.116095"]
8    ...
9    s_0x17831a38->s_0x14627dd8 [label = "0"]
10   s_0x14627dd8->s_0x17831a38 [label = "0"]
11   s_0x10d757f8->s_0xc2fcfc0 [label = "1"]
12   s_0xc2fcfc0->s_0x38cf4d68 [label = "1"]
13   ...
14 }
```

Figure 3.6: An example control genotype encoded in graphviz. (Some data removed for clarity.) As with the morphology genotype, the graphviz format encodes connectivity, and additional attributes are stored as comments. Lines 2-7 describe nodes, and lines 9-12 describe wires connecting nodes. This encoding is depicted as a graph in Figure 3.7.

multiply: Outputs the product of two inputs.

divide: Divides first input by second input.

sum: Outputs the sum of two inputs.

difference: Subtracts second input from first.

derivative: Outputs difference between current and previous input, scaled to units of change per tenth of a second, with evolvable direction flag.

threshold: Outputs one if input is above evolvable threshold, zero otherwise.

switch: If first input is at or above evolvable threshold, output second input, otherwise output zero.

delay: Applies an evolvable delay to input signal.

absolute difference: Outputs the absolute difference between two inputs.

This set of nodes is inspired by the similar collection of Sims, who had 23 node types of similar style and power [46], and is intended to offer a broad variety for evolution to employ in ways that might be hard to predict.

Just as for the morphology genotype, the control genotype is stored in the graphviz format (Figure 3.6), with the graph defining topology, and extra attributes stored in comments. Unlike the morphology, however, the phenotype control system is directly encoded in the genotype—it is essentially a direct copy of the genotype's control graph (Figure 3.7).

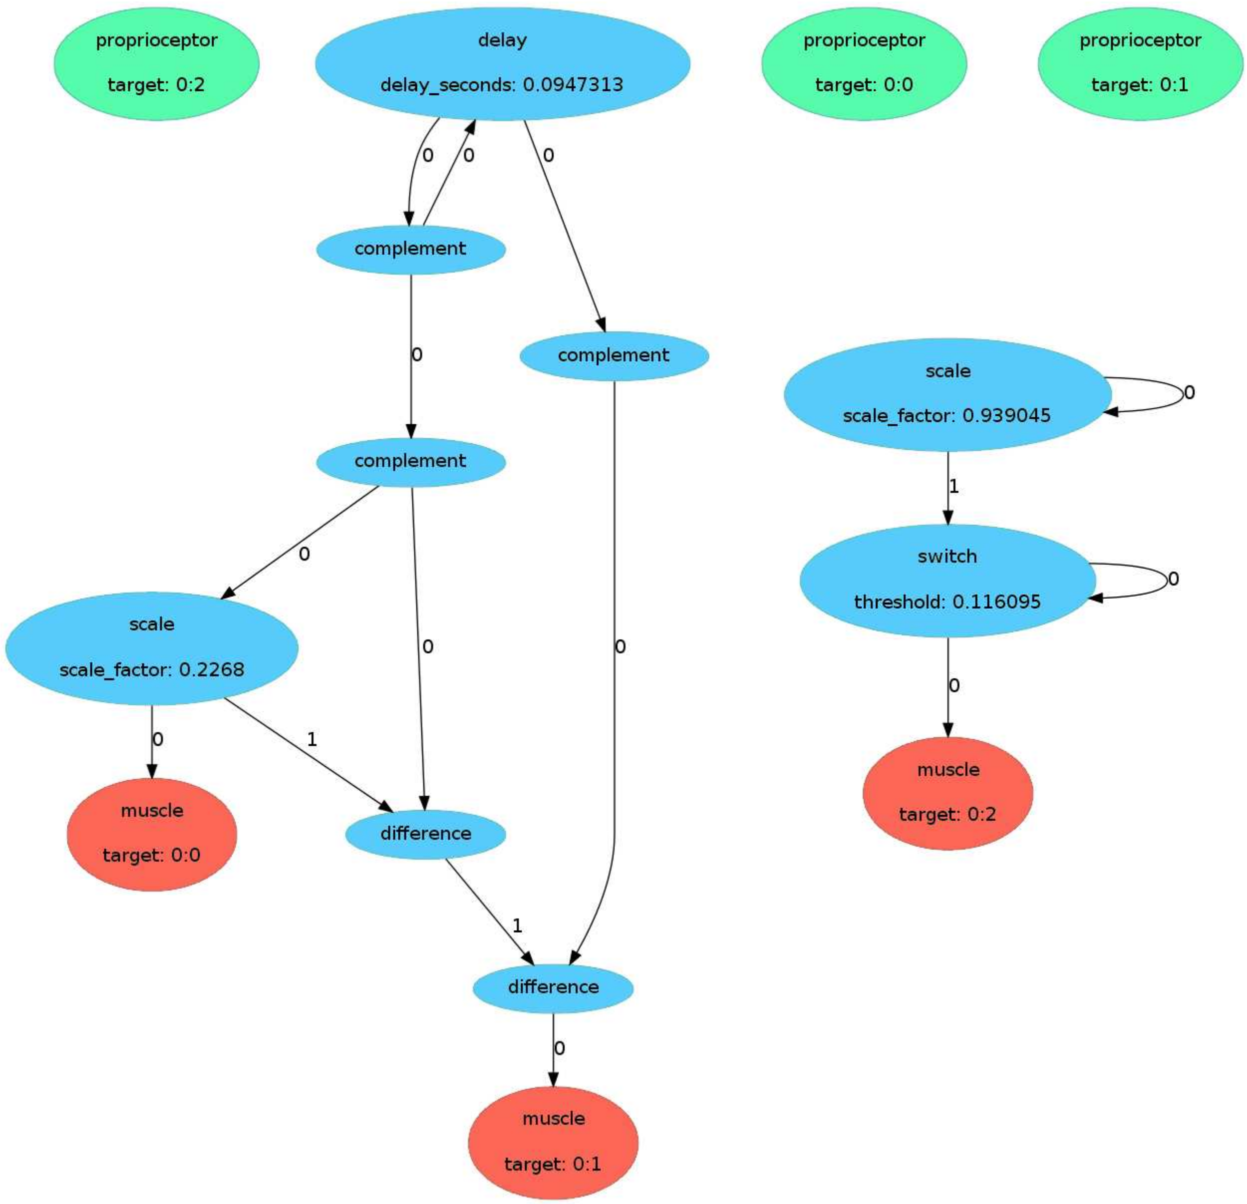

Figure 3.7: The control genotype graphviz of Figure 3.6 rendered as a graph. Muscle nodes are colored red, proprioceptors are colored green, and all other nodes are colored blue. Unlike morphology, control is encoded directly, with the phenotype control graph being essentially a direct copy of the genotype control graph.

3.4 Photoreceptors

For tasks involving light sensing, creatures are allowed to develop simple photoreceptors ((a) in Figure 3.8), 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. In contrast, in Sims' implementation [46], each segment had exactly one photoreceptor, always aligned with the segment, eliminating the possibility of physical intelligence embodied as eye placement.

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, and multiple lights are allowed. Let i be the light's intensity, θ the angle between the direction to the light and the sensor's orientation, d the distance to the light, and H the *falloff half-distance* (used so that the light's effect is halved at this distance from the sensor). Then, when a light is unoccluded, its contribution s to the sensor's signal is

$$s = \frac{i\cos(\theta)}{1 + (d/H)}. \tag{3.1}$$

For each photoreceptor in the body, a corresponding brain node is added which makes the receptor's output signal (clamped to the range [0,1]) available to the rest of the brain.

3.5 Muscles

In a break with traditional EVC systems, which typically use forces exerted directly at joints, the Basic EVC System uses simulated muscles as actuators. Each muscle ((b) in Figure 3.8) 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. If the activation signal (in [0,1]) is a, the evolved maximum spring constant is k_{max}, and the muscle's current length is x, the resulting force F applied by the muscle is

$$F = (ak_{max})x. \tag{3.2}$$

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, accepting an input to set the muscle's activation, and another node is added that makes the muscle's proprioceptive output signal available to the rest of the brain. The proprioceptive signal is linear with respect to muscle length, 0.5 at the muscle's initial length, and clamped to [0,1]. So, if the muscle's initial length is l_{init} and its current length is l_{curr}, the proprioceptive signal p (before clamping) is

$$p = 0.5\frac{l_{curr}}{l_{init}}. \tag{3.3}$$

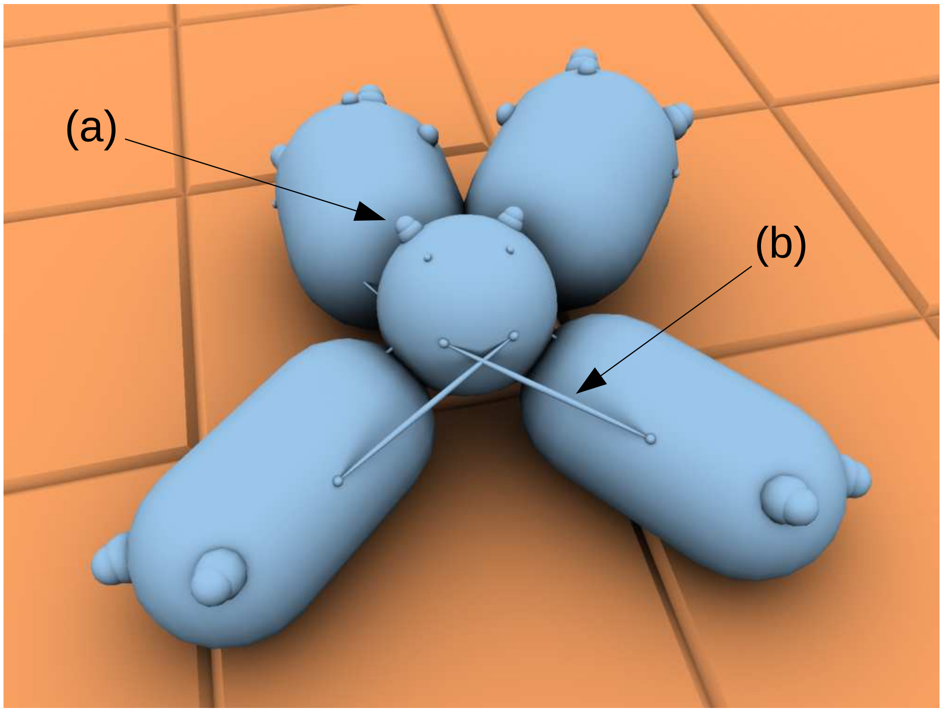

Figure 3.8: Photoreceptors (a) and muscles (b) bring sensing and actuation to creatures in the Basic EVC System. For both, function depends upon placement, so creature form develops meaningfully as capabilities are evolved.

Muscle drives benefit EVCs in several ways. They are flexible, as they can be used even on creatures without joints. They are efficient, since effectors need only exist where useful, not at every degree of freedom of every joint. And they are beautiful, tapping into the human affinity for elegant, functional body structure. These benefits of the evolved musculature system are explored in depth in Chapter 6.

3.6 Evolving Locomotion

In this section, the Basic EVC System is applied to a standard benchmark task for evolved virtual creatures: forward locomotion. These results were produced with the use of shaping, with creature fitness evaluated using a sequence of goals leading up to locomotion: the addition of joints, the addition of useful muscles, jumping, then horizontal motion.

The ultimate fitness for this task was defined by interleaving 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 (within [0, 1]), the combined fitness f is

$$f = \frac{\sigma(\lfloor \frac{s}{\sigma} \rfloor + \epsilon)}{s_{max}}. \tag{3.4}$$

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

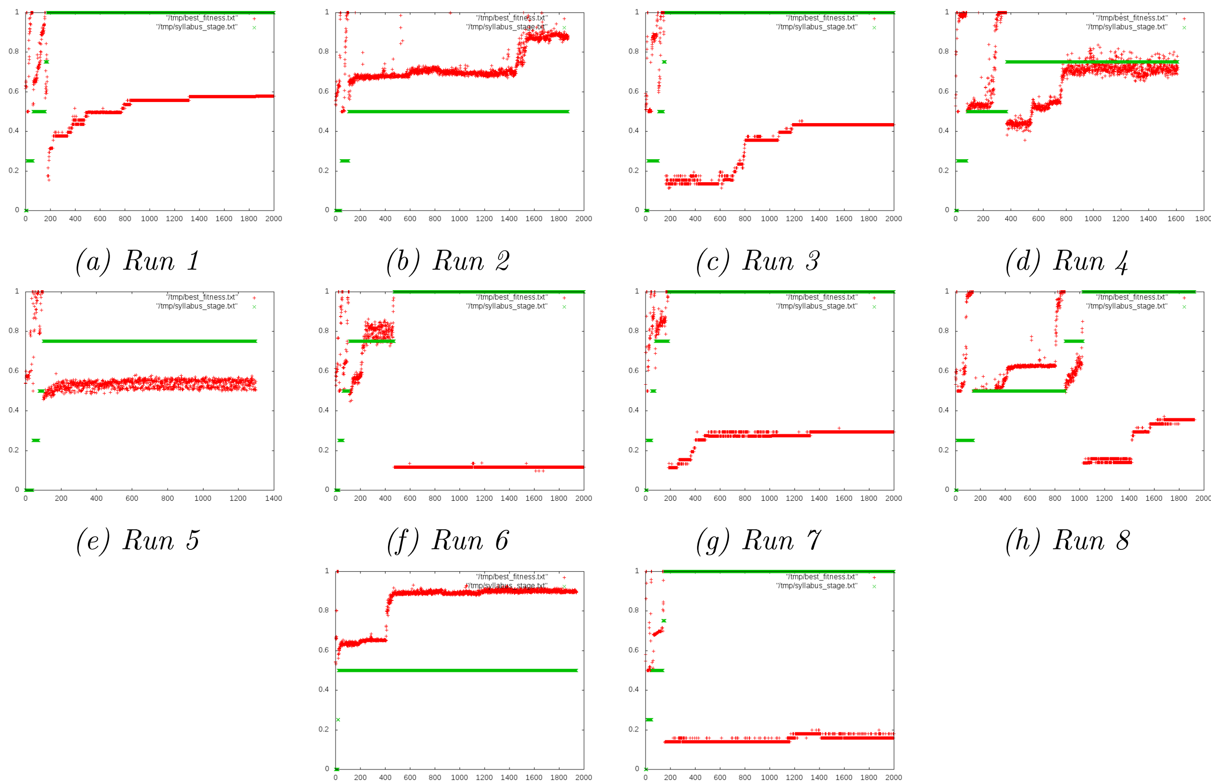

(a) Run 1 (b) Run 2 (c) Run 3 (d) Run 4

(e) Run 5 (f) Run 6 (g) Run 7 (h) Run 8

(i) Run 9 (j) Run 10

Figure 3.9: The fitness graphs (in red) of all 10 runs from which the locomotion result was selected. Within each graph, the horizontal axis measures generations of evolution, and the vertical axis indicates fitness (for the red marks). For the green marks, vertical position indicates the stage of fitness shaping for that generation (as described in Section 6.5). Seven of the 10 were successful, as shown in Figures 3.10.

In this experiment, 10 independent instances of the evolutionary algorithm were run, each with a unique random seed and a random starting population of single-segment creatures with empty brains. In the end, the best was chosen based on a number of criteria, including both numerical fitness and aesthetic concerns. Each instance had a population size of 200 and was allowed to evolve for up to 2000 generations, although earlier generations (as low as a few hundred in some cases) sometimes produced better results. The fitness graphs of all 10 runs are shown in Figure 3.9. Out of these 10, seven produced useful results. Although no diversity promoting mechanisms were used, the champions were highly diverse in both morphology and locomotion techniques. In fact, each of the seven runs resulted in a different method of locomotion, as illustrated in Figure 3.10.

(a) By oscillating its simple, two-segment body, this creature (Run 1, generation 2000) produces a highly effective form of locomotion.

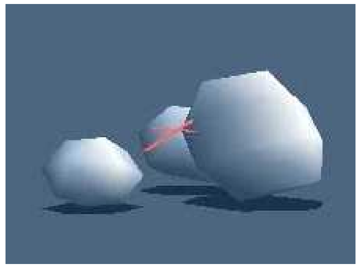

(b) With two small limbs in front alternately pulling its body forward, the creature of Run 3's 2000th generation has found a new way to propel itself.

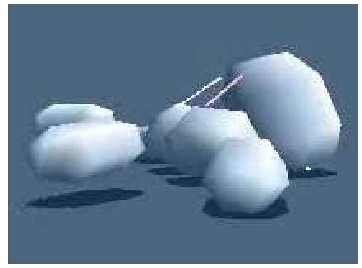

(c) Although slower than some other results, generation 300 of Run 5 produced a creature that travels sideways in a manner reminiscent of a maraca dancer.

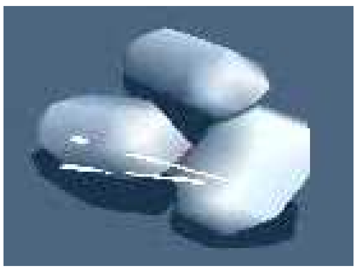

(d) This creature (Run 6, generation 2000) swings two limbs back to move forward in a swimming-like motion.

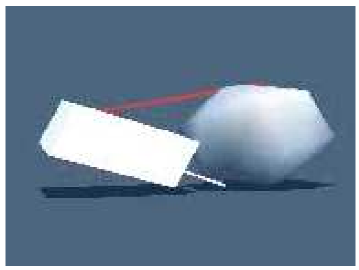

(e) This creature (generation 350 of Run 7) deliberately drags itself along with a single box-shaped forelimb.

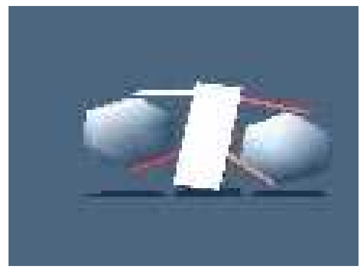

(f) This creature (Run 8, generation 1900) uses two limbs held off the ground for balance, as it achieves an extremely fast, stable, and visually appealing form of locomotion.

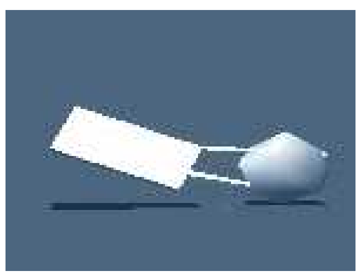

(g) The creature of Run 10, generation 2000 takes the unusual approach of kicking itself in the back to produce its forward locomotion.

Figure 3.10: Successful forward locomotion results produced by the Basic EVC System. These creatures illustrate the seven successful results of the 10 evolutionary runs used for this experiment. Note that each one employs a different method of locomotion, despite the lack of any diversity-promotion mechanism.

3.7 Conclusion

This chapter described the particular EVC system implementation which is the foundation of this dissertation's three primary contributions. In the next three chapters, the Basic EVC System is extended and evaluated in three ways: (1) constructing a creature that learns more complex tasks than seen before for EVCs, (2) evolving morphologies adapted to multiple tasks, and (3) demonstrating how muscle drives can take on part of the intelligence burden normally borne by the brain.

Chapter 4

Evolving Complex Behavior with ESP

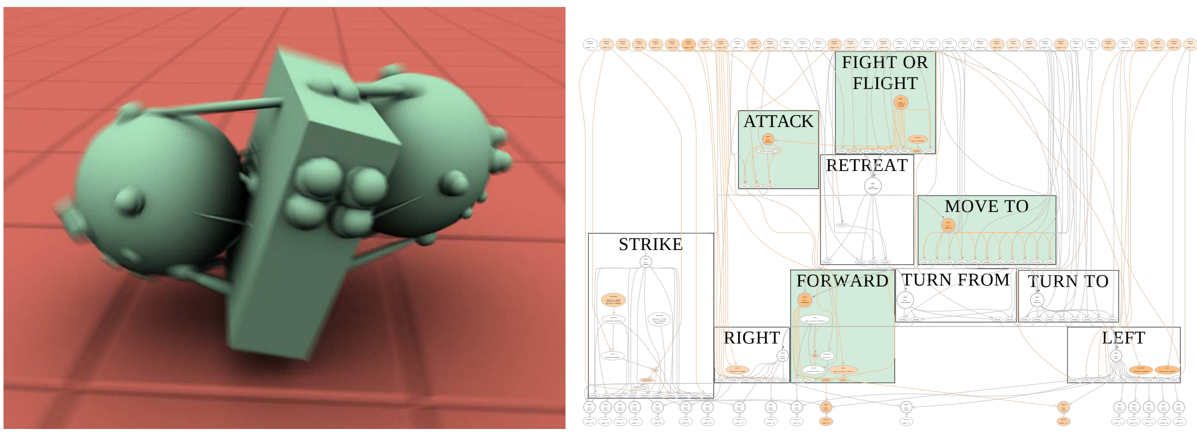

Figure 4.1: The body and brain of a creature evolved using the ESP method to learn a complex fight-or-flight behavior. This creature has achieved a level of behavioral complexity that is approximately double the previous state of the art for evolved virtual creatures. Previously, the most complex behavior in EVCs was the ability to move to a light source. This creature can move to a light source, strike it once it arrives, and switch to a flight behavior when appropriate, based on its perception of the environment.

At this point, with the underlying EVC system specified, it is possible to describe this dissertation's first major contribution: ESP, a mechanism to exceed the nearly two-decades-old ceiling in behavioral complexity for evolved virtual creatures [27].

4.1 Motivation

As discussed in Section 1.1, behavioral complexity for evolved virtual creatures has not increased for 19 years. This stagnation is surprising given that more complex behaviors would greatly improve creature content. The right kinds of complex behavior inspire in the viewer a sense of inner life and motivation in even the simplest creature forms (i.e., perceptual animacy). This effect makes complex behaviors a promising avenue for improving EVCs as content.

But how can more complex behaviors be created? *ESP* (Figure 4.1) answers this question by applying real-world instructional ideas that have a proven record of bringing otherwise impractical learning goals within reach.

Consider, for example, the task of hovering a helicopter. This is a challenging task even for humans to learn. It requires proper inputs in four or five dimensions of control, each of which typically affects the others: lift the collective to increase altitude, and the change in power requires a corresponding increase in anti-torque pedal application; use longitudinal cyclic input to begin forward motion, and complex aerodynamic effects require changes in collective input to maintain altitude.

This challenging task is usually learned by decomposing it into a sequence of simpler challenges. First the instructor might control all but the collective, so that the student can learn to control altitude in isolation; then the instructor might control all but the cyclic, so the student can focus only on controlling the helicopter's tilt; then, with those two mastered, they can be attempted together. Proceeding in this manner, the student can acquire smaller skills independently, and in the proper order, which eventually allows the skills to be combined to accomplish the full hovering behavior successfully. This type of hierarchical syllabus-based task decomposition is the core concept behind ESP.

Could a student accomplish the same goal without such guidance? For example, could a student learn martial arts simply by repeatedly entering competitions and trying to improve based on his or her score? While it is certainly possible for human students to learn a complicated topic independently, their development is typically faster and surer with an expert-designed syllabus. Simple individual skills—punch high, kick low, block left, block right, etc.—are learned one at a time, then eventually combined to accomplish more complex goals: When you perceive an attack from your right, block right; when your opponent's legs are vulnerable, kick low. Ultimately, this approach brings even the highest-level skills—such as participating in a competition—within reach. Thus, the syllabus acts as a sequence of waypoints through the space of possible solutions, decomposing the larger learning task into a succession of more manageable steps between the waypoints.

If such assistance is useful—or even required—for a learner as powerful as a human student, it should be even more useful for the algorithmic mechanical learners available to artificial life researchers today. It is also important that the human input required to create the syllabus is relatively abstract and human-relatable. In fact, such input should be simpler than the fitness functions currently required for evolving virtual creatures. In comparison to those highly technical specifications, the syllabus should require only the encoding of high-level learning concepts such as "before you learn to move to a target, first learn to move forward, left, and right". The method for doing that will be described next.

4.2 Method

ESP adds three new elements to the underlying EVC system: encapsulation, syllabus, and pandemonium. In this section, each of these components is described in detail. As the primary component of the ESP system, the syllabus is presented first.

4.2.1 Syllabus

In the ESP system, the *syllabus* consists of an ordered sequence of intermediate goals used to reach the ultimate, larger goal. This collection of 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.

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. Left turn and right turn are of a similarly manageable difficulty, and will be attempted next. Then, with the turns mastered and with an additional ability to develop photoreceptors, it is relatively straightforward to maintain orientation toward a light source. With the ability to face a light and the ability to move forward, navigating to that light is a similarly achievable goal, and so on. Proceeding in this manner, a knowledgeable human designer might produce the following sequence of subskills to be learned; each subskill is attainable with basic EVC methods, and 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)

This information may be conveniently summarized in a graph, encompassing subskills to be learned, dependency between subskills, learning order, and competition (Section 4.2.3), as seen in Figure 4.2.

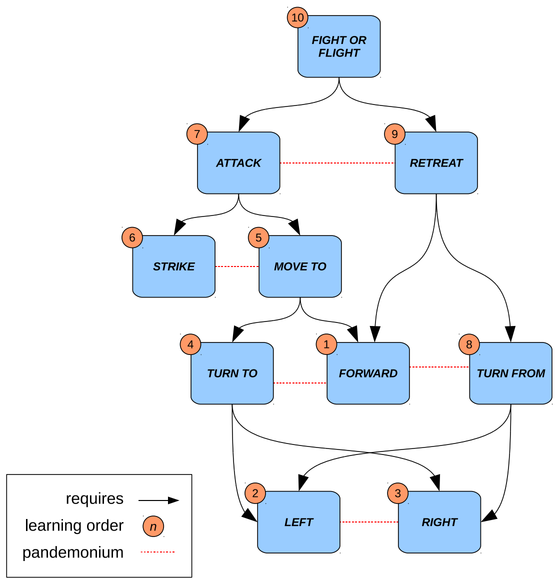

Figure 4.2: An example syllabus as a graph. Graph nodes represent individual subskills to be learned, directed edges indicate dependencies between subskills, and the numbering indicates a learning order that satisfies the dependency requirements. Pandemonium (i.e., competitive) relationships are indicated by dashed red lines.

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? Encapsulation is a mechanism that makes it possible.

4.2.2 Encapsulation

The second element of the ESP system is a mechanism to *encapsulate* previously learned skills. This element accomplishes two goals: It ensures that previously learned skills (and the body components on which they rely) are preserved, and it makes these skills easily accessible to future evolutionary development. Both of these goals are achieved through the automated encapsulation process illustrated in Figure 4.3.

Figure 4.3a depicts a brain evolved for forward locomotion, and Figure 4.3b shows the result of encapsulation. Note that, first, the nodes that compute the old skill have

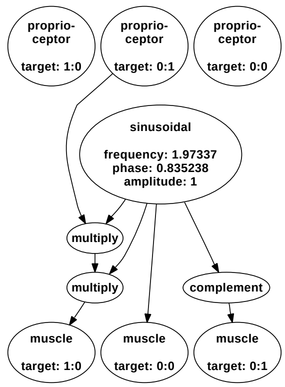

(a) Before encapsulation.

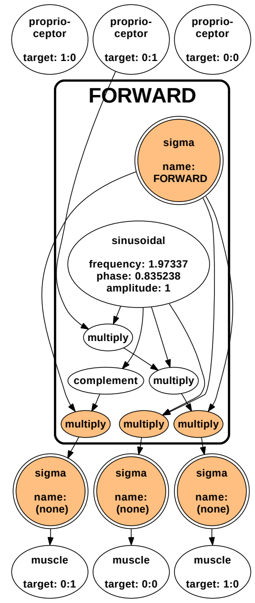

(b) After encapsulation (with new nodes shaded).

Figure 4.3: Encapsulation. The encapsulation of an evolved skill—in this case, forward locomotion—ensures that it will persist throughout future evolution, while also allowing it to be activated easily as a unit by future skills. Before encapsulation, in (a), the brain nodes that constitute the forward locomotion skill connect directly to muscles and sensors in the body, and may be changed by future evolution. In (b), the newly added nodes (shaded) implement the encapsulation of that skill. The multiply *nodes at the bottom throttle all signals leaving the skill, allowing its effect to be dialed up or down. The* sigma *node at the top acts as a single point of control for all of the throttling multiply nodes, allowing all outputs from the skill to be blocked or allowed out simultaneously. All nodes within the box labeled with the skill name* FORWARD *are protected from future evolutionary changes. (The* sigma *nodes at the bottom are not directly related to the encapsulation of the new skill, but rather are required so that future skills can share control over new muscles that have been recently added.)*

been preserved and frozen (meaning that future evolution cannot change them). Second, a new *multiply* node has been inserted into every output wire leaving the encapsulated skill. The internals of the skill will continue to function as before, always trying to perform their forward locomotion task, but now, a second signal sent to each new multiply node will modify those outgoing forward-locomotion control signals, scaling them by a number within [0,1]. Third, a single controlling node (called a *sigma node* for its function as a summation of zero or more inputs) is added, sending 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.

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 frozen, 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 acquired cumulatively, but a potential problem still remains: How will competing signals from the multiple sub-brains within a single creature be resolved? The pandemonium component will achieve this goal.

4.2.3 Pandemonium

Consider the following example based on the syllabus graph of Figure 4.2. A creature evolved through this syllabus will ultimately have a part of its brain devoted to left and a part to right turns. But it is 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 competitive system like this, sub-brains within the creature can compete for overall control, with little risk of sabotaging the rest of the brain. In Figure 4.2, 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.

4.3 Evolving a Fight or Flight Behavior

The first contribution to be described in this dissertation is an application of the ESP method, using the syllabus of Figure 4.2, to evolve a virtual creature through a sequence of ten learning tasks. The first five of these tasks approximately match the previously demonstrated behavioral-complexity limit for EVCs, and the second five approximately double it. These results are best viewed in the accompanying video at `http://youtu.be/dRLNnJlT8rY`.

4.3.1 FORWARD LOCOMOTION

In the first step, a FORWARD LOCOMOTION result from the basic EVC system was chosen, and its control abilities encapsulated. The creature selected for continued evolution (Figure 4.4) was produced at generation 1900 of run number eight of the Basic EVC System experiment described in Section 3.6. Receiving a fitness score of 0.355, it was chosen from among the successful results based on qualities beyond the locomotive speed required by the fitness function: a smooth, coordinated style of motion; stability and reliability; and a simple brain.

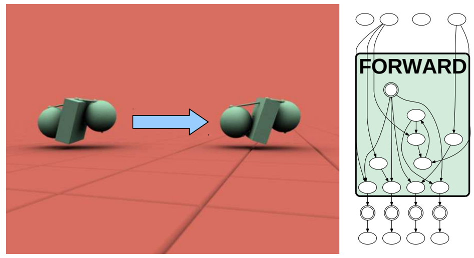

Figure 4.4: The chosen FORWARD LOCOMOTION *result (run eight, generation 1900) after encapsulation. This creature was selected not only for its fitness score, but also for its aesthetically pleasing motion style, reliable locomotion, and simple brain.*

At this point in its progression through the syllabus, 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.

4.3.2 LEFT TURN

With the LOCOMOTION skill encapsulated, a new run of evolution begins, this time with the fitness function rewarding the ability to rotate counterclockwise while largely maintaining position. Specifically, core fitness was determined by degrees per second of rotation (out of a maximum of 90), with a fitness of zero assigned if the creature's center of gravity moved more than its approximate diameter during a single evaluation (which lasted for one second, after a one-second spin-up period).

To prepare for each LEFT TURN run, the selected locomotion result (with LOCOMOTION encapsulated and thereby excluded from further evolution) was duplicated 200 times. This provided each run with an initial population whose brains appeared (to the current round of evolution) empty, just as they were before locomotion was evolved. The same technique is used for all steps in the syllabus.

A left-turning creature was produced using five such independent runs, each of which lasted less than 300 generations (Figure 4.5). Such reduced evolutionary resources (compared to locomotion) were sufficient since this skill is easier (aside from the risk of physics cheating,

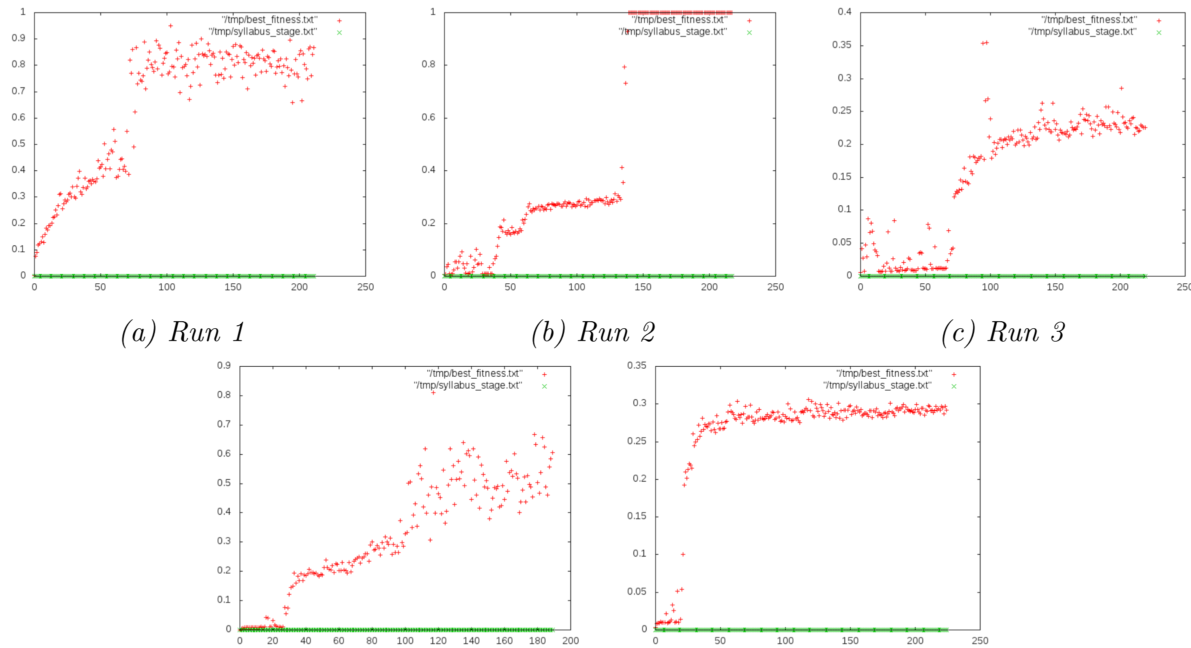

(a) *Run 1* (b) *Run 2* (c) *Run 3*

(d) *Run 4* (e) *Run 5*

Figure 4.5: The fitness graphs of all five runs from which the LEFT TURN *result was selected. All runs were successful, making it possible to select a creature for further evolution using visual criteria.*

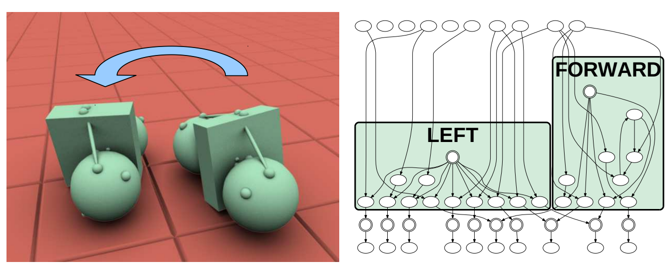

Figure 4.6: The chosen LEFT TURN *result (generation 319, run two). In addition to the fact that it achieved the highest fitness—the maximum possible for this task—this creature was selected for fast, fine turns; reliable skill transitions; and the visible contrast between its turn and locomotion behaviors.*

discussed in Section 4.5). Each of these five starts produced a usable result. The addition of new muscles was allowed during this stage of evolution.

The chosen skill is shown (after encapsulation) in Figure 4.6. It was selected at generation 319 from run number two for three reasons beyond having the highest fitness score (clamped at 1.0): (1) This creature displayed a fast skittering turn, able to stop at a finely selected angle, which improves its ability to achieve a given orientation.(2) It had the ability to transition reliably between skills. (3) It demonstrated a visually pleasing contrast between the left turn and the previously acquired forward locomotion skill.

4.3.3 RIGHT TURN

With the first two skills encapsulated, a clockwise turn is evolved in the same way as the counterclockwise turn (but with the fitness score negated) and the result is encapsulated (Figure 4.8). The fitness scores from all five runs are illustrated in Figure 4.7. The individual at the end of the highest-scoring run (generation 450 from run number five) was examined and found to perform the right turn well. This creature was used for continued evolution, and the results from the other runs were not examined in detail.

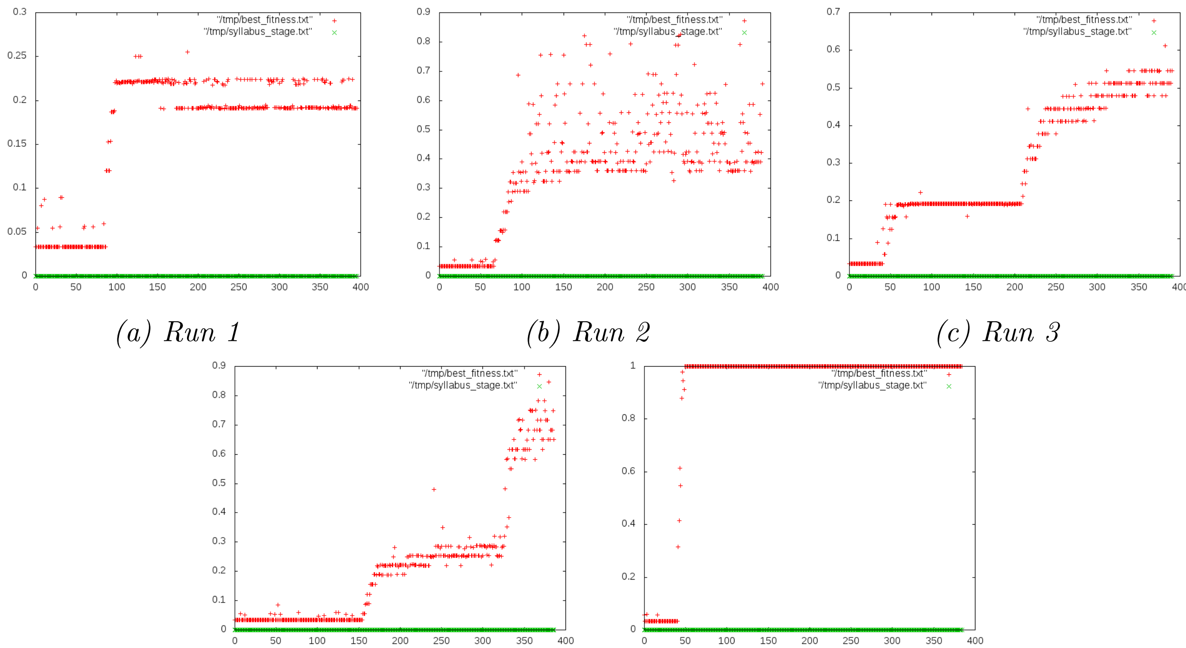

(a) Run 1 *(b) Run 2* *(c) Run 3*

(d) Run 4 *(e) Run 5*

Figure 4.7: The fitness graphs of all five runs for the RIGHT TURN *skill.*

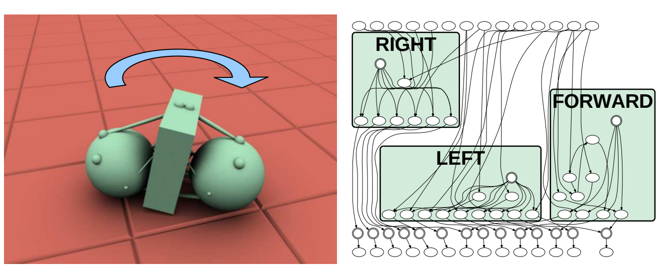

Figure 4.8: The selected RIGHT TURN *result. This creature (generation 450 from run five) achieved the highest fitness score in all five runs for this skill.*

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. Next, one such skill, TURN TO LIGHT, will be developed.

4.3.4 TURN TO LIGHT

For this skill, the creature is tested on its ability to orient its direction of locomotion toward a target perceived as a point light source. It is allowed to evolve photoreceptors (described in Section 3.4), and use the previously encapsulated LEFT TURN and RIGHT TURN skills, which are placed in a pandemonium relationship, as indicated in the syllabus (Figure 4.2). To encourage a general solution, the fitness evaluation is an average over four tests, each with a fixed light source at a different heading from the creature (northeast, northwest, southeast, or southwest). For each direction, fitness is the average throughout the evaluation time of the number of degrees off heading from the target (normalized for the maximum of 180). For this task, physics cheating was not a factor, and a relatively small amount of evolutionary resources were sufficient. Successful solutions were evolved in each of the five runs, using a population size of 35 and less than 300 generations, with scores stabilizing before approximately 100 generations.

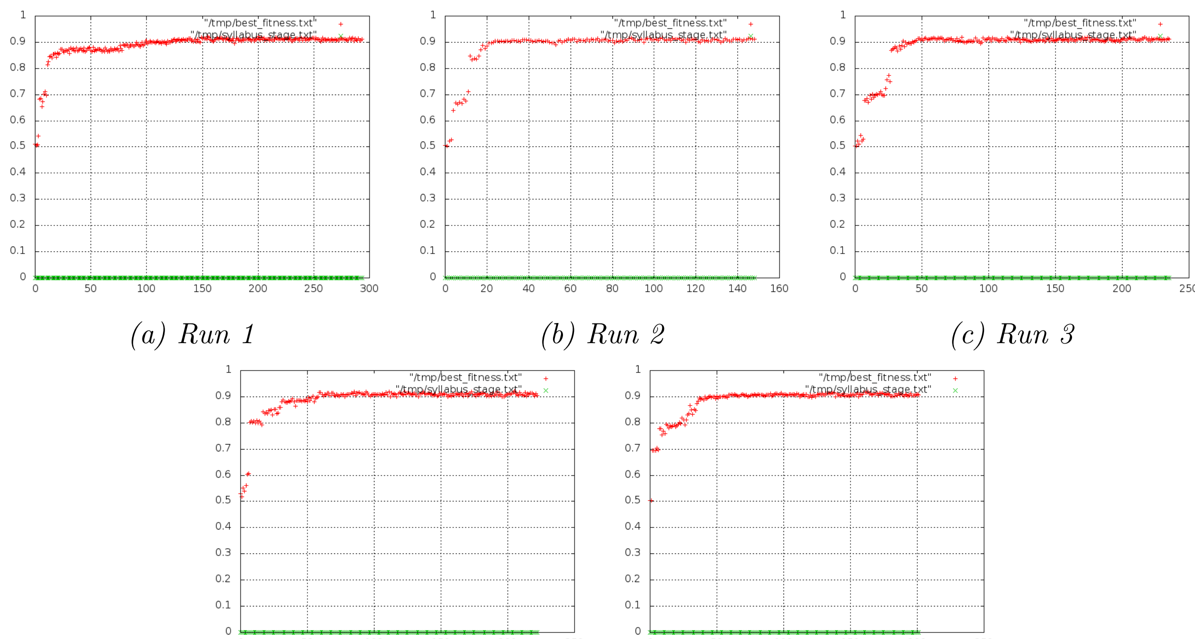

(a) Run 1 *(b) Run 2* *(c) Run 3*

(d) Run 4 *(e) Run 5*

Figure 4.9: The fitness graphs of all five runs for the TURN TO LIGHT *skill. Again all runs were successful, making it possible to select the winner based on aesthetic concerns.*

The full set of fitness graphs is depicted in Figure 4.9. Because consistently high scores were achieved in all runs, the creature from generation 349 of run one was selected based on aesthetic concerns: crisp stops between turns, a lack of blind spots, a simple brain, and relatively few photoreceptors. Figure 4.10 shows the completed and encapsulated result, which is able to consistently aim its locomotion direction at a user-controlled target.

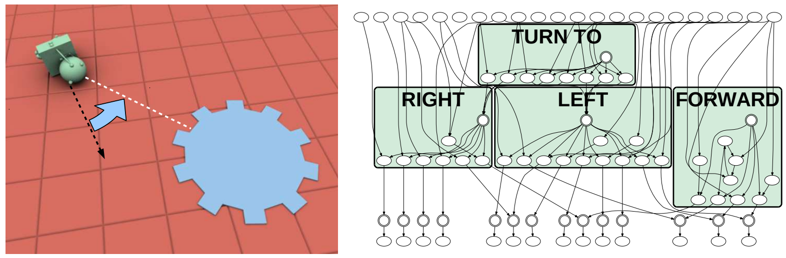

Figure 4.10: The chosen TURN TO LIGHT *result. The* TURN TO LIGHT *skill keeps the locomotion direction (black dashed arrow) oriented toward a target (depicted here as a large disc, but perceived by the creature as a single omnidirectional light source at the disc's center). The winning creature for this task (generation 349, run one) is depicted. It was chosen for multiple aesthetic reasons beyond its high fitness score: stops between turns, a lack of blind spots, a simple brain, and a relatively small number of added photoreceptors.*

4.3.5 MOVE TO LIGHT

With TURN TO LIGHT and FORWARD LOCOMOTION available, and with the evolution of more photoreceptors allowed, the creature was then evaluated on its ability to navigate to a light source. As with TURN TO LIGHT, fitness is averaged over multiple evaluations (in this case five), again with a fixed light source at a different relative angle each time. In this case, the light positions were the same as for the previous skill, but with an additional evaluation for a light directly behind the creature to reduce the occurrence of blind spots. Fitness for a single evaluation was defined as the average distance from the target during simulation.

As in the previous task, only a relatively small amount of evolutionary resources were required. Five runs with a population size of 35 produced successful results in all runs, with all scores leveling off by approximately 100 generations (Figure 4.11. Despite the almost identical fitness, some minor variations were observed across runs in attributes not encoded in the fitness function. For example, some results focused more on choosing an accurate locomotion direction first then moving straight to the target, while others achieved a similar score by adjusting course more while approaching the target, giving the impression of a somewhat less confident and capable creature (e.g., run 4, generation 286, with fitness 0.762; illustrated in Figure 4.12). The creature selected for continued evolution from these runs (run 5, generation 308, with fitness 0.765) was chosen for being most consistent in using forward locomotion in long uninterrupted stretches. This behavior gave the creature the strongest appearance of deliberate intentionality.

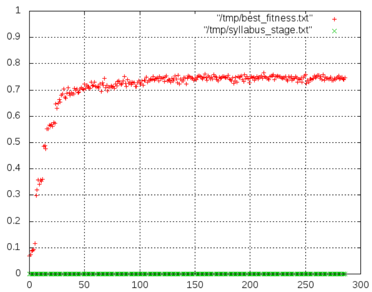

(a) Run 1

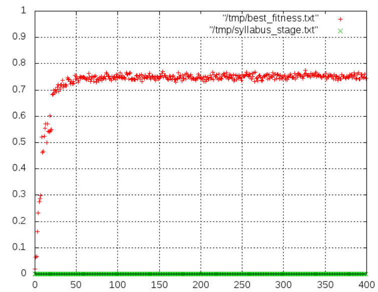

(b) Run 2

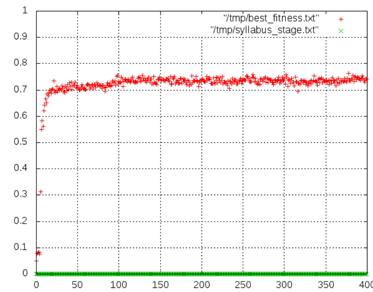

(c) Run 3

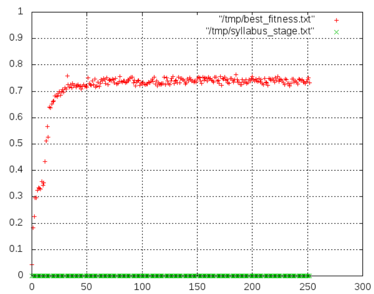

(d) Run 4

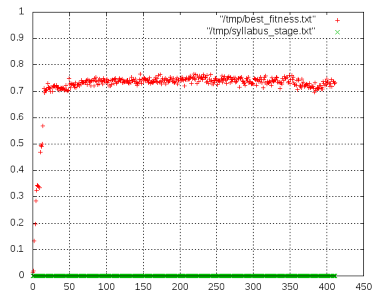

(e) Run 5

Figure 4.11: The fitness graphs of all five runs for the MOVE TO LIGHT *skill. All runs were successful, and the winner was chosen based on aesthetic concerns.*

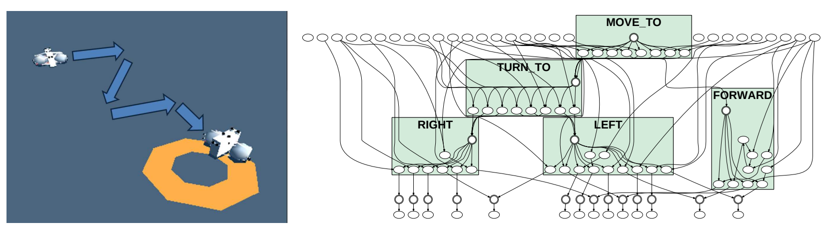

Figure 4.12: A high-scoring, but less appealing MOVE TO LIGHT *result. Despite meeting the requirements of the fitness function approximately as well as the selected creature, this creature's style of motion gave it a less confident appearance.*

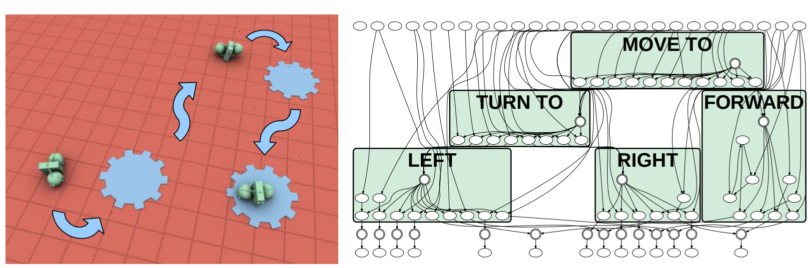

Figure 4.13: The chosen MOVE TO LIGHT *result. The creature shown in this image (generation 308, run five) has acquired the* MOVE TO LIGHT *skill, allowing it to follow a target along a curving path, catching the target when it finally stops. This creature was chosen for its appearance of deliberate intentionality.*

The creatures produced for this task demonstrated behavioral complexity that approximately matched the state of the art. The selected result is illustrated in Figure 4.13.

4.3.6 STRIKE

In anticipation of the upcoming ATTACK task (Figure 4.2), the creature must first learn to deliver a strike to the ground underneath it. Fitness is primarily computed as total such strike force in each one-second interval (averaged across five intervals in sequence), with small factors added to reward beginning and ending contacts as well as vertical center-of-gravity movement during an interval (to provide a degree of shaping).

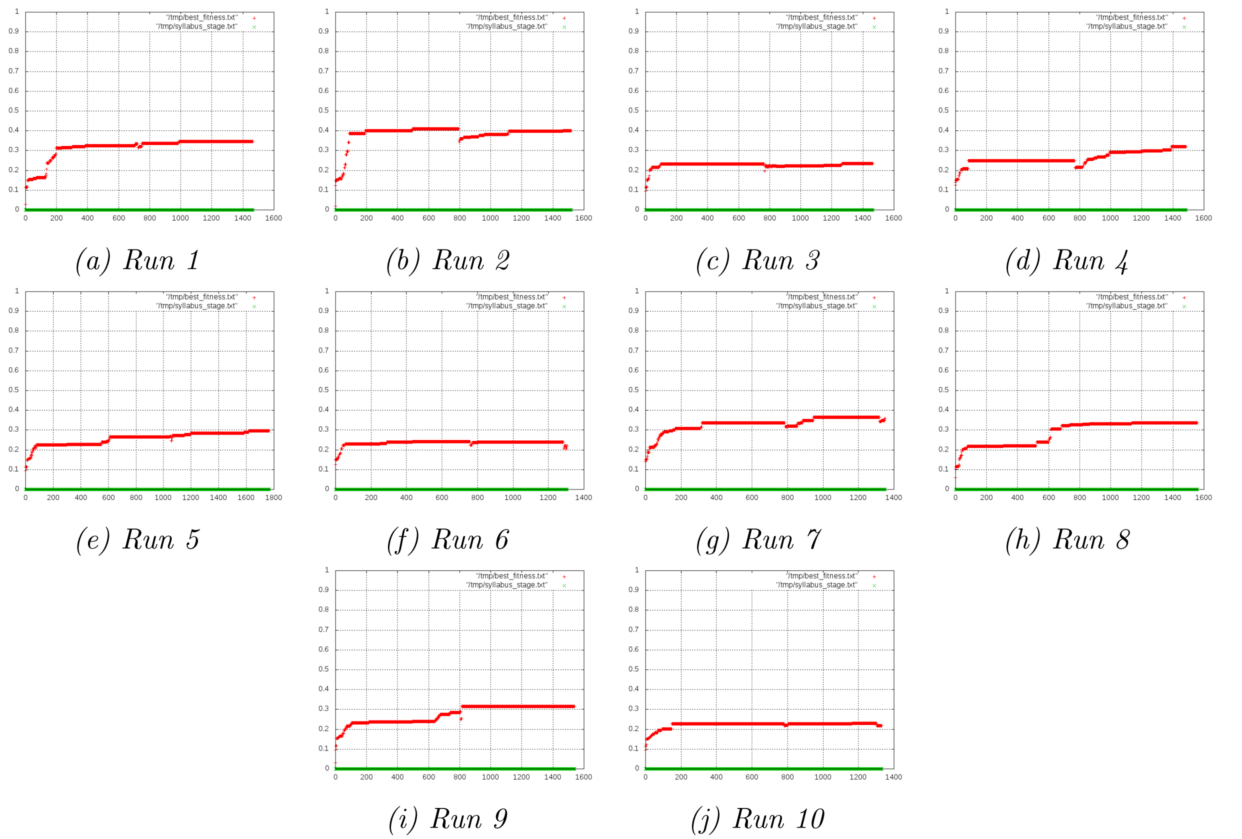

(a) Run 1 (b) Run 2 (c) Run 3 (d) Run 4

(e) Run 5 (f) Run 6 (g) Run 7 (h) Run 8

(i) Run 9 (j) Run 10

Figure 4.14: The fitness graphs of all 10 runs for the STRIKE *skill. Five of the runs produced useful results. They varied significantly in their approach, and the winner was chosen based on aesthetic concerns.*

To facilitate the evolution of this new low-level skill, evolution of new muscles is allowed. As might be expected, this task requires greater evolutionary resources, both because of the muscle evolution and because it controls the body directly rather than switching between a few pre-existing skills. Using a population of 100, five of 10 runs produced useful results. Some of them emerged after as few as 100 generations, but some improvements still appeared after more than 1000 generations (Figure 4.14).

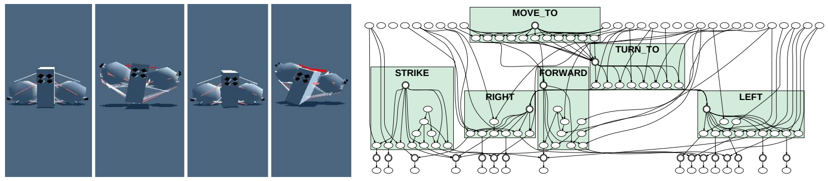

Figure 4.15: A high-scoring, but less appealing STRIKE *result. Despite a high fitness score, this creature's technique consisting of many small jumps similar to locomotion produced a less rewarding visual effect.*

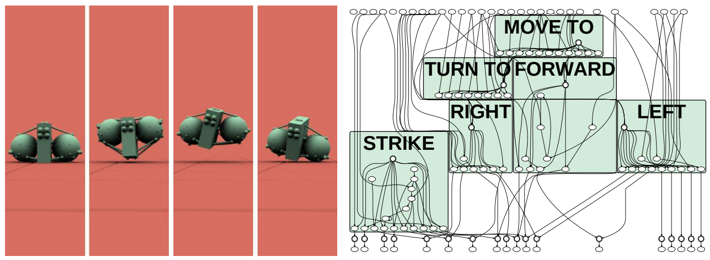

Figure 4.16: This creature's STRIKE *solution (generation 1313 of run seven) employs a vertical jump, and was chosen for its high score, deliberate appearance, and visual contrast with other behaviors.*

As often occurs, even results with similar fitness scores varied in how useful they were because important aesthetic factors that are not encoded in the fitness function. For this skill, for example, it was possible to score well with many small strikes rather than few larger ones. This approach looks weak and is visually less distinguishable from locomotion, producing a less rewarding overall effect when switching between the two behaviors (as in run 2, generation 781, with fitness 0.409; Figure 4.15). In contrast, the selected winner for this skill (run seven, generation 1313, with fitness 0.364; Figure 4.16) was chosen not only for its high score and ground impact, but also for its moderately high and less frequent jumps. These produced a visually appealing look of deliberate action, as well as a clear visual contrast with other behaviors.

4.3.7 ATTACK

Having learned MOVE TO LIGHT and STRIKE, it is now possible to produce an ability more complex than simply moving to a target. By first moving to the target, *then* striking, this creature takes another small step toward the behavioral complexity of compelling creature content from the real world. For this task, fitness is an average across the four cardinal 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).

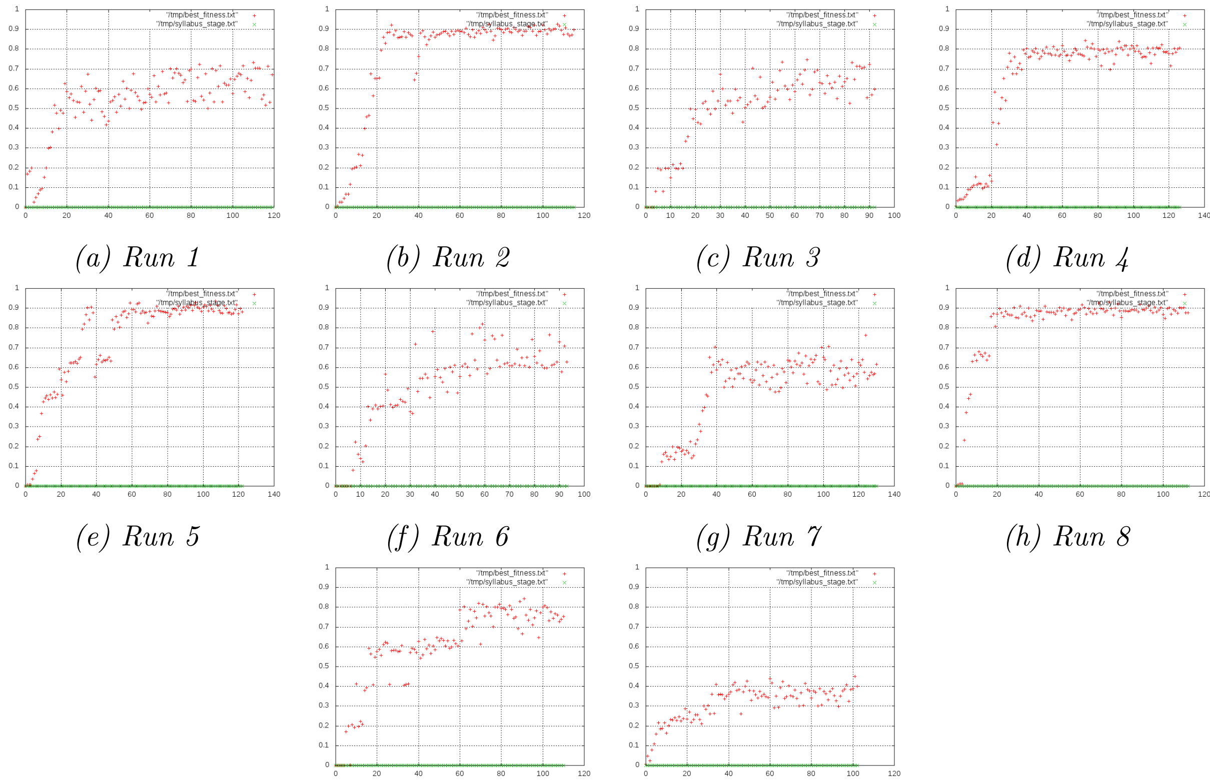

(a) Run 1 (b) Run 2 (c) Run 3 (d) Run 4

(e) Run 5 (f) Run 6 (g) Run 7 (h) Run 8

(i) Run 9 (j) Run 10

Figure 4.17: The fitness graphs of all 10 runs for the ATTACK *skill. Three of these runs produced good results, making it possible to select a winner based on aesthetic concerns.*

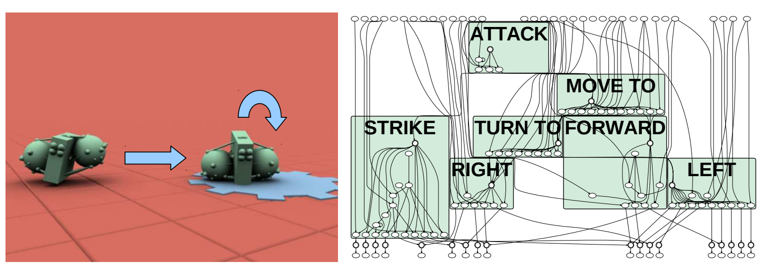

Figure 4.18: In the newly added ATTACK*, the creature navigates to the target, then strikes it. The winning creature shown here (run two, generation 24) was chosen for having few new eyes, accurate strikes, and clean switches between skills.*

This task builds upon existing skills only, with no direct control over the body, and it thus required relatively few evolutionary resources. With a population size of 50, three of the 10 runs produced results of sufficient quality, and the highest scores were achieved within 50 generations (Figure 4.17). While this skill's chosen creature (generation 24 of run two, with fitness 0.86; Figure 4.18) did not have the highest numerical score, it stood out for the following reasons: few new eyes added (and none on the top of the head), very consistent and accurate initial strikes, and very clean switching between skills.

4.3.8 TURN FROM LIGHT

In preparation for the upcoming RETREAT skill, the creature must learn to turn away from a light source. Although similar to TURN TO LIGHT, this task also requires a fitness term to discourage an initial turn in the wrong direction, in order to achieve reasonable results for targets near the creature's front. Also, significantly more evaluation directions (13) were used, particularly near the front, to produce reasonable reactions in those cases. (Evaluations were performed at directions every 60 degrees from the front, as well as at plus and minus 1, 2, 5, and 10 degrees.) When turning *to* a target near the front, choosing the wrong direction would be detrimental to fitness and this response would be eliminated quickly in the evolution. When turning away, however, the fitness cost of an initial wrong-direction turn with a near-front target is low (because a turn in the wrong direction quickly becomes a turn in the correct direction), so such behaviors are more likely to persist. Unfortunately, this kind of incorrect reaction with a target near the front looks particularly unnatural, and therefore it is worth spending a bit more evolutionary resources to eliminate it.

Other than those increased testing requirements, however, this task is not very demanding. All five runs (Figure 4.19) produced usable results in less than 60 generations, using a population size of 50. From these, the creature of generation 46 from run two (Figure 4.20) was selected not only for having the highest score (0.874), but also for aesthetic reasons. Specifically, the lack of eyes on top of the head made this creature more visually appealing, and its technique of remaining still after turning away from the target gave it a greater appearance of purposeful action.

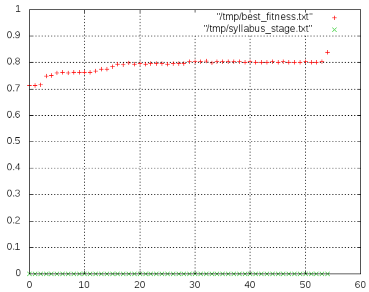

(a) Run 1

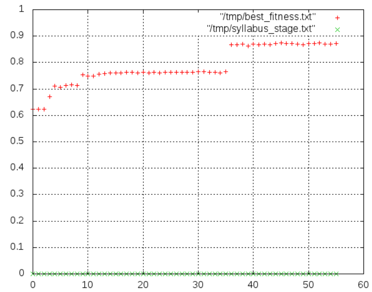

(b) Run 2

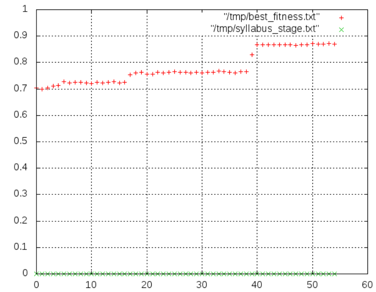

(c) Run 3

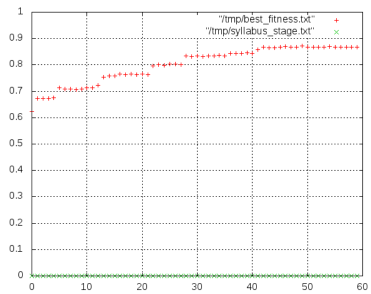

(d) Run 4

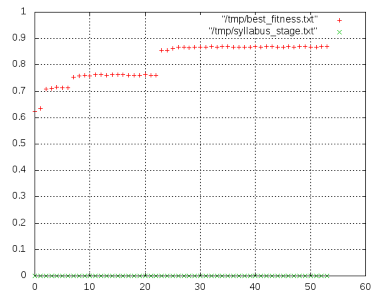

(e) Run 5

Figure 4.19: The fitness graphs of all five runs for the TURN FROM LIGHT *skill. All five runs produced usable results, allowing the winner to be selected based on aesthetic concerns.*

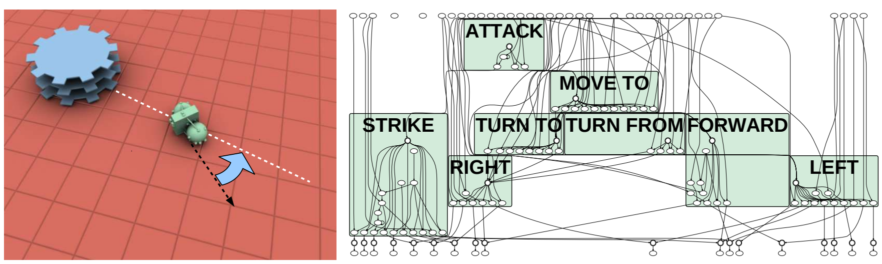

Figure 4.20: The selected TURN FROM LIGHT *result. The* TURN FROM LIGHT *behavior keeps the locomotion direction (black dashed arrow) oriented away from the target. The target (depicted as as a stack of three spinning discs) is perceived by the creature as point light sources at the center of each disc. Later, this will provide the opportunity to distinguish between the two target types based on both vertical placement of lights and overall light intensity. The winning creature for this skill (run two, generation 46) is depicted in this figure. In addition to having the highest score, it had an aesthetically pleasing eye placement and apparent intentionality.*

4.3.9 RETREAT

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.

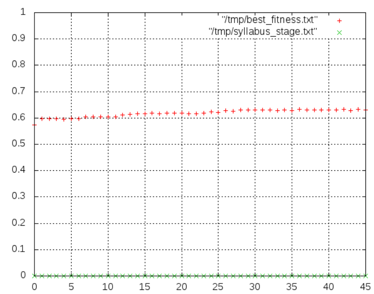

(a) *Run 1*

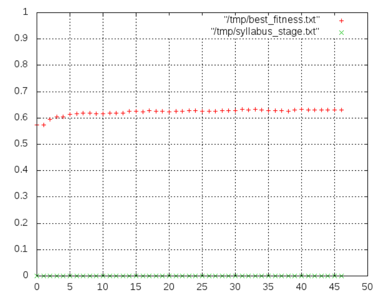

(b) *Run 2*

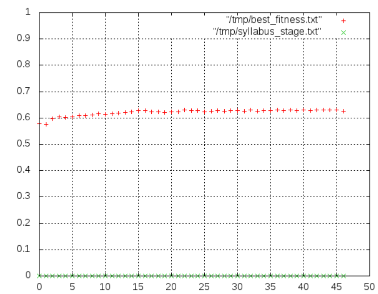

(c) *Run 3*

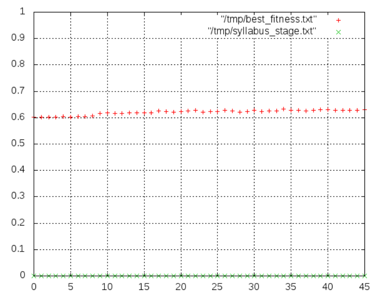

(d) *Run 4*

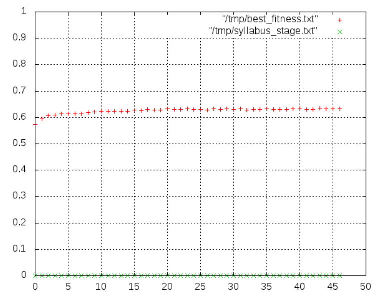

(e) *Run 5*

Figure 4.21: The fitness graphs of all five runs for the RETREAT *skill. All five runs produced usable results, allowing the winner to be selected based on aesthetic concerns.*

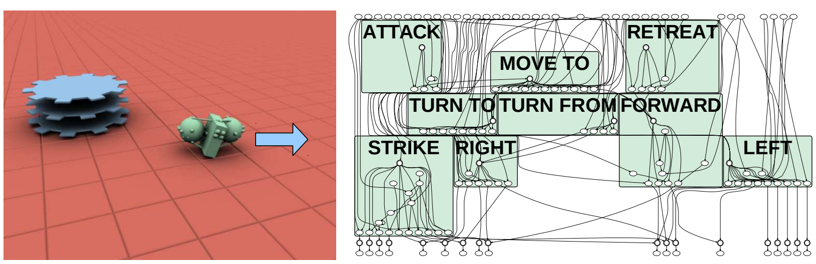

Figure 4.22: The winning creature for the RETREAT *skill (run five, generation 41). It was chosen for having the highest score and few added eyes.*

As for the previous skill, a population size of 50 was used along with five runs. All starts quickly produced high-quality results, with fitness scores for all runs (Figure 4.21) leveling off at their approximate best within a mere 15 generations (as observed within the 40 or more generations allowed for each run). From these runs, the winning creature (generation 41, run five, fitness 0.635; Figure 4.22) was selected for having the highest score along with a minimum of newly added eyes.

With this skill complete, the necessary components are in place for the final top-level skill of the syllabus.

4.3.10 FIGHT OR FLIGHT

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, or a dangerous target (a spinning vertical stack of three such discs), which will destroy the creature if it touches it. These two target types may be distinguished in the creature's perception based on the differences in light elevation and the overall intensity of the light.

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 from such a pair of evaluations 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 goal, 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 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, with f^+ the average *maximum* score across all n test directions, and f^- the average *minimum* score across all n test directions, the final overall fitness f is computed as

$$f = \frac{f^+ + 2n \cdot f^-}{2n+1}. \tag{4.1}$$

Without these additional motivations, solutions emerge that choose a single (higher-scoring) hard-coded reaction to be used for each light position—regardless of target type. The fitness function of equation 4.1 encourages the leap to the increased scores available by discerning between the two types of target.

As with a number of previous examples, five runs and a population size of 50 were used (Figure 4.23). However, unlike the typical results for internal-only tasks, this skill required significantly more evolutionary computation. Although all five runs eventually achieved similarly high scores, more generations were required—in one case, almost 200.

Figure 4.24 illustrates a highly successful, visually appealing and alive-looking result for this task (generation 163 of run two, fitness: 0.368), 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.

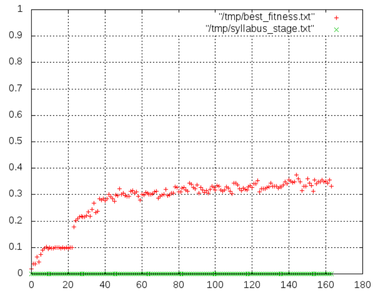

(a) *Run 1*

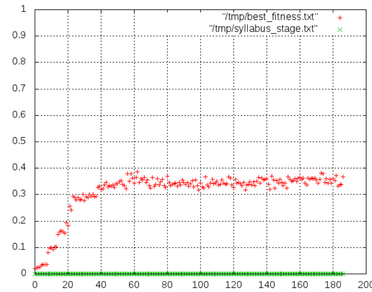

(b) *Run 2*

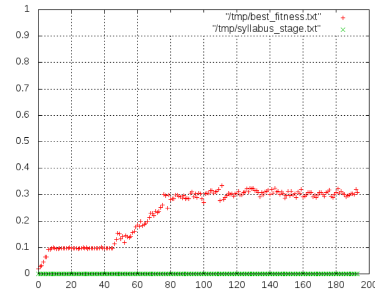

(c) *Run 3*

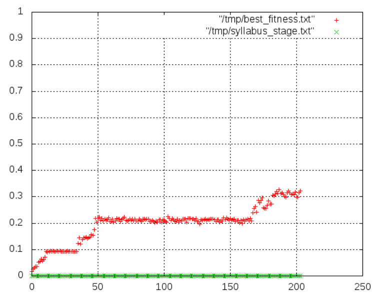

(d) *Run 4*

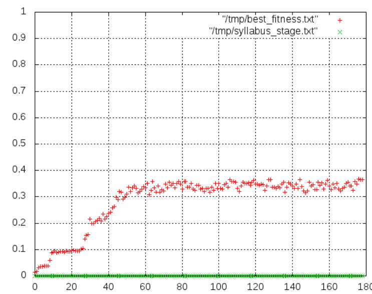

(e) *Run 5*

Figure 4.23: The fitness graphs of all five runs for the FIGHT OR FLIGHT *skill. All five runs achieved similarly high scores.*

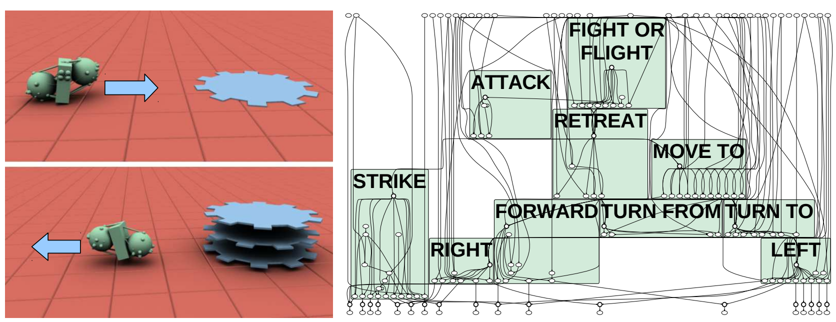

Figure 4.24: The winning creature (run two, generation 163) with the FIGHT OR FLIGHT *ability. This creature completes the full progression through the syllabus, resulting in behavior two levels more complex than the prior state of the art.*

4.4 Diversity of Solutions

The primary results described in Section 4.3 demonstrate the level of behavioral complexity that ESP makes possible. With that result established, it is informative to evaluate the diversity of complex-behavior creatures this system can produce. Throughout Section 4.3, a variety of solutions were generated for each skill (some of them very different from each other), and one selected for further development. An interesting question is, what would have happened if different choices had been made? In particular, what if a creature with a very different morphology would have been selected from the very beginning?

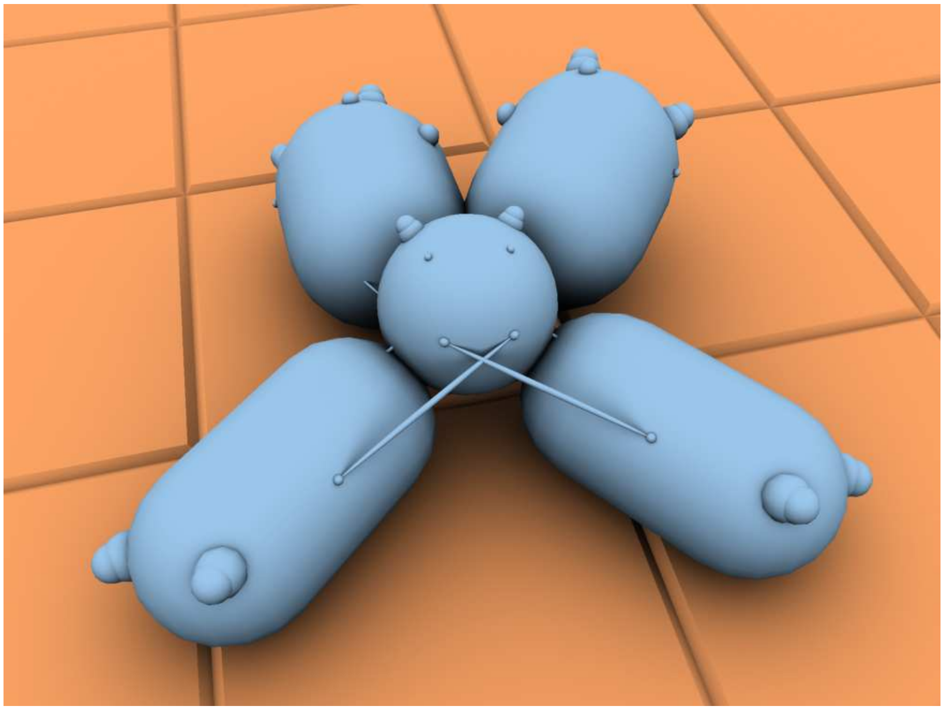

Figure 4.25: A creature with significant differences in body plan, physical techniques, and character from the creature evolved in Section 4.3. This creature (despite being produced using the same evolutionary technique and same selection criteria) has a different number of limbs, different types of body segments, and a diverse style of action. Despite these differences, this second creature nonetheless succeeded in all skills attempted, including the ATTACK *behavior—a level of complexity beyond the previous limit for EVCs. This result indicates that it is possible to produce significant variation using ESP, which makes it a good tool for creating virtual content.*

To answer this question, the creature in Figure 4.25 was chosen for study in this section. It originates from a different evolutionary run on the locomotion task, and has very different physical characteristics from the first selected creature (Figure 4.4). This second creature was further evolved in the same syllabus as the first creature. It is encouraging to see that, despite significant differences in morphology and locomotion style, the second creature was able to succeed in all of the skills which were attempted, including the ATTACK behavior—a level of complexity beyond the previous limit for EVCs. For video of this creature, see `http://youtu.be/dRLNnJlT8rY`.

Whereas the first creature has two short limbs that do not typically touch the ground, the second creature has four elongated limbs that do make contact with the ground. Whereas the first creature uses the momentum of its limbs to effect changes to the root segment's motion, the second creature uses its limbs to push against the ground for locomotion, turning, and strikes. Whereas the first creature's best strike is produced by jumping through upward swings of its limbs, the second creature's useful strikes are produced by direct impacts by its limbs on the ground. In fact, the second creature developed a variety of these entertaining and effective strikes, from whole-body momentum delivered through a single limb's impact, to a flurry of alternating blows from multiple limbs.

One way in which the second creature fares worse is in turn precision. Whereas the first creature's short and fast turn-movement cycles allow it to control its orientation to a fine degree, the second creature's turns are produced by a much slower and angularly larger unit of movement, offering fewer choices of orientation. While such imprecision may not be immediately apparent in isolated turning tasks, when those turns are employed to aim the creature for locomotion toward a target, it becomes a disadvantage. It often leads to multiple false starts and en-route adjustments to which the original creature was far less prone. With awareness of this potential pitfall, fine turning precision might be a useful element of fitness in the evolution of future creatures for similar tasks.

It is also worth noting that each of these creatures has potential value as content due to their contrasting styles of movement. The first creature's rapid and precise motions convey a very different character than the second creature, with its slow, loping, and perhaps even clumsy ways. ESP is able to create such diversity through natural variation in its evolutionary search, which makes it a good candidate for creating content automatically or various virtual world applications.

4.5 Discussion

As is common in evolutionary computation in a complex environment, evolution attempts to "cheat", often leading to a kind of arms race between evolved solutions and carefully defined fitness functions. In particular with the ESP system, cheating was a significant factor in the evolution of turns: Creatures tended to evolve that exploited small errors in physical simulation. Through the proper addition and application of muscles, it was relatively easy for creatures to apply forces to themselves which (presumably due to small physics errors) produced a physically inaccurate asymmetrical effect—a net turning force that would not have existed in reality. This exploit was rapidly and consistently discovered by evolution, preventing the physically realistic solutions from being discovered. This obstacle was eventually overcome by adding a sequence of evaluations in which the creature was required to turn in a reasonable way for multiple combinations of active and inactive brain and muscles. In addition, the full fitness scoring for a single individual required five repetitions of that evaluation sequence, with an overall score of zero if any of the repetitions indicated cheating. (This form of cheating produced such high fitness scores that it was considered a risk to allow even intermittent cheaters to survive, since even a single successful cheat might be enough to quickly take over the population.) Although this solution was found, it was only produced through many failed experiments, and resulted in a fitness function that is significantly more complex and costly to evaluate than would otherwise have been required.

It is also important to note that, while there are particular challenges in applying ESP-style task decomposition to EVCs, it has been used in multiple related fields for many years. As was outlined in Chapter 1, Selfridge's pandemonium, Minsky's society of mind [31], and Brooks' subsumption architecture [6] are prominent examples of such use in artificial intelligence and robotics. In reinforcement learning and evolutionary computation, layered learning and hierarchical task decomposition [49, 51, 11] explore similar concepts. With EVCs, however, no previous system has demonstrated how such an approach can be used to increase behavioral complexity beyond existing limits.

Note that in order to demonstrate how the long-standing ceiling on behavioral complexity can be exceeded, only a limited case of morphological evolution was considered in this chapter. Specifically, changes to body segments and joints (essentially the creature's skeleton) were prevented after the first skill was complete (although the brain, muscles, and eyes continued to evolve). While this constraint may have been an appropriate choice for the first step, it is a serious limitation. The next chapter describes the extended ESP system, in which this constraint is removed.

4.6 Conclusion

This chapter described the first of this dissertation's major contributions: ESP, a mechanism for exceeding the behavioral complexity ceiling in EVCs that had existed for almost two decades. In addition, useful techniques for overcoming expected but unpredictable cheats by evolution were developed. Third, an evaluation of a second creature with a number of diverse characteristics established that ESP can create a variety of useful content, which makes it promising as a technique for virtual world applications.

The following chapter presents this dissertation's second major contribution: an extended version of the ESP system that allows full morphological adaptation to multiple tasks, while keeping the original system's ability to decompose a complex learning task into a hierarchical sequence of simpler goals.

Chapter 5

Evolving Complex Morphology with Extended ESP

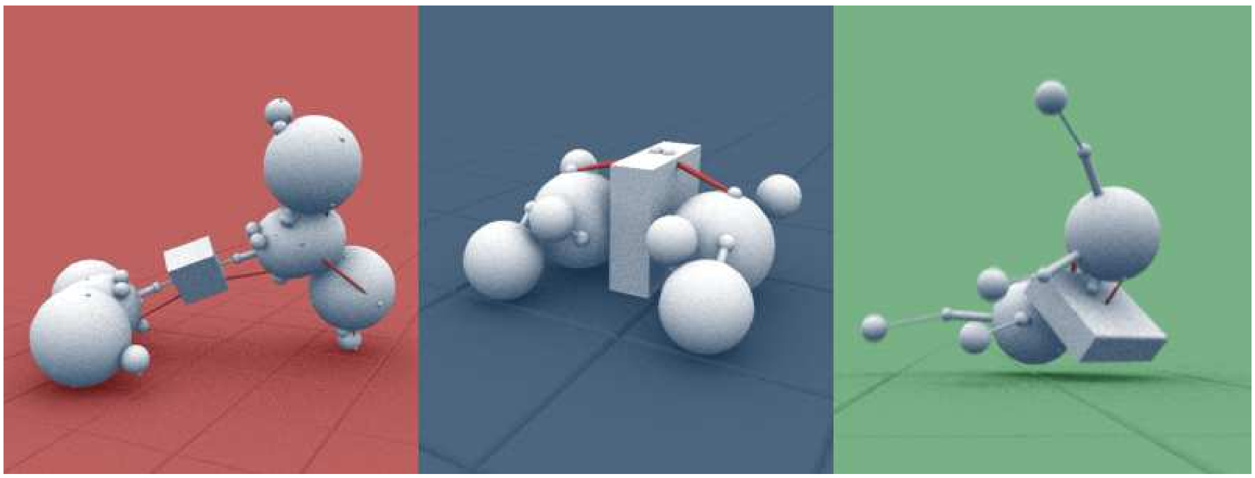

Figure 5.1: A selection of creatures produced using Extended ESP. These results illustrate some of the useful variety and multi-skill morphological adaptation produced by this system.

The previous chapter described the ESP system and demonstrated that it can evolve complex behavior. In this chapter, that system will be generalized to a significantly less constrained version that allows morphological changes to continue to evolve beyond the first skill. This extension improves both the fitness and the useful diversity of results, and constitutes this dissertation's second major contribution.

5.1 Motivation

The initial ESP implementation did achieve its goal of breaking the behavioral-complexity barrier for EVCs. However, it applied only to a significantly restricted case—one in which the most important morphology (i.e., the creature's skeletal segments and joints) were fixed after the first skill's evolution was complete.

Such a system is limited in its ability to discover interesting creatures. For example, what if a creature is evolved for an initial skill such as locomotion, then is asked to adapt to a largely orthogonal skill such as reaching up to a high target? That creature may or may not have the required morphological capacity for performing the second task, depending on the accidents of evolution.

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

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

5.2 Method

The extended ESP method is presented in two parts, the first of which describes its underlying concept, and the second its implementation.

5.2.1 Replacing Morphological Constraints with Retesting

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

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

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

Fortunately, the retesting can be reduced considerably by focusing it where it matters. Consider the syllabus graph shown in Figure 5.2. The skills that have a direct influence on the creature's body are shaded, and will be referred to as *leaf* skills. These are: FORWARD

LOCOMOTION, LEFT TURN, RIGHT TURN, and STRIKE. Once these skills are successfully established, the remaining non-leaf skills can be evolved independently (in an order that meets dependency requirements), without the need for any retesting. This approach stops the $\mathcal{O}(n^2)$ growth in testing requirements significantly earlier than would otherwise be possible—in this syllabus, for example, after four skills instead of 10 (assuming all leaf skills are learned first).

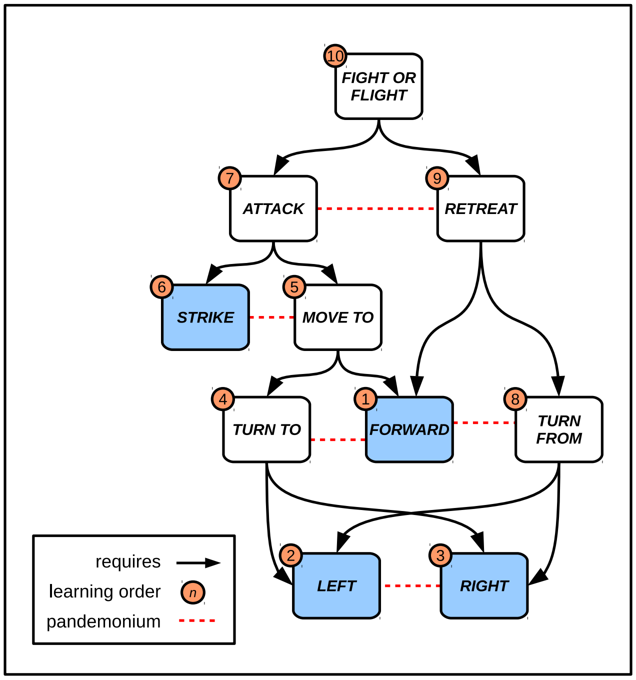

Figure 5.2: In this representation of the syllabus graph, shaded nodes are leaf nodes*, which act only on the body, rather than other nodes, and constitute the focus of the extended ESP system discussed in this chapter. While the leaf nodes may benefit from the extended system's continued morphological development, non-leaf nodes are not expected to require these additional evolutionary resources, and can be evolved using the original system, which produces only a linear growth in required testing.*

5.2.2 The Extended ESP Algorithm

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

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

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

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

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

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

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

Experiments with the extended ESP system demonstrate the advantages of the continuing morphological evolution enabled by this new algorithm. In Section 5.3 (*Strike Results*), an experiment from the original ESP system is reproduced in the extended ESP system, with dramatically different results. In Section 5.4 (*High-Reach Results*), a learning challenge designed to highlight the extended system's advantages is presented, and detailed benefits are described. Note that the extended ESP maintains original ESP's ability to construct complex hierarchical behaviors, and that ability is inherited largely without modification in the new system. Therefore, instead of simply replicating the fight-or-flight behavior of Chapter 4, the experiments in this chapter demonstrate the extended system's success in more challenging applications that were impossible in the original system. Video illustrating both of the result sections of this chapter can be viewed online at `http://youtu.be/fyVr7gdGEPE`.

5.3 Strike Results

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

5.3.1 Strike in Original ESP

Figure 5.3a depicts a creature evolved for locomotion in the underlying EVC system (which is common to both the original and extended versions of ESP). In Original ESP (as presented in Section 4.3.6), that creature consistently solved the challenge of producing a striking behavior by using its existing skeletal structure to either jump up and down or smash the ground with its limbs (best fitness in ten runs: 0.358), without any opportunity to explore the potential for new strategies or better adaptation that might result from continuing full morphological development.

5.3.2 Strike in Extended ESP

When the morphology is allowed to continue to evolve, however, new strategies become possible, and even old strategies may be better executed with morphological changes adapted to their specific needs. The extended ESP system develops a variety of such solutions, as can be seen in Figures 5.3b through 5.3f.

These creatures were produced using 20 parallel runs, with each employing a population size of 200 (Figures 5.4 and 5.5). As described in Section 5.2.2, each run began with a fixed number of generations (in this case 500), during which the new strike skill was rewarded as morphology and control were completely free to evolve. The minimum requirement for a nonzero fitness score was that the original locomotion skill continue to function to an acceptable level within the changing body (Algorithm 1). Specifically, 50% of the original locmotion score was required for 15 of the runs, with only 10% required for the other 5 runs. Useful results were obtained in the majority of runs with both of these settings. During the second stage of each run (Algorithm 2), an additional 250 generations were allowed for the first skill (locomotion) to reconcile itself to any morphological changes made during the first stage.

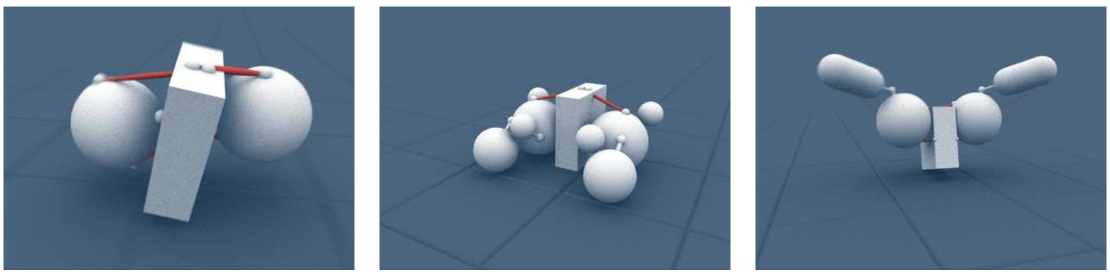

(a) Initial locomoting creature. *(b) Heavy smashing arms (Run 4).* *(c) Smashing flail arms (Run 20).*

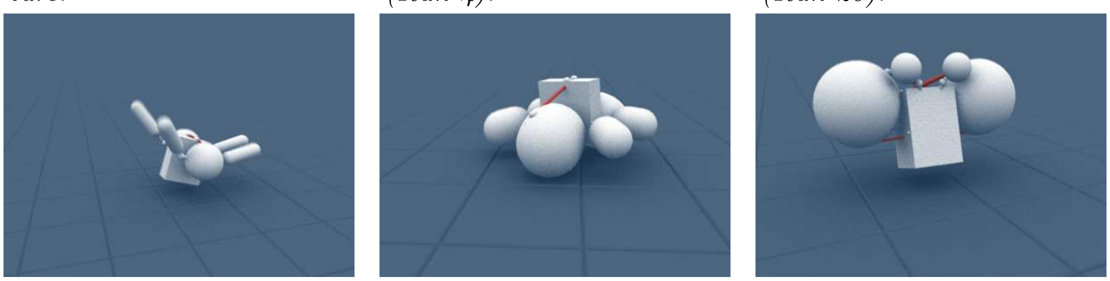

(d) Jump with anti-tip limbs (Run 5). *(e) Smashing tail, stabilizers (Run 8).* *(f) Jump with heavier body (Run 10).*

Figure 5.3: Further evolution of the hopper locomotion morphology in the strike task. (a) A creature adapted for locomotion. From this creature, creatures (b) through (f) were evolved using the extended ESP method described in this chapter. Each of them has developed a new technique (with corresponding morphological changes) for accomplishing an additional task—in this case, delivering a strike to the ground—while still maintaining the ability to perform the initial skill (locomotion) to prescribed levels. The extended ESP system makes such adaptations possible, resulting in morphology that supports multiple distinct skills.

Morphological changes seen in Figure 5.3b (run four, fitness 0.254) have produced a strong new strike technique and an appearance to match. The arms are now significantly heavier and can easily reach the ground to deliver impressive simultaneous smashing strikes directly. In Figure 5.3c (run 20, fitness 0.230), additional long segments are attached by hinges onto the end of the arms, allowing a striking technique reminiscent of a flail from

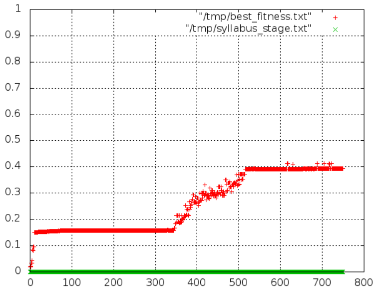

(a) Run 1

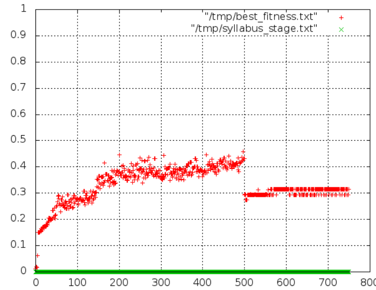

(b) Run 2

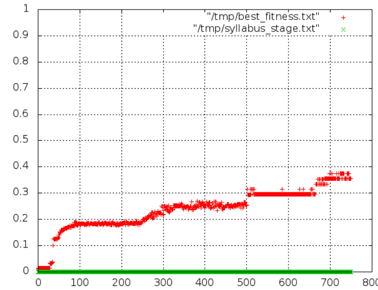

(c) Run 3

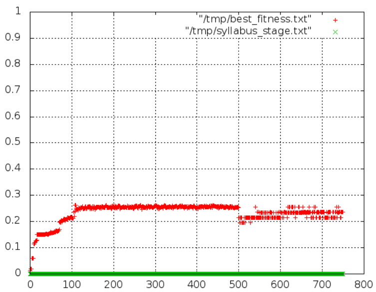

(d) Run 4

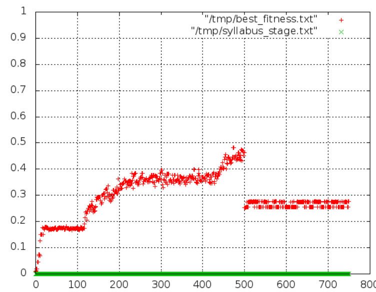

(e) Run 5

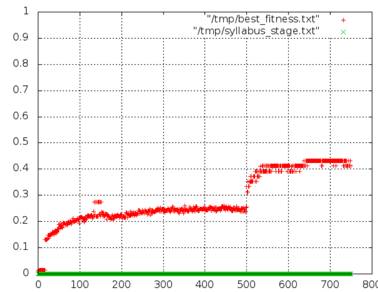

(f) Run 6

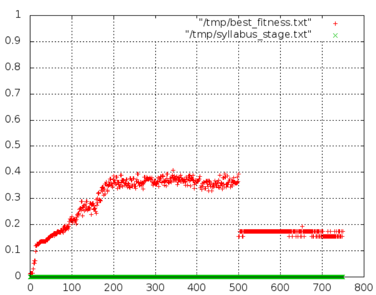

(g) Run 7

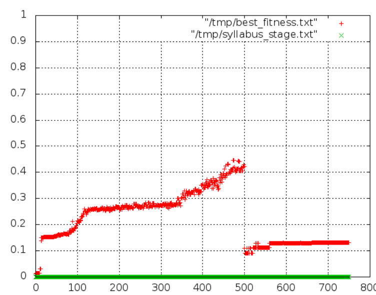

(h) Run 8

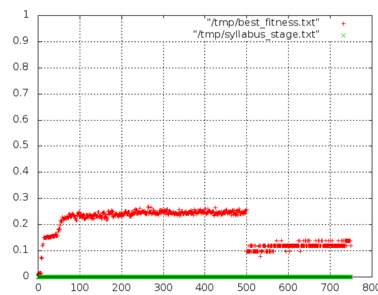

(i) Run 9

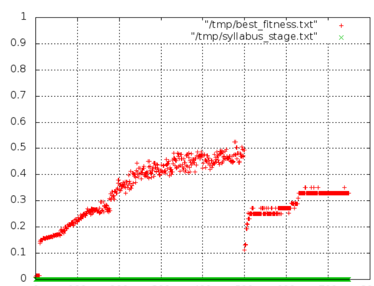

(j) Run 10

Figure 5.4: Fitness graphs for the first ten of 20 runs of Extended ESP. The five selected for Figure 5.3 (Runs 4, 5, 8, 10, and 20) were chosen because they demonstrated the most diverse, interesting, and successful solutions. Within each graph, the horizontal axis measures generations of evolution, and the vertical axis indicates fitness. As in all fitness graphs in this chapter, the first 500 generations are devoted to the development of body and brain for the new skill (in this case, strike), and the last 250 generations are used to allow the initial skill (locomotion) to reconcile itself to any changes in morphology.

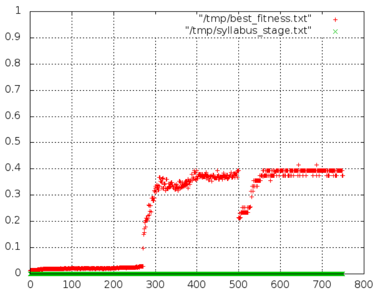

(a) Run 11

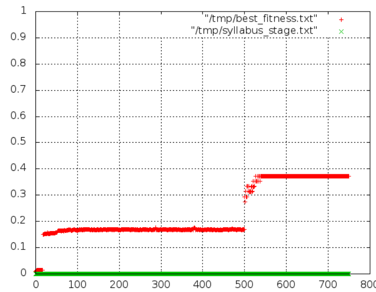

(b) Run 12

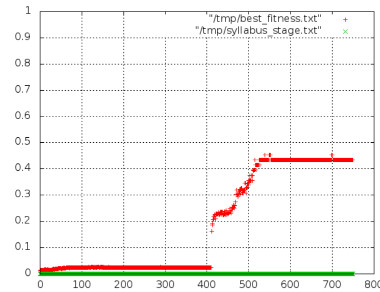

(c) Run 13

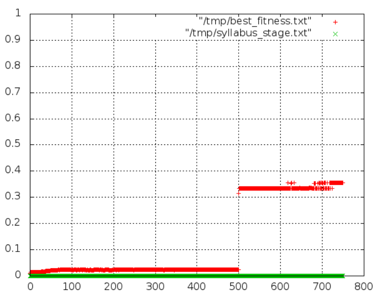

(d) Run 14

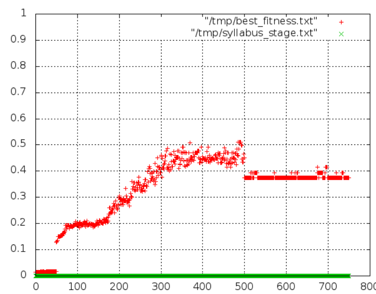

(e) Run 15

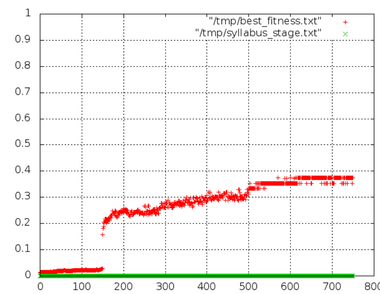

(f) Run 16

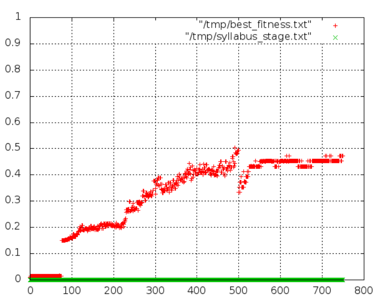

(g) Run 17

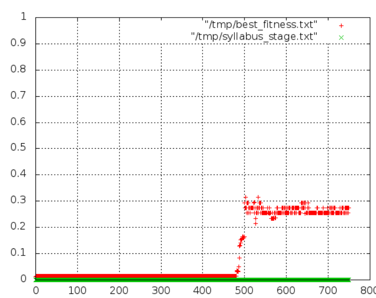

(h) Run 18

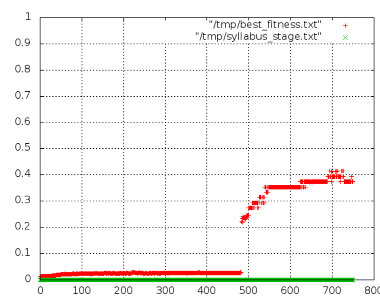

(i) Run 19

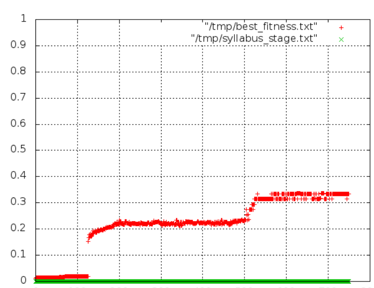

(j) Run 20

Figure 5.5: Fitness graphs for the last ten of 20 runs from which the strike results in Figure 5.3 were selected.

medieval Europe or Japanese nunchaku. The creature of Figure 5.3d (run five, fitness 0.462) employs a jumping strike similar to one produced in the original ESP system, but this time, long extensions to the limbs make it almost impossible for the creature to tip over, even during extremely energetic attacking leaps. A completely new technique is made possible by morphological changes in Figure 5.3e (run eight, fitness 0.423). Here, four low stabilizing legs have been added, allowing one of the original limbs to deliver a focused tail strike much like that of the ankylosaurus. Finally, the creature of Figure 5.3f (run ten, fitness 0.493) uses the same technique as the example from the original system, but with further morphological adaptations, such as a heavier body and heavier arms.

5.4 High-Reach Results

In this chapter's first experiment (Section 5.3.2), the goal was to demonstrate the extended system's benefits over the original system, when applied to the previously seen strike task in the hopper morphology (Figure 5.3a). In this section, the experiment is specifically designed to highlight the extended system's advantages by selecting a task that is more different from locomotion than the strike task: that of reaching a high target. Specifically, a selection of three different locomoting creatures was evolved, using both the original and extended ESP systems, and the differences in results were examined in detail. For this experiment, fitness was defined as the maximum height reached by any part of the creature in a two-second evaluation interval (relative to the creature's start height), averaged across five successive intervals.

5.4.1 Original Creature 1: Hopper

In this experiment, the locomoting creature of Figure 5.3a was evolved toward new high-reach goal. As in the previous experiment, the results are from 20 independent runs (Figures 5.7 and 5.8), each with a population size of 200, using 500 generations for full morphological and control development of the new skill (high reach) followed by 250 generations to allow the original skill (locomotion) to reconcile itself to the new body.

In the Original ESP system, only two strategies were observed, within which the results were extremely uniform. Using core morphology unchanged from the original locomotion result, all such creatures developed to either jump as high as possible (Figure 5.9), or reach a limb up by tipping over onto the other limb (as seen in Figure 5.6 (a) and (e), but without Extended ESP's beneficial morphological adaptations). In both cases, the results (Figure 5.10, with best fitness in ten runs: 0.182) were limited by the inability of skeletal morphology to adapt to this new task.

In the extended ESP system, in contrast, a wide variety of results was observed, in which a number of novel strategies were used, often to great effect. The creature of

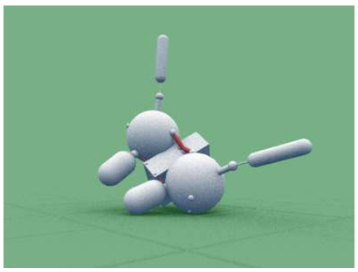

(a) Tipping, long new limbs (run 2).

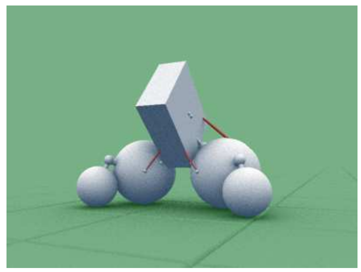

(b) Push-up, extended limbs (run 16).

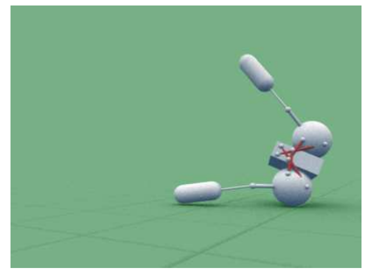

(c) Telescoping limbs (run 1).

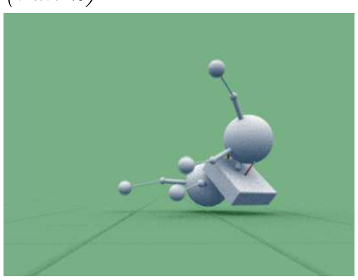

(d) Telescoping, anti-tip limbs (run 14).

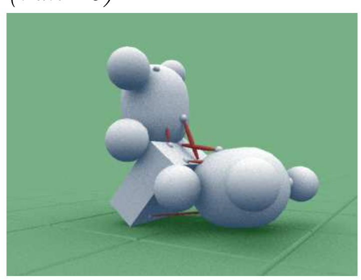

(e) Tip with enlarged limbs (run 11).

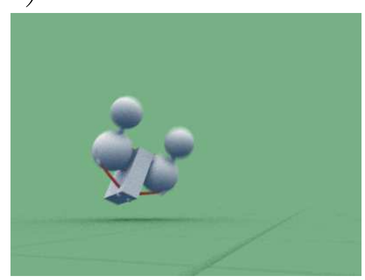

(f) Jump, swing extensions up (run 19).

Figure 5.6: Further evolution of the hopper morphology in the high-reach task. The locomoting creature of Figure 5.3a was further evolved using the extended ESP system to adapt to a high-reach task. The results demonstrate the potential of continued morphology evolution to produce a great degree of useful variety.

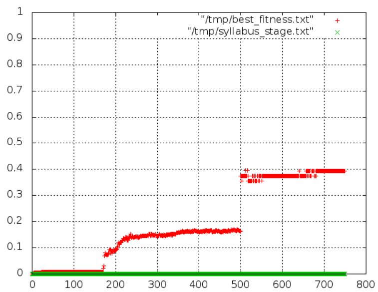

(a) Run 1

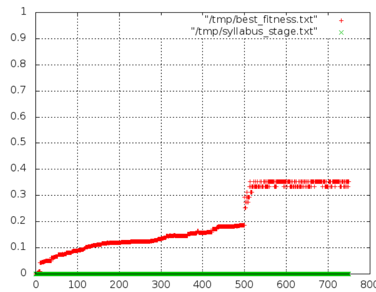

(b) Run 2

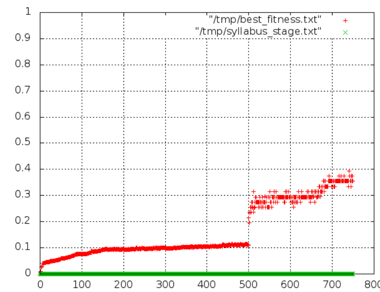

(c) Run 3

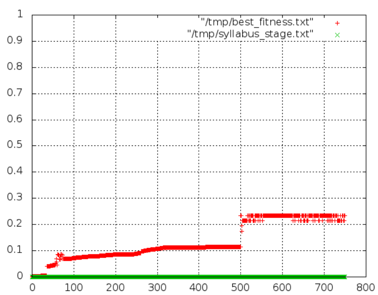

(d) Run 4

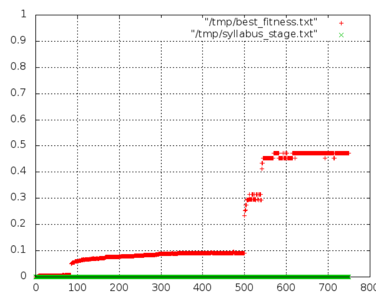

(e) Run 5

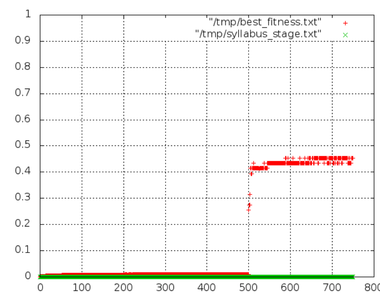

(f) Run 6

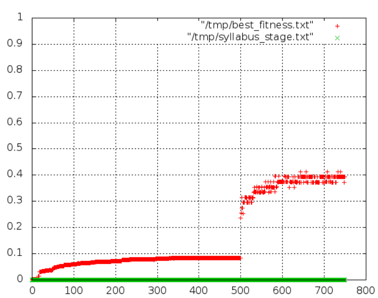

(g) Run 7

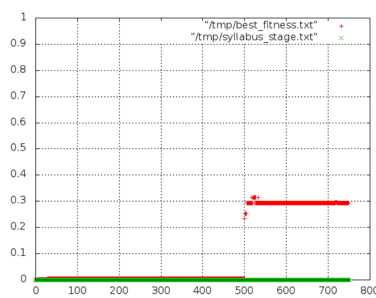

(h) Run 8

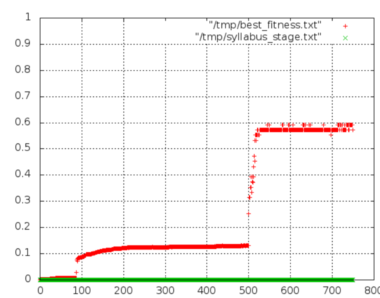

(i) Run 9

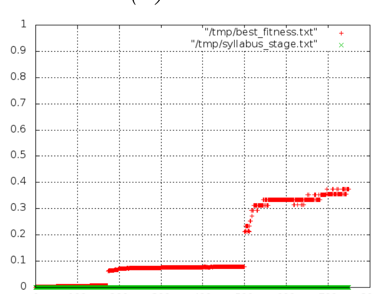

(j) Run 10

Figure 5.7: Fitness graphs for the first ten of 20 runs of Extended ESP. The six selected for Figure 5.6 (Runs 1, 2, 11, 14, 16, and 19) were chosen because they demonstrated the most diverse, interesting, and successful solutions.

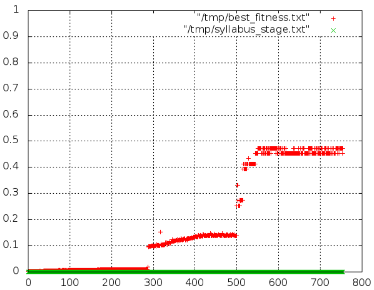

(a) Run 11

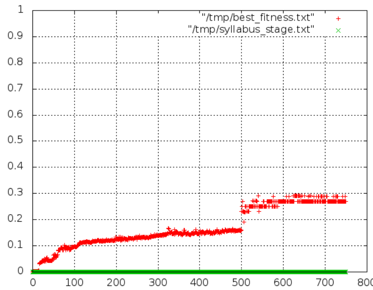

(b) Run 12

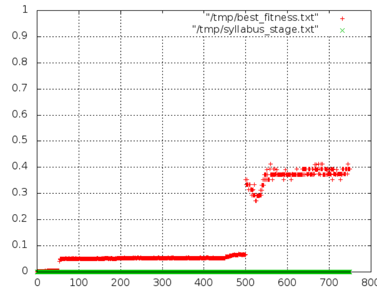

(c) Run 13

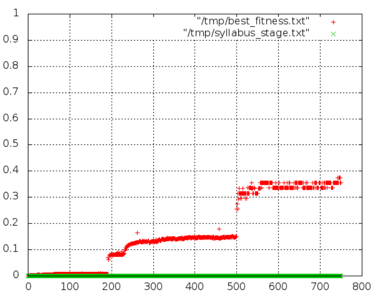

(d) Run 14

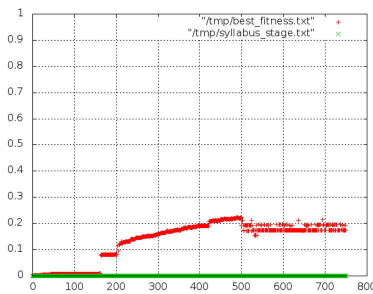

(e) Run 15

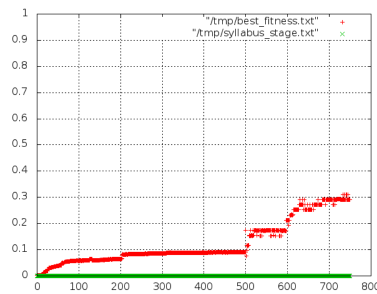

(f) Run 16

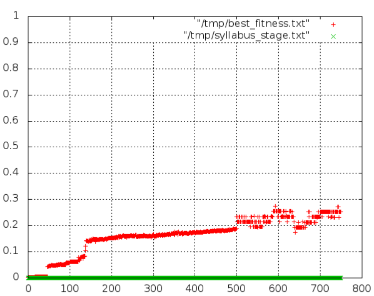

(g) Run 17

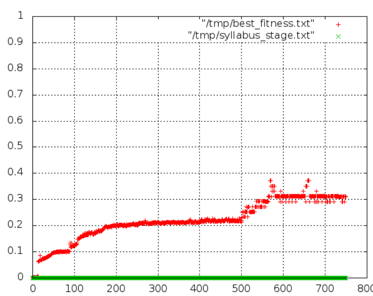

(h) Run 18

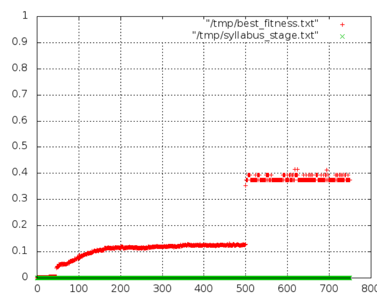

(i) Run 19

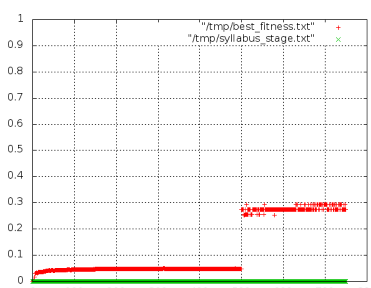

(j) Run 20

Figure 5.8: Fitness graphs for the last ten of 20 runs of Extended ESP from which the high-reach results in Figure 5.6 were selected.

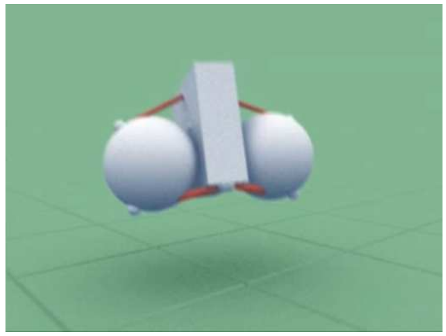

Figure 5.9: A simple jump for height with the existing morphology was produced in eight of ten runs of the Original ESP system on the learning task of Section 5.4 (high reach, hopper morphology).

Figure 5.6a (run two, fitness 0.186) employs a tipping strategy similar to that produced in the Original ESP system, but the capacity for morphological changes permits long limb extensions that greatly increase the creature's reach as well as allowing greater freedom to tip without falling over. In Figure 5.6b (run 16, fitness 0.090), a completely new strategy is seen, in which the augmented limbs are used to push the creature's root segment up for a high reach that again would not have been possible in original ESP. In Figure 5.6c (run one, fitness 0.163), another new strategy was discovered: long telescoping limbs are cast upward as part of a tipping action for an extremely high reach. (Note that, while this morphological adaptation is made in such a way that locomotion is preserved, this type of body plan would be unlikely to develop in the original system's locomotion-only morphological evolution.) The creature depicted in Figure 5.6d (run 14, fitness 0.150) combines the strategies seen in Figures 5.6a and c, with the addition of both telescoping and anti-tip limbs. Figure 5.6e (run 11, fitness 0.140) demonstrates a solution similar to the original ESP's tipping strategy, but this time with greatly enlarged limbs for a higher reach, along with smaller limb additions that prevent the creature from falling over during the reaching process. The creature of Figure 5.6f (run 19, fitness 0.128) produces a novel take on the original system's jumping strategy with the addition of swinging limb extensions that increase its upward reach at the height of its leap.

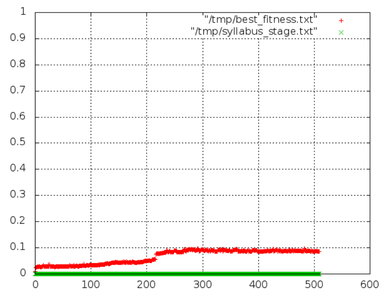

(a) *Run 1*

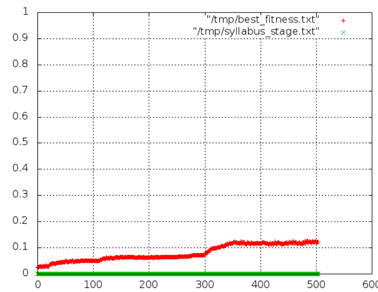

(b) *Run 2*

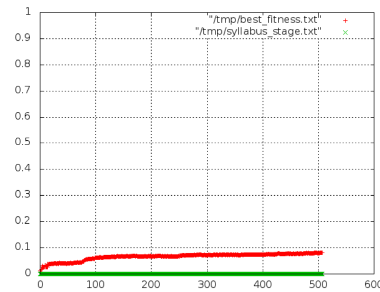

(c) *Run 3*

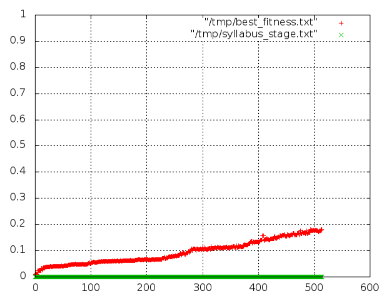

(d) *Run 4*

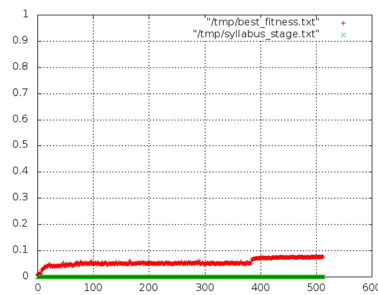

(e) *Run 5*

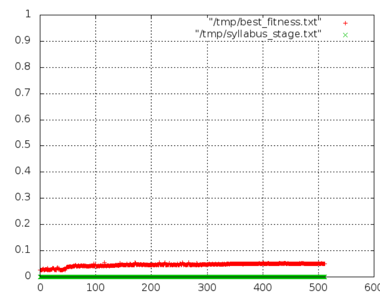

(f) *Run 6*

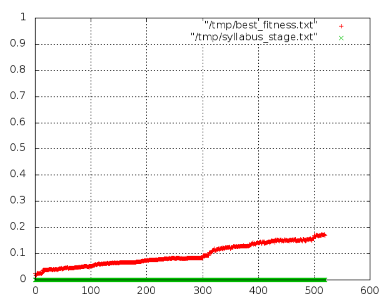

(g) *Run 7*

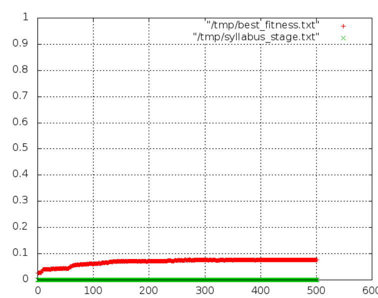

(h) *Run 8*

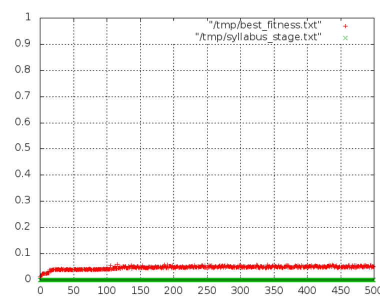

(i) *Run 9*

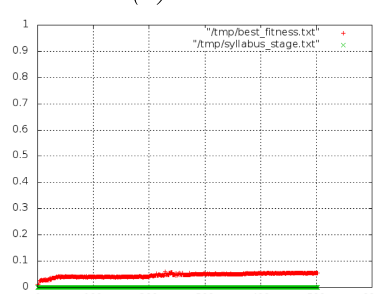

(j) *Run 10*

Figure 5.10: Fitness graphs for ten runs of Original ESP on the task of Section 5.4.1 (high reach with hopper morphology). Due to Original ESP's constraints on morphology, these runs produced extremely uniform results, with only two strategies observed.

5.4.2 Original Creature 2: Snake

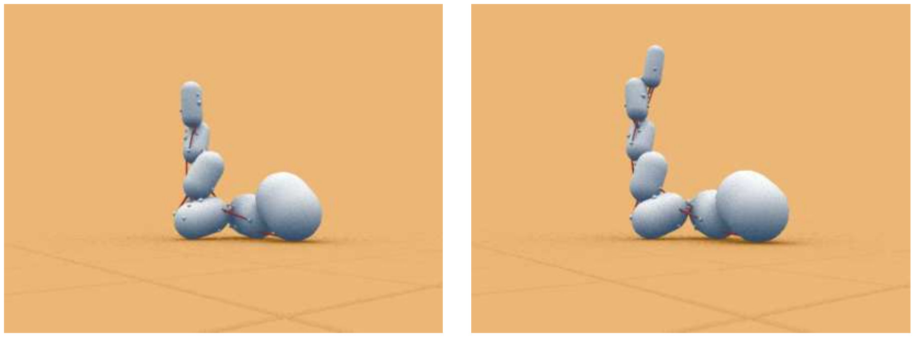

(a) *Original ESP result.* (b) *Result in new ESP (run 1).*

Figure 5.11: Further evolution of a snake morphology in the high reach task. These results demonstrate how the extended ESP system (b) can produce better fitness values (i.e., a higher reach) than the original ESP system (a) by allowing the addition of new body segments.

Another successful solution to the locomotion task produced by the Original ESP is shown in Figure 5.11a. This snake-like creature achieved a high reach by extending one end of its long morphology, while the rest of the body maintained balance. This creature's performance in the high-reach task provides an especially clear example of how Extended ESP can provide improved results over Original ESP.

As in the preceding two experiments, a population size of 200 was applied to 500 generations of new-skill body-and-brain adaptation, followed by 250 generations of first-skill reconciliation. In this case, even the Extended ESP system produced rather uniform results across multiple runs, and therefore the results were obtained from only 10 runs.

The Original ESP produced a high fitness of 0.174 in ten runs (Figure 5.13). Extended ESP improved upon this creature by adapting its morphology to the secondary task, while its strategy remained unchanged, as seen in Figure 5.11b (Run 1 of Figure 5.12, fitness 0.267). It grew an additional body segment that enabled the higher reach, while allowing it still to perform locomotion to acceptable standards. In this manner, further evolution can adapt existing morphology while preserving prior function.

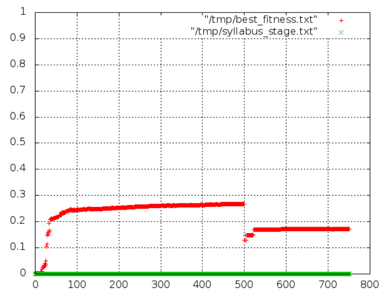

(a) Run 1

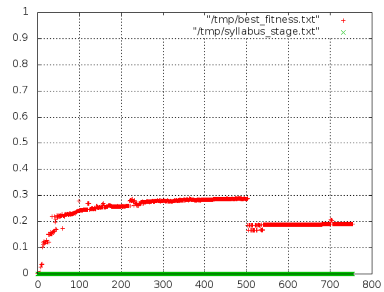

(b) Run 2

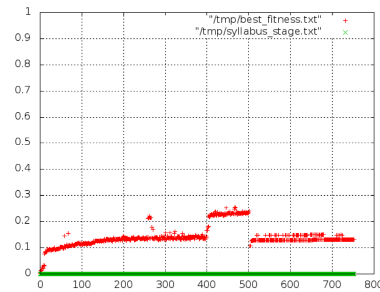

(c) Run 3

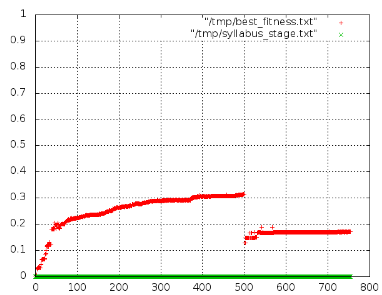

(d) Run 4

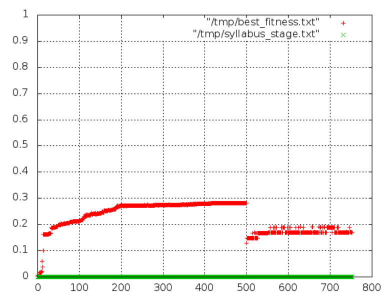

(e) Run 5

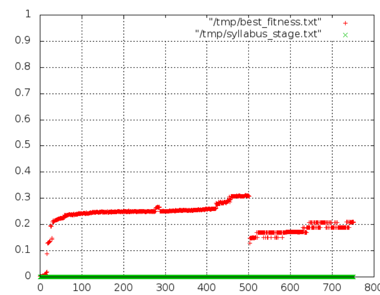

(f) Run 6

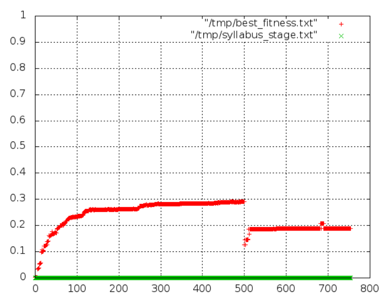

(g) Run 7

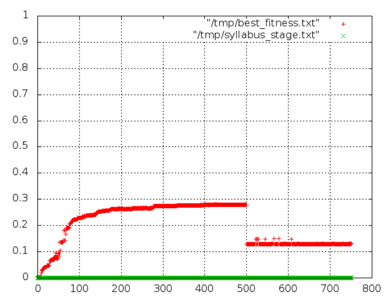

(h) Run 8

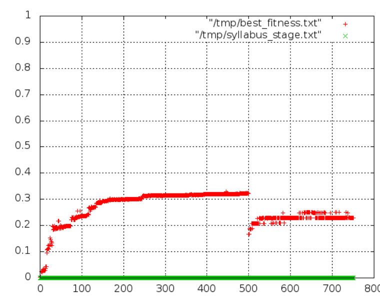

(i) Run 9

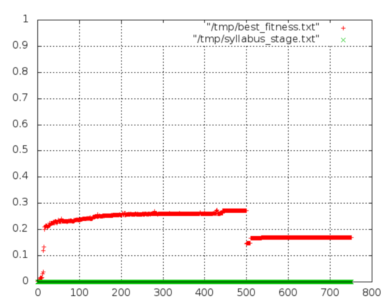

(j) Run 10

Figure 5.12: Fitness graphs for all 10 runs of Extended ESP from which the high-reach result in Figure 5.11b (high reach with snake morphology) was selected. Extended ESP's ability to change morphology produced a distinct increase in fitness over Original ESP (see Figure 5.13).

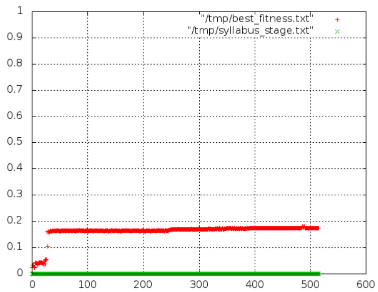

(a) *Run 1*

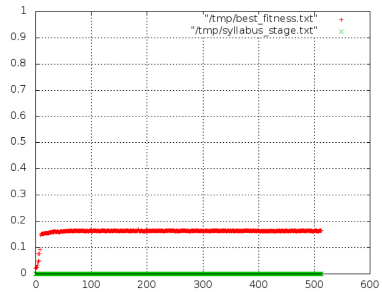

(b) *Run 2*

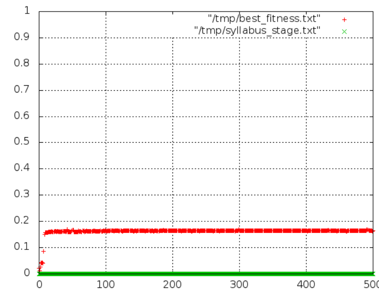

(c) *Run 3*

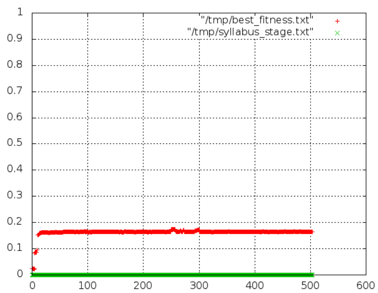

(d) *Run 4*

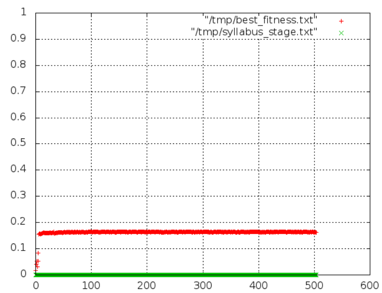

(e) *Run 5*

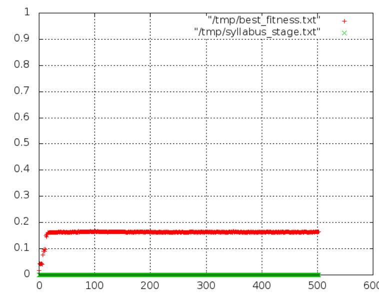

(f) *Run 6*

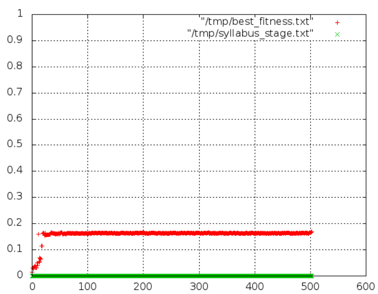

(g) *Run 7*

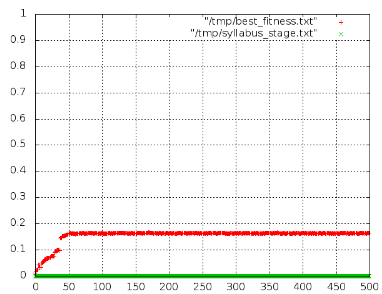

(h) *Run 8*

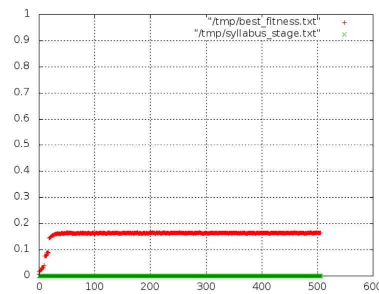

(i) *Run 9*

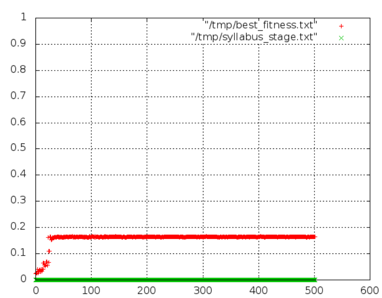

(j) *Run 10*

Figure 5.13: Fitness graphs for 10 runs of Original ESP on the task of Section 5.4.2 (high reach with snake morphology). This system's inability to fully adapt morphology produced significantly less fit results than Extended ESP.

5.4.3 Original Creature 3: Quadruped

The relatively complex quadruped seen in Figure 5.14a was a third type of solution developed by the underlying EVC system for the locomotion task (with all experimental details as in the previous examples and the use of 20 parallel starts; Figures 5.15 and 5.16). In continued evolution of the high-reach task in the Original ESP system, this creature's results were again extremely uniform in approach and fitness. They all reached up with a single limb (Figure 5.17), and all with approximately equal success (producing a top fitness of 0.164 in ten runs, shown in Figure 5.18).

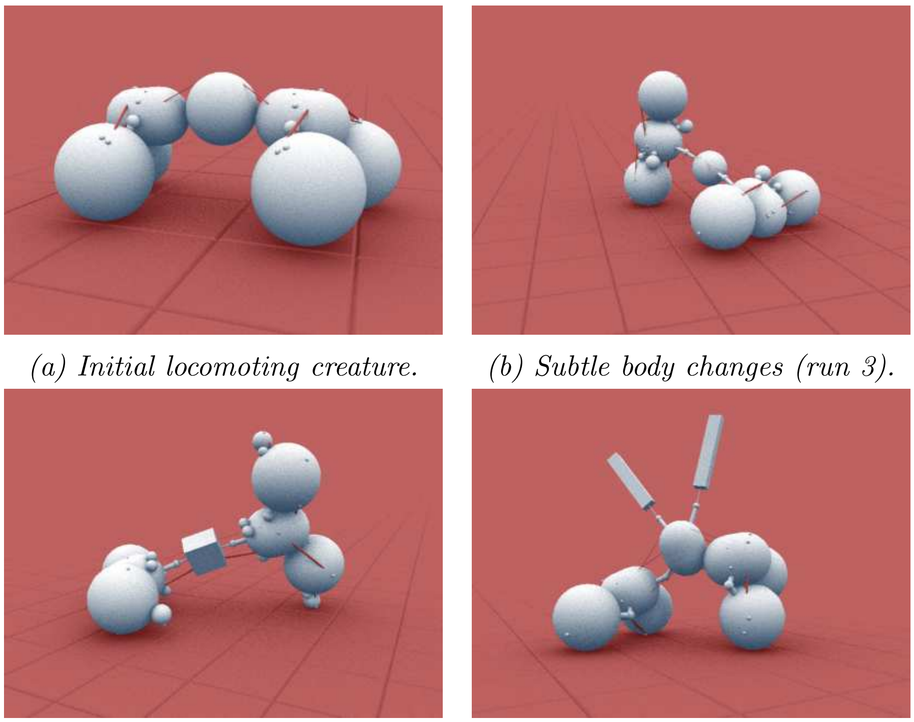

(a) Initial locomoting creature.

(b) Subtle body changes (run 3).

(c) More obvious body changes (run 5).

(d) Dramatic changes in morphology (run 14).

Figure 5.14: Further evolution of a quadruped morphology in the high reach task. The initial locomoting quadruped (a) is evolved for high reach in the Extended ESP system (b)-(d). Through a variety of strategies, each of the Extended ESP creatures shown scores better on this new task than any creature from the Original ESP system.

In the Extended system, the ability to continue to adapt morphology to this new task led to a diverse set of useful results. All of those depicted in Figure 5.14 also were more fit than those produced with the Original ESP. For example, Figure 5.14b (Run 3, fitness 0.209) illustrates a creature that pursues the same strategy as the creature in Figure 5.14a, yet does so more effectively due to subtle morphological adaptations, such as changes in segment dimensions. In Figure 5.14c (Run five, fitness 0.294), more obvious changes have

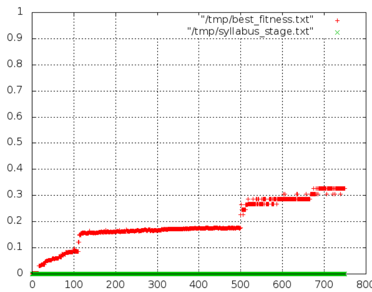

(a) Run 1

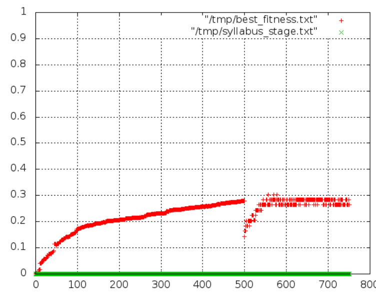

(b) Run 2

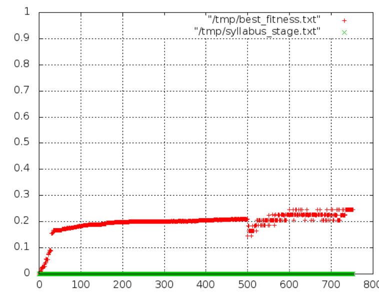

(c) Run 3

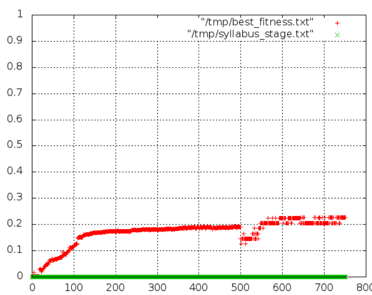

(d) Run 4

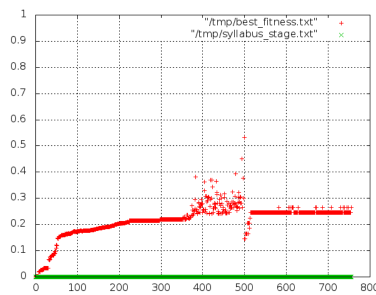

(e) Run 5

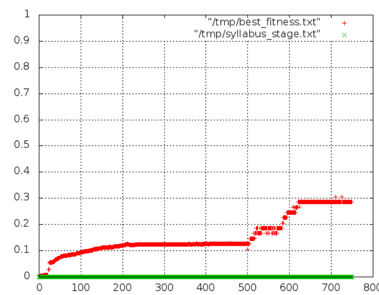

(f) Run 6

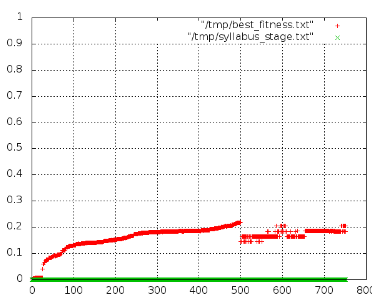

(g) Run 7

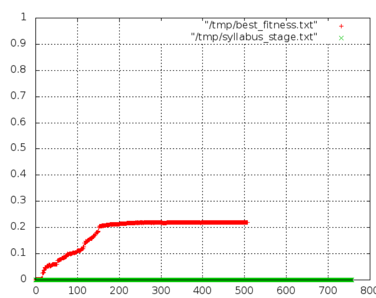

(h) Run 8

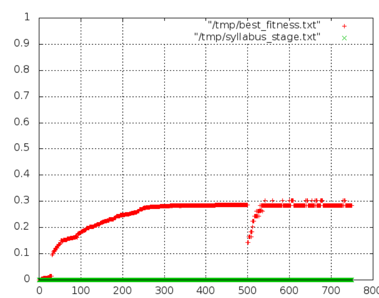

(i) Run 9

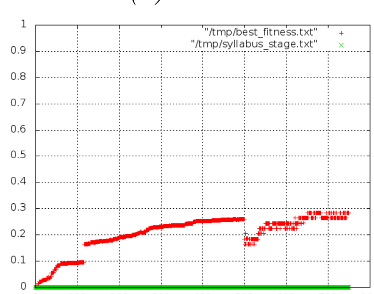

(j) Run 10

Figure 5.15: Fitness graphs for the first ten of 20 runs from which the quadruped high-reach results in Figure 5.14 were selected. Extended ESP's ability to fully adapt morphology produced results which were both more varied and more fit than those of Original ESP.

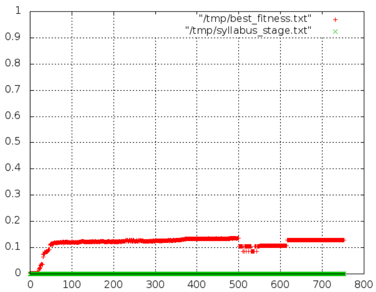

(a) Run 11

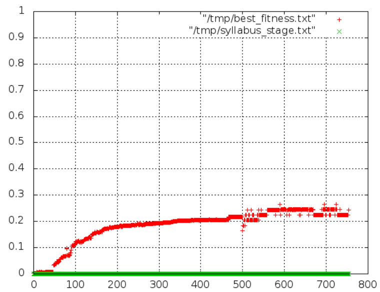

(b) Run 12

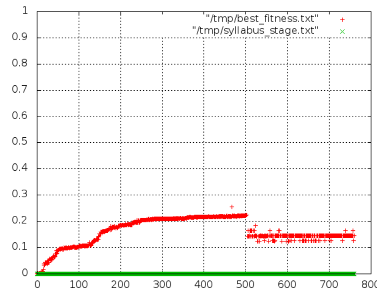

(c) Run 13

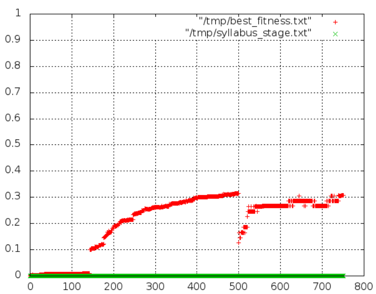

(d) Run 14

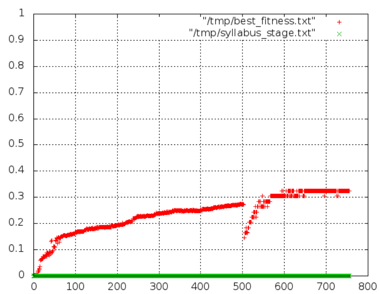

(e) Run 15

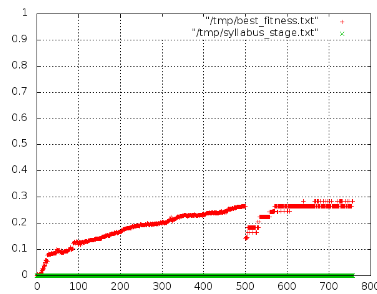

(f) Run 16

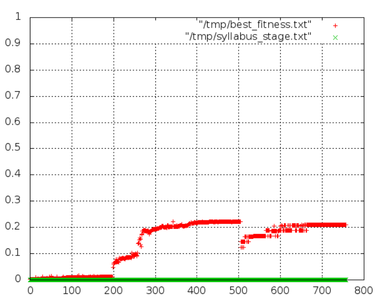

(g) Run 17

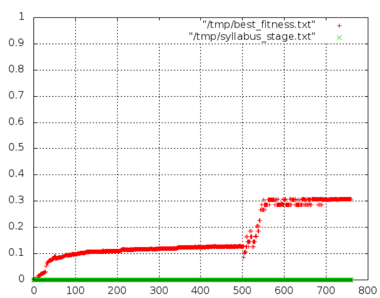

(h) Run 18

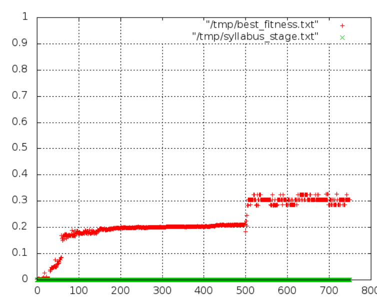

(i) Run 19

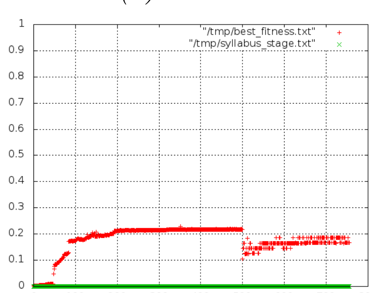

(j) Run 20

Figure 5.16: Fitness graphs for the last ten of 20 runs from which the high-reach results in Figure 5.14 were selected.

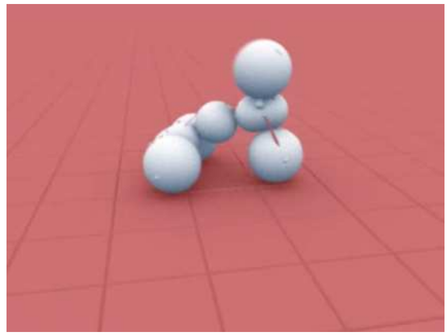

Figure 5.17: The only technique developed in ten runs of the Original ESP system on the learning task of Section 5.4.3 (high reach, quadruped): reaching up with one limb. Due to the fixed morphology in Original ESP, this resulted in almost exactly the same score for all ten runs.

been made to the body. In particular, new segments have been added to the ends of the limbs which provide an increased reach. These changes made it possible to further exceed the uniform performance limit of the original creature, while still employing the same basic high-reach technique. In Figure 5.14d (run 14, fitness 0.314), even more dramatic changes to morphology provide a new way of accomplishing the high reach task: This creature employs a new pair of tall, dedicated limbs to even further exceed the previous system's performance—another clear example of an evolutionary path to improved fitness that was unavailable in the original ESP system.

Interestingly, in the case of 5.14c , these body-plan changes were accompanied by a very different but effective new form of locomotion in a direction perpendicular to the original. This result demonstrates how evolution in further tasks can result in novel solutions in earlier tasks as well. In other words, further evolution does not only add more structure to existing solutions, but it can fundamentally change these solutions as well.

5.5 Discussion

The results in this chapter demonstrate that (1) it is useful to continue to adapt the morphology when adding a new task, and (2) such continued adaptation can leverage the diversity in the initial morphology and produce a variety of interesting solutions.

Although the Extended ESP algorithm removed the original system's explicit limitations on body changes after the first skill, development of morphology throughout the acquisition of complex skills is still not fully general and completely unlimited. First, the retesting requirements would make morphological development impractical if continued through too many steps of leaf skill addition. To mitigate this issue in the future, it may be possible to do the retesting periodically rather than universally, and to run the tests in parallel. Also, the more leaf skills there are, the more likely it is that the morphological change required by

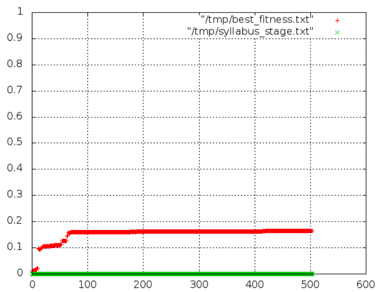

(a) Run 1

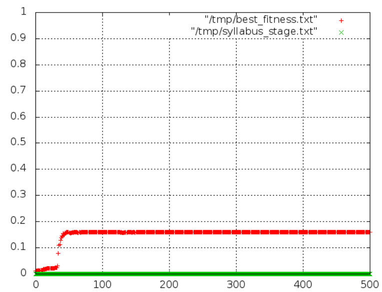

(b) Run 2

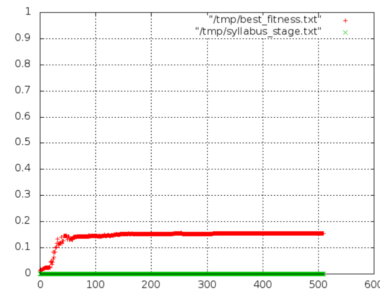

(c) Run 3

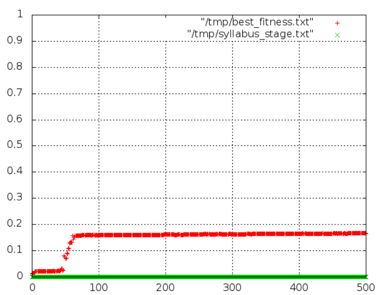

(d) Run 4

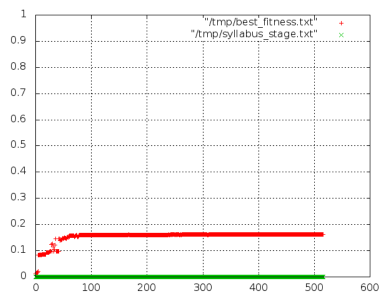

(e) Run 5

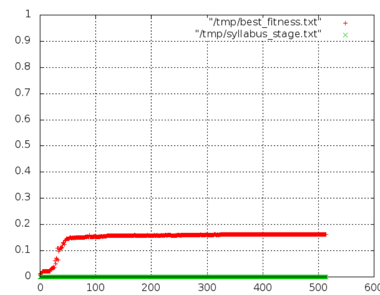

(f) Run 6

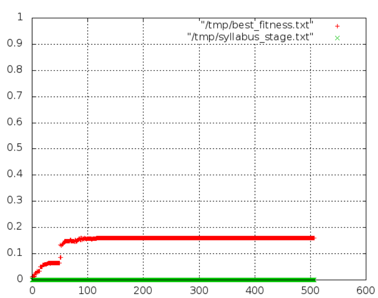

(g) Run 7

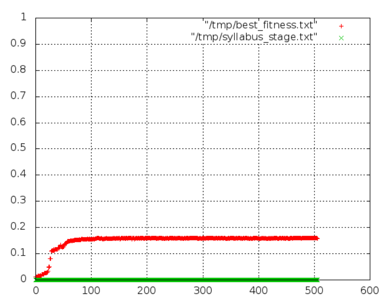

(h) Run 8

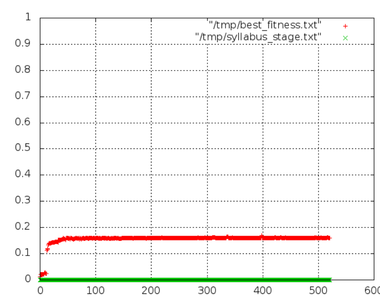

(i) Run 9

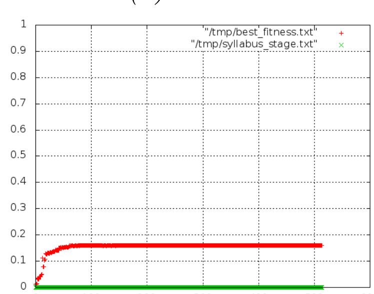

(j) Run 10

Figure 5.18: Fitness graphs for ten runs of Original ESP on the task of Section 5.4.3 (high reach with quadruped morphology). In this case, Original ESP's constraints on morphological adaptation produce results which are both extremely uniform and less fit than those of Extended ESP.

one skill will be harmful to the others. This limitation may be more difficult to overcome, because it reflects the inherently conflicting demands that any creature would face in such an environment. As in the real world, some compromise is expected to evolve, or perhaps a set of solutions that compromise in different ways, or create different niches. Such variety exists in biology and it would also be expected in artificial creatures.

5.6 Conclusion

This chapter described this dissertation's second major contribution: a useful generalization of the original ESP system's specialized implementation into one that allows morphological adaptation to multiple tasks, improving both performance and diversity. The next chapter presents the third major contribution of this dissertation: a novel actuator system for evolved virtual creatures that results in an increase in meaningful morphological complexity and a significant reduction in cognitive load.

Chapter 6

Muscle Drives

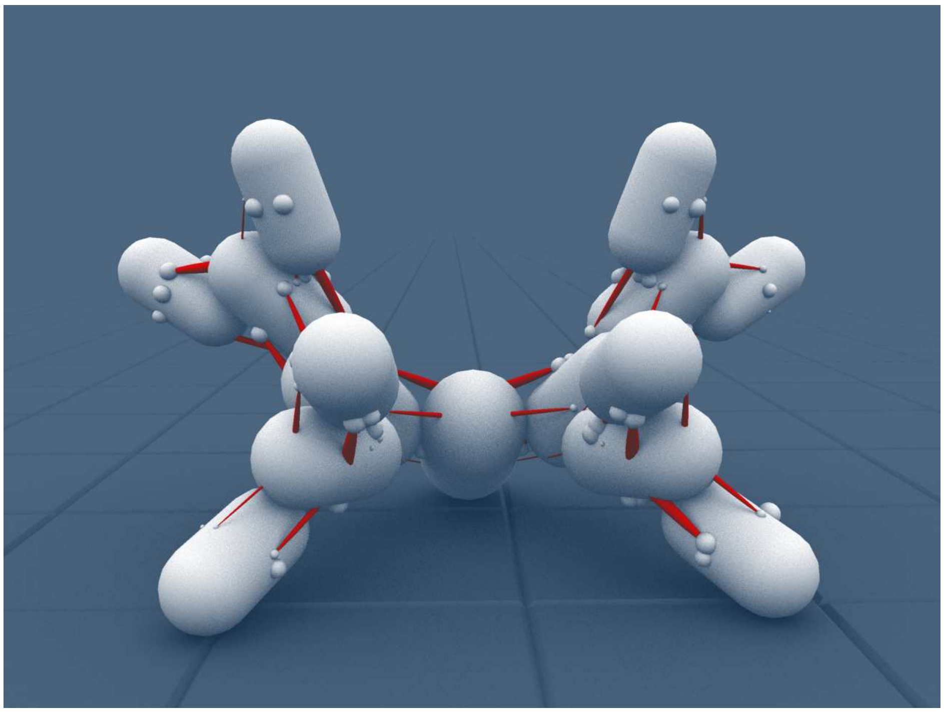

Figure 6.1: A creature evolved for jumping (Section 6.4.6) using the method described in this chapter, demonstrating the morphological complexity that results from replacing implicit joint-motor drives with an evolvable musculature. As with all other examples in this chapter, the physical intelligence embodied by these muscle drives enables this creature to perform a useful task essentially without control intelligence (Section 6.3). Video of this and all other results from this chapter can be viewed at http://youtu.be/csZ9JZcuBfE.

The preceding two chapters described the Original and Extended ESP systems, which served to increase the behavioral complexity of evolved virtual creatures. This chapter describes this dissertation's third major contribution: a novel muscle drive system that increases morphological complexity in a meaningful, bio-mimetic fashion, while simultaneously reducing the control requirements imposed on creature brains.

6.1 Motivation

In addition to behavioral complexity, morphological complexity is an important goal for evolved virtual creatures (Figure 6.1) [4]. How can it be increased to approach the morphological complexity of creatures evolved in the real world? Traditional segmented EVCs [46, 7, 30, 26] 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 [4, 22, 8] demonstrated a different—and arguably greater morphological complexity, albeit based on indirect developmental mechanisms. In contrast, the technique described in this chapter 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., attached to every free axis of every joint, and fixed over time. They are also typically unseen, perhaps because they are all the same, and therefore displaying them provides no useful information.

However, the ESP system 4 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 example, the muscles were controlled by a complex brain, one particularly interesting property of such drives is that they do not always require this control complexity. As will be shown in this chapter, these muscles can 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 locomotion.

Besides simplicity, there's another beneficial result of this shifting of intelligence from brain to body: where the control intelligence was invisible, the physical intelligence that replaces it is visible, in the morphological complexity of the muscles. 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 of muscles at a joint, their size (with rendered thickness indicating strength), orientation (indicating direction of force), and their attachment points.

Many 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 this chapter, the implementation of the muscle drives is first presented in detail. The extremely minimal control that it enables is specified, followed by experimental results in creating jumping and locomotion for EVCs with varying morphologies.

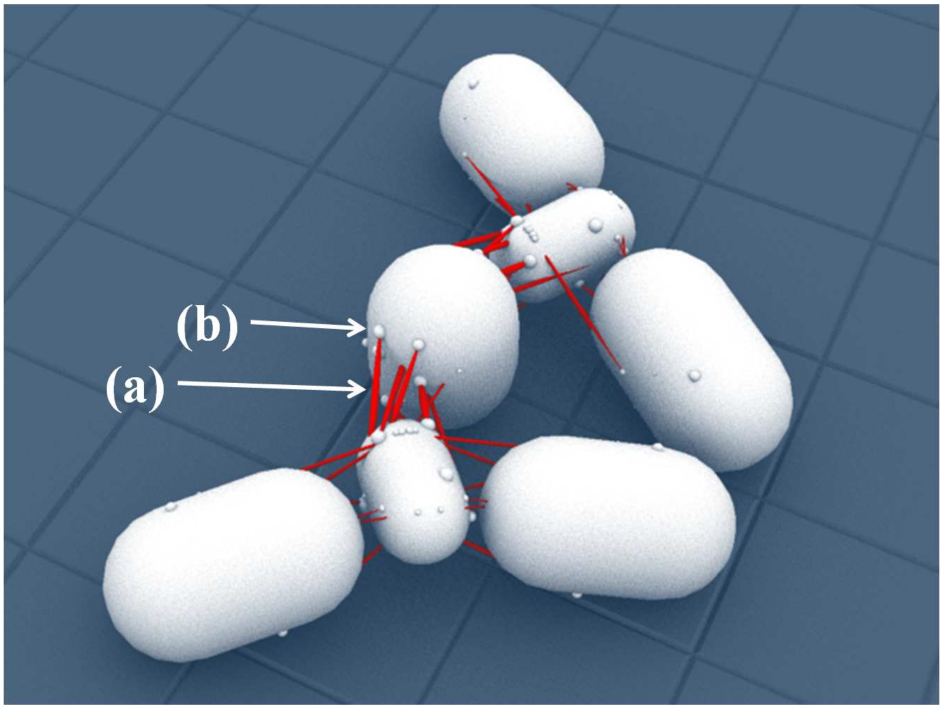

Figure 6.2: Evolvable musculature, with example muscle body (a) and attachment point (b) indicated. The density of muscles at a joint, their thickness (indicating current force), orientation, and attachment points all contribute meaningfully to the creature's morphological complexity.

6.2 Muscle Implementation

The muscle drives are implemented as simple linear springs. Each muscle (Figure 6.2(a)) is completely described by its attachment points and maximum strength. An attachment point (Figure 6.2(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. Note that this implementation—with numerous joints of varying types between a single pair of rigid body segments—is not typical for PhysX, and indeed initial results with normal settings resulted in simulations that were not sufficiently stable. All experiments presented here rely on a much smaller simulation step—1/240th of a second—as well as other configuration settings, all of which can make the simulation significantly more expensive.

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 can potentially embody sufficient physical intelligence to perform basic behaviors with only simple control intelligence required, as will be described next.

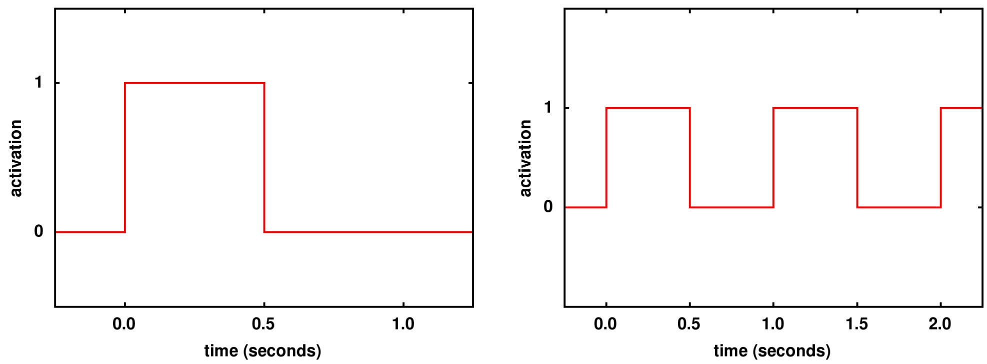

(a) Jump activation (single square pulse). *(b) Locomotion activation (repeating square wave).*

Figure 6.3: The fixed global muscle activations that replace the typical EVC's relatively complex brain for all experiments in this work are illustrated. With the muscle drives' capacity for physical intelligence, simple but useful behaviors can be performed effectively without control intelligence.

6.3 Minimal Control

As a demonstration of approximately how much physical intelligence the evolvable musculature can embody, the examples in the subsequent Results sections all function almost entirely without control intelligence (Figure 6.3).

In a conventional evolved virtual creature, control intelligence is implemented as a neural network [30] or a directed graph of simple computing nodes [46] (Chapters 4 and 5), as shown in Figure 2.11. In the creatures of this chapter, far less is required. For these 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. 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 6.3.

6.4 Jump Results

In this section (6.4) and the next (6.5), 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 video at `http://youtu.be/csZ9JZcuBfE`.

6.4.1 Experimental Setup

For all of the results in this section, the population size was approximately 100, initially filled with single-segment genotypes of random dimensions, and the results were obtained with between 221 and 500 generations. Forty independent runs of the experiment were executed, each with its own random seed. The champions of these 40 runs are diverse, although certain variations of morphological themes tend to recur. Illustrative examples are presented in the sections below. The variety demonstrated in these results suggests that the approach should scale well to more challenging tasks, as discussed in Section 6.6.

The following five examples 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. Each stage is complete when a sufficient fraction of the population—on the order of 5%—has achieved full fitness. At that point, the individuals with full fitness are 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. The shaping progressed in three steps.

1. **static AABB top**
   For this step, full fitness is achieved by producing a creature so tall (at rest) that it cannot contain only one segment. In this way, it is ensured that creatures will contain joints, which will permit the addition of muscles in future steps.

2. **static AABB top + highest AABB top**
   In this step, half of the fitness (static AABB top) depends on maintaining sufficient static size to ensure that joints are retained. The other half of the fitness score (highest AABB top) is used to encourage the addition of muscles which produce upward motion.

3. **highest AABB top + highest AABB bottom**
   In this step, half of the fitness (highest AABB top) encourages the retention of upward-motion-producing muscles developed in the preceding step. In addition, the other half of the fitness score (highest AABB bottom) encourages the ultimate goal: getting the creature's lowest point as high off the ground as possible during the jump.

For all jump evolution experiments, control consists solely of the single fixed global activation signal depicted in Figure 6.3(a), with all other required intelligence residing entirely within the body, including the evolvable musculature.

In each result illustration (Figures 6.4-6.8), the left and right sides show the creature before and during its jump, respectively.

6.4.2 Jump Result 1: Two-Armed Swing (Repeatable)

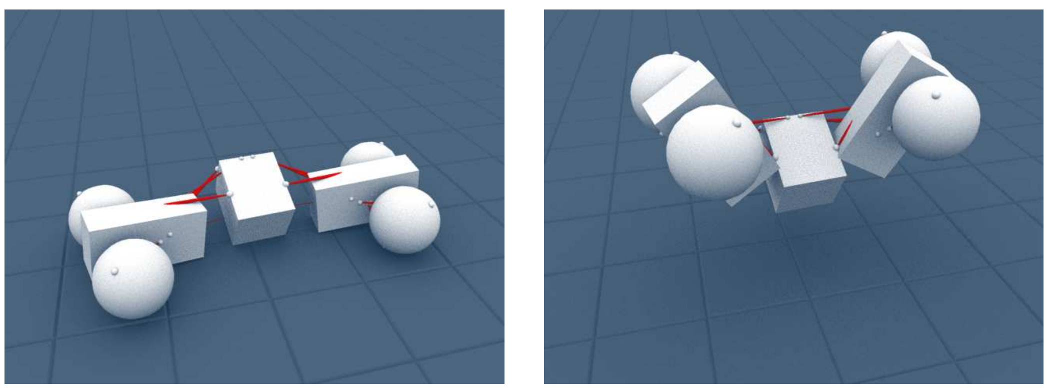

Figure 6.4: Two-armed swing (repeatable), from generation 300 of Run 4 of the jump task.

The creature in Figure 6.4 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. This result is from generation 300 of Run 4 (Figure 6.9) and achieved a fitness of 0.830.

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.

6.4.3 Jump Result 2: Two-Armed Swing (Non-Repeatable)

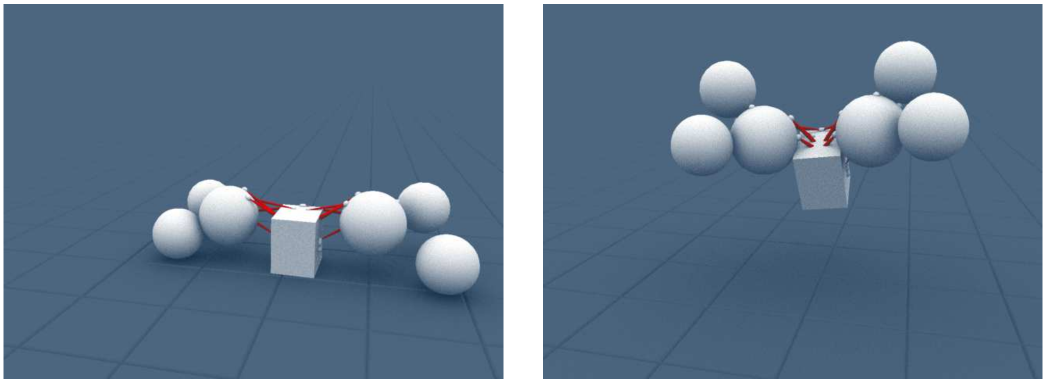

Figure 6.5: Two-armed swing (non-repeatable), from generation 360 of Run 12 of the jump task.

The creature in Figure 6.5 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. This result is from generation 360 of Run 12 (Figure 6.9) and achieved maximal fitness (1.000), which was the best among all results in this experiment.

6.4.4 Jump Result 3: One-Armed Swing

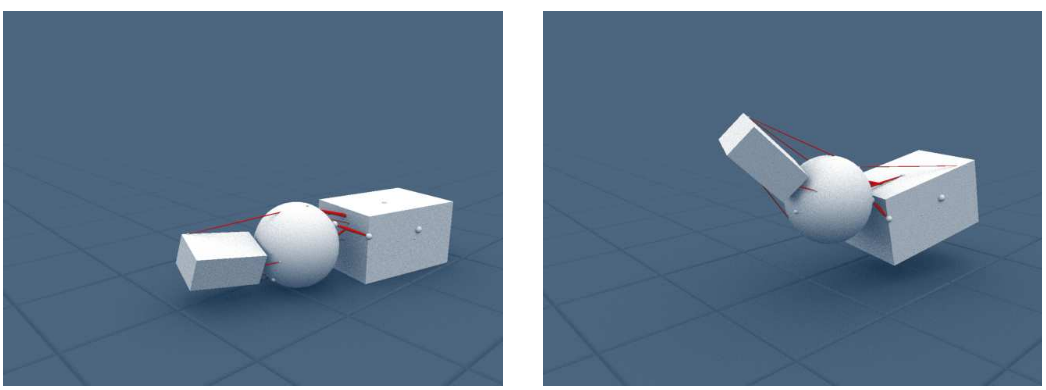

Figure 6.6: One-armed swing, from generation 320 of Run 19 of the jump task.

The strategy in Figure 6.6 is similar to that of the previous two, but it works with a single limb instead of a symmetrical pair of limbs. This result is from generation 320 of Run 19 (Figure 6.10) and achieved a fitness of 0.162.

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 (Section 6.5.3).

6.4.5 Jump Result 4: Four-Legged Push

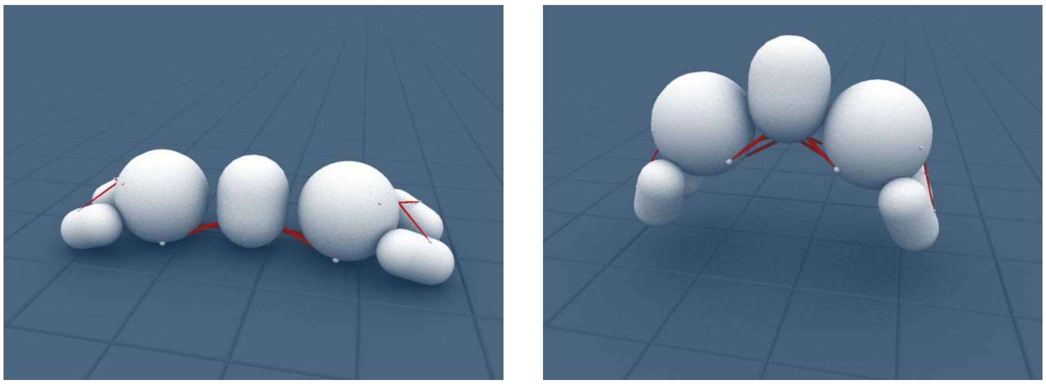

Figure 6.7: Four-legged push, from generation 221 of Run 25 of the jump task.

The creature in Figure 6.7 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. This result is from generation 221 of Run 25 (Figure 6.10) and achieved a fitness of 0.626.

6.4.6 Jump Result 5: Complex-Arm Swing

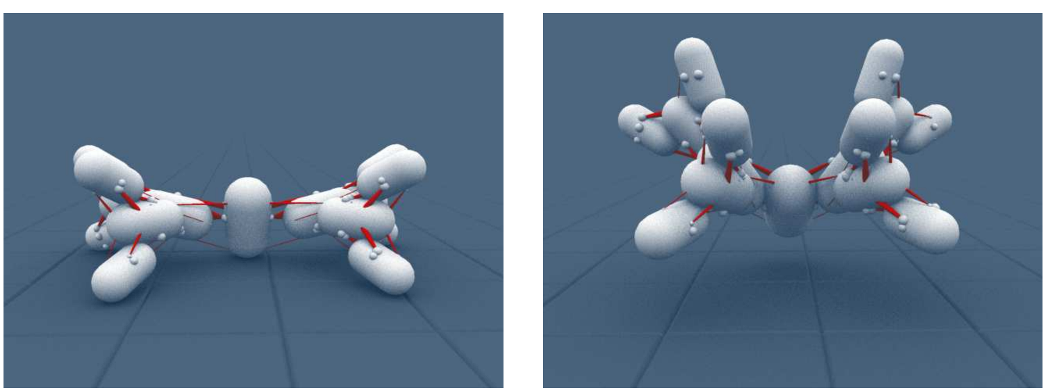

Figure 6.8: Complex-arm swing, from generation 500 of Run 35 of the jump task.

In Figure 6.8—the most morphologically elaborate of this chapter's 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 6.1 for a more detailed illustration.) This result is from generation 500 of Run 35 (Figure 6.11) and achieved a fitness of 0.346.

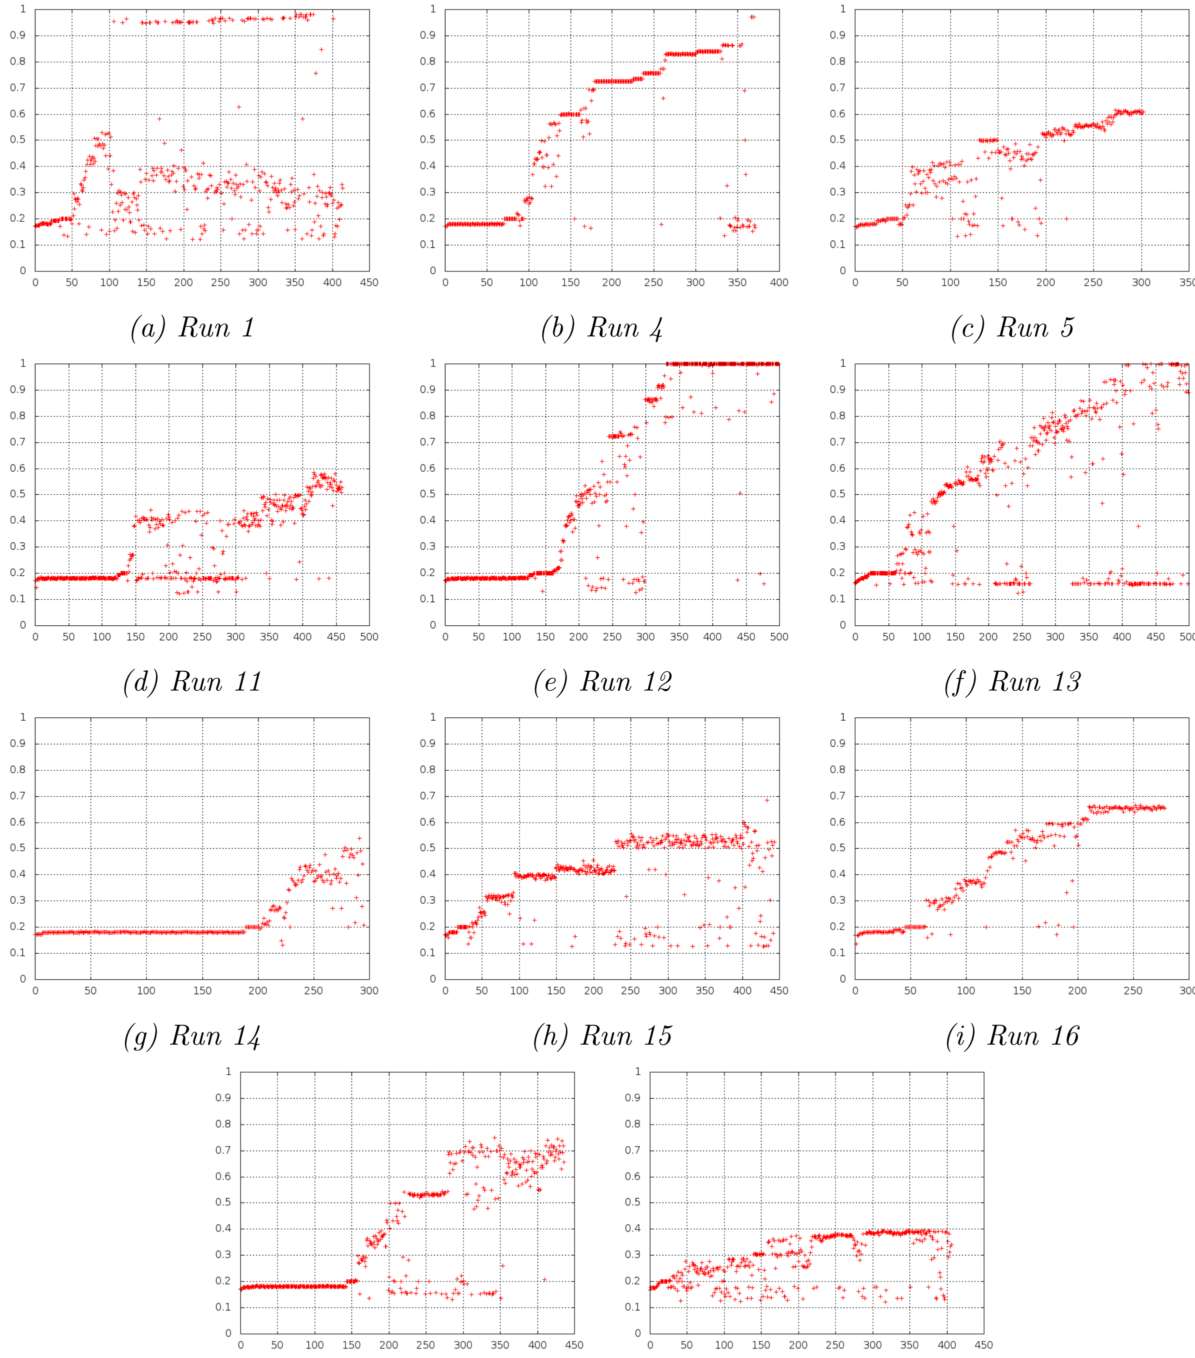

(a) Run 1 (b) Run 4 (c) Run 5

(d) Run 11 (e) Run 12 (f) Run 13

(g) Run 14 (h) Run 15 (i) Run 16

(j) Run 17 (k) Run 18

Figure 6.9: Fitness graphs for Run 1 through 18 of the jump task. (Note: The data from Runs 2-3 and 6-10 has been lost.) This set of runs contained the repeatable two-armed swing of Figure 6.4 and the non-repeatable two-armed swing of Figure 6.5.

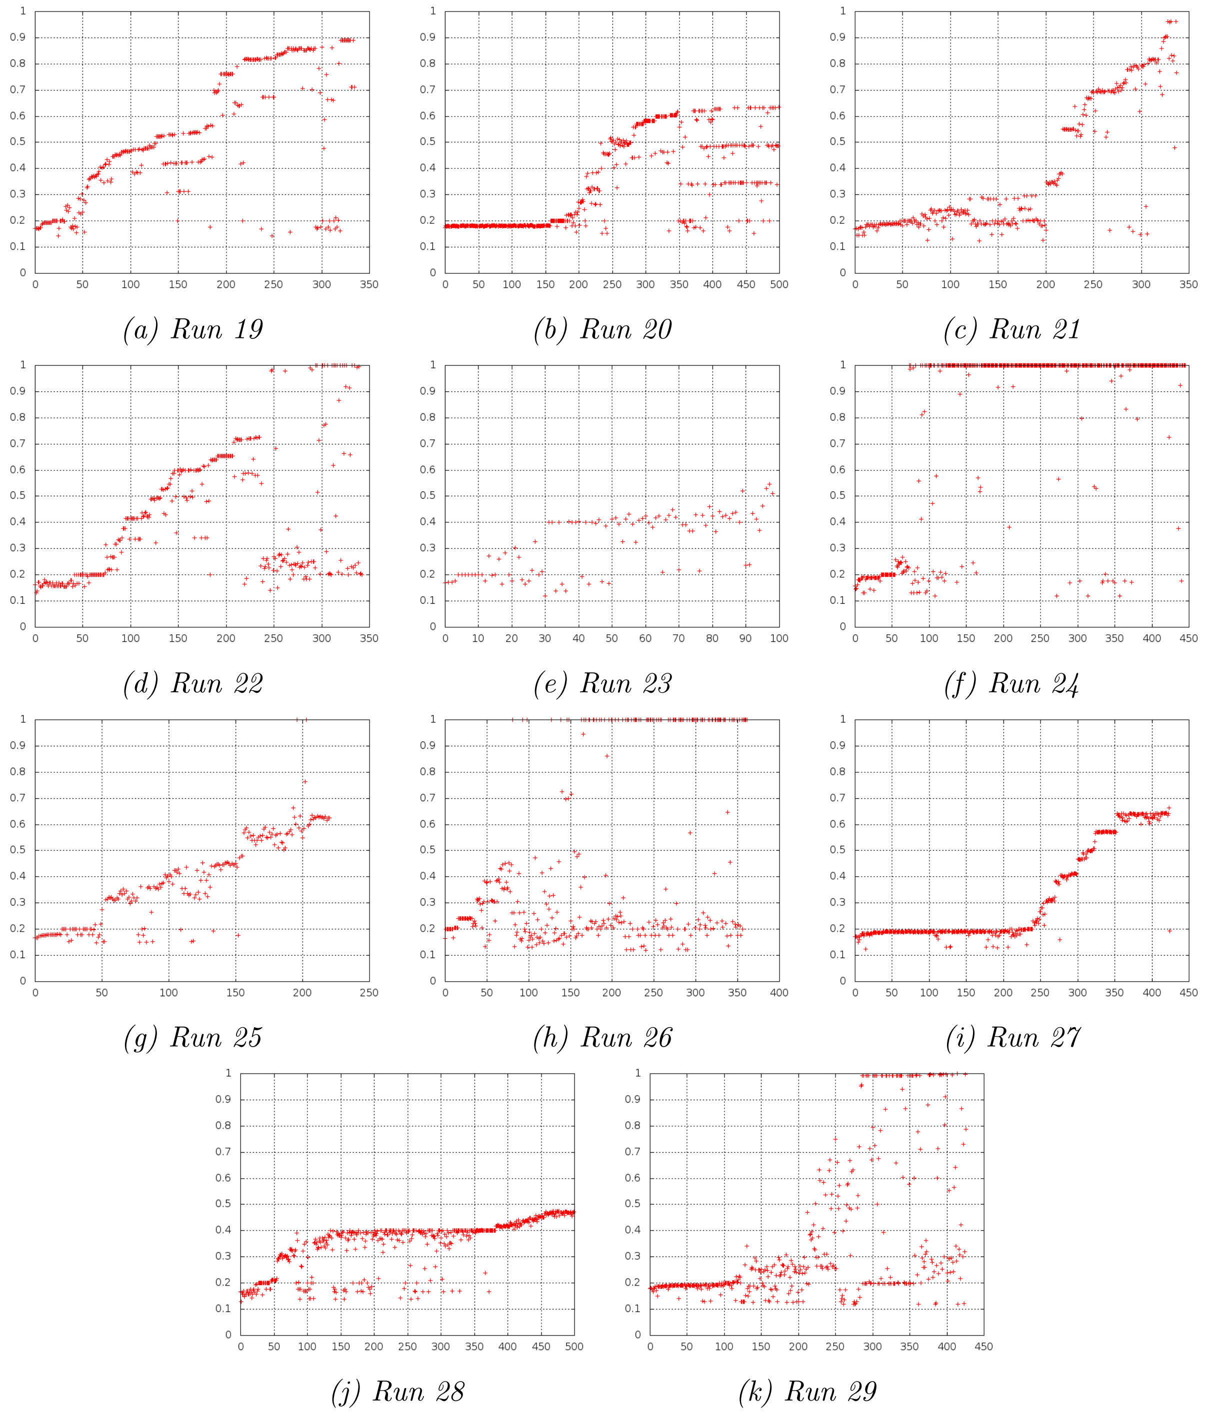

(a) Run 19 (b) Run 20 (c) Run 21
(d) Run 22 (e) Run 23 (f) Run 24
(g) Run 25 (h) Run 26 (i) Run 27
(j) Run 28 (k) Run 29

Figure 6.10: Fitness graphs for Run 19 through 29 of the jump task. This set of runs contained the one-armed swing of Figure 6.6 and the four-legged push of Figure 6.7.

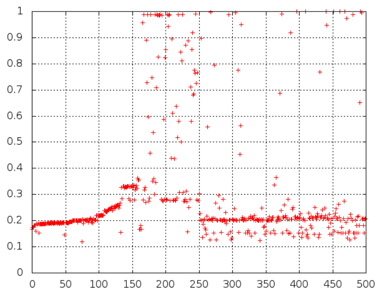

(a) *Run 30*

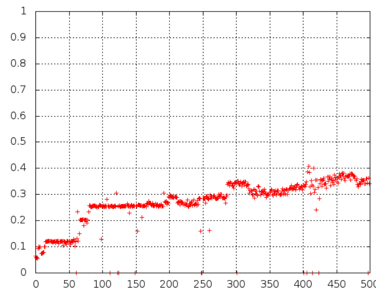

(b) *Run 31*

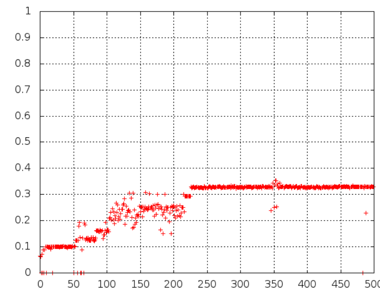

(c) *Run 32*

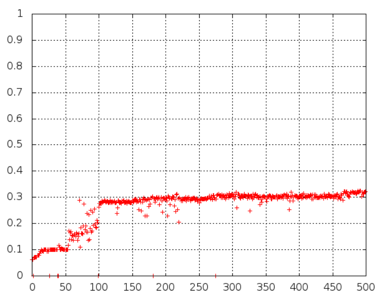

(d) *Run 33*

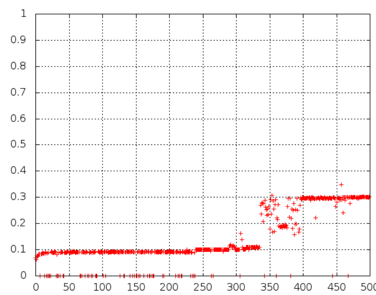

(e) *Run 34*

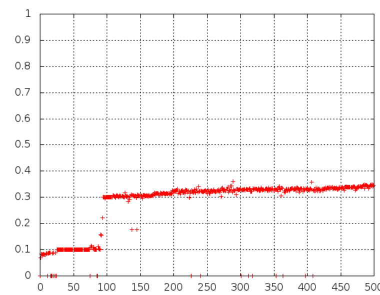

(f) *Run 35*

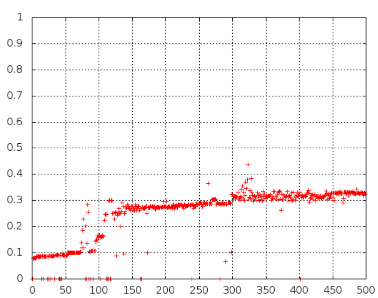

(g) *Run 36*

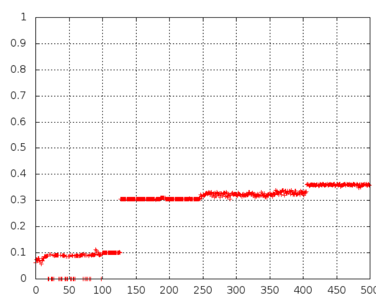

(h) *Run 37*

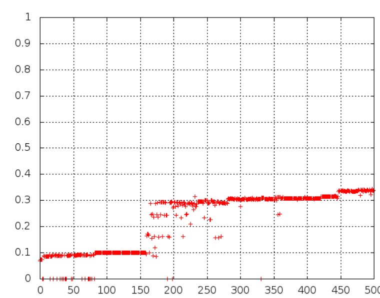

(i) *Run 38*

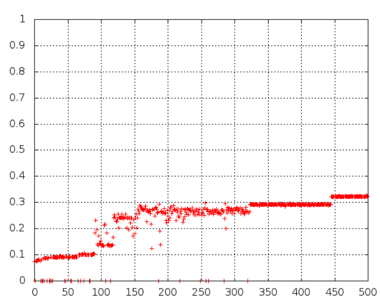

(j) *Run 39*

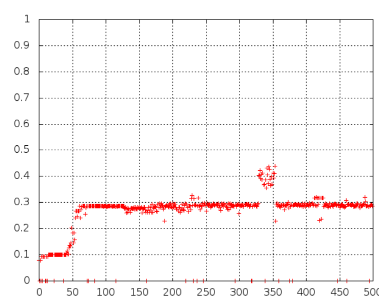

(k) *Run 40*

Figure 6.11: Fitness graphs for Run 30 through 40 of the jump task. This set of runs contained the complex-arm swing of Figure 6.8.

6.5 Locomotion Results

In this section, results from a locomotion experiment are presented, in which the single fixed square-pulse control signal of Figure 6.3(a) was replaced with the repeating fixed square-wave signal of Figure 6.3(b), and the ultimate fitness function was changed from jump height to distance traveled in a given amount of time.

6.5.1 Experimental Setup

In locomotion experiments, the population size was 100, and the results were obtained between 1000 and 2000 generations. As for the jump experiments, 40 independent runs were executed. For locomotion, the shaping schedule was more extensive: the first three goals are the same as in the jump task, and the last two utilize these results to construct locomotion. In the fourth step, in addition to ground clearance, a modest amount of horizontal travel is required. In the final stage, 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**
   For this step, full fitness is achieved by producing a creature so tall (at rest) that it cannot contain only one segment. In this way, it is ensured that creatures will contain joints, which will permit the addition of muscles in future steps.

2. **static AABB top + highest AABB top**
   In this step, half of the fitness (static AABB top) depends on maintaining sufficient static size to ensure that joints are retained. The other half of the fitness score (highest AABB top) is used to encourage the addition of muscles which produce upward motion.

3. **highest AABB top + highest AABB bottom**
   In this step, half of the fitness (highest AABB top) encourages the retention of upward-motion-producing muscles developed in the preceding step. In addition, the other half of the fitness score (highest AABB bottom) encourages getting the creature's lowest point as high off the ground as possible during the jump.

4. **highest AABB bottom + (modest) horizontal distance traveled**
   In this step, half of the fitness (highest AABB bottom) helps maintain the ground clearance achieved in the previous step. The other half is devoted to rewarding a small amount of horizontal travel. Thus, the previous step's jump is encouraged to become a jump with translation.

5. **horizontal distance traveled (unlimited)**
   In the final step, locomotion alone is encouraged, with fitness being defined simply as

horizontal distance traveled. The target distance for achieving full fitness in this step is so large as to be effectively unattainable, so that locomotion development can be pursued indefinitely.

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.

6.5.2 Locomotion Result 1: Double Front-Armed Swing Hop

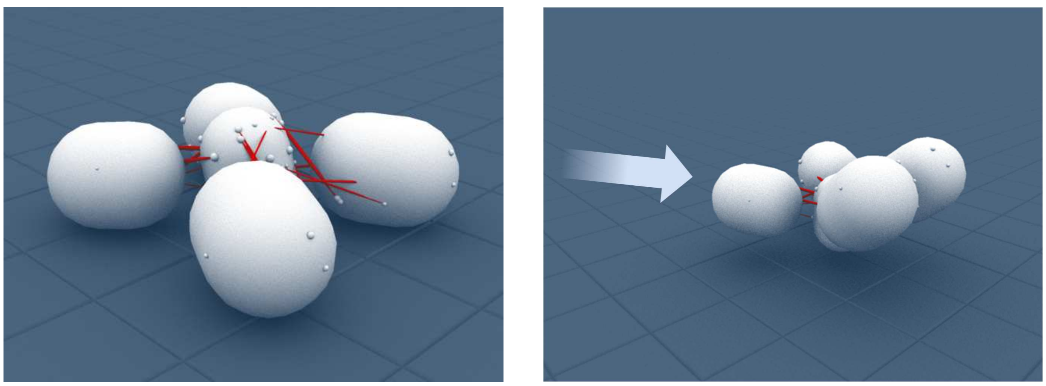

Figure 6.12: Double front-armed swing hop, from generation 1000 of Run 6 of the locomotion task.

In Figure 6.12, the square-wave activation of muscles is used to swing the front limbs up, accumulating momentum, which produces forward translation during repeated jumps. This result is from generation 1000 of Run 6 (Figure 6.20) and achieved a fitness of 0.017.

6.5.3 Locomotion Result 2: Single Front-Armed Swing Hop

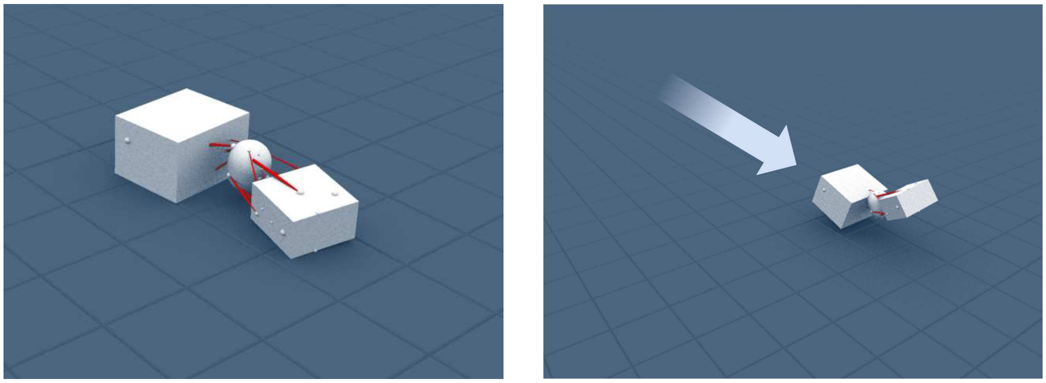

Figure 6.13: Single front-armed swing hop, from generation 1000 of Run 7 of the locomotion task.

With morphology and action very similar to that of Jump Result 3, the creature in Figure 6.13 also employs a repeating forward-translating jump for simple but highly effective locomotion. This result is from generation 1000 of Run 7 (Figure 6.20) and achieved a fitness of 0.038.

6.5.4 Locomotion Result 3: Front-Armed Swing Step

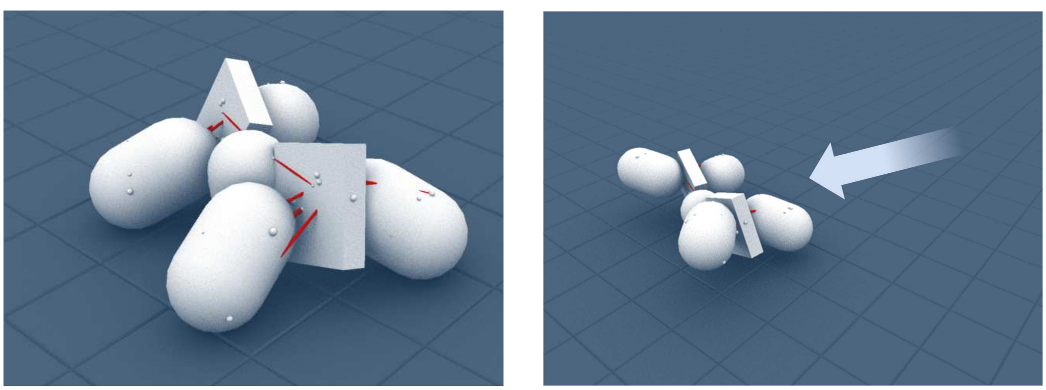

Figure 6.14: Front-armed swing step, from generation 1000 of Run 12 of the locomotion task.

The mode of locomotion of the creature in Figure 6.14 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.

This result is from generation 1000 of Run 12 (Figure 6.21) and achieved a final-stage fitness of 0.016.

6.5.5 Locomotion Result 4: Delta Wheelbarrow

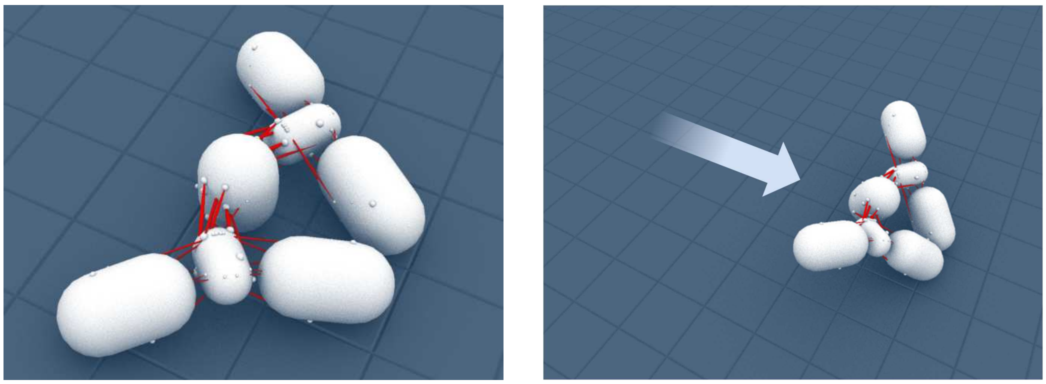

Figure 6.15: Delta wheelbarrow, from generation 2000 of Run 12 of the locomotion task.

The creature in Figure 6.15 employs a dense concentration of muscles at its central joints to produce upward and forward momentum, which results in a wheelbarrowing forward slide. This result is from generation 2000 of Run 12 (Figure 6.21) and achieved a final-stage fitness of 0.015.

6.5.6 Locomotion Result 5: Front-Hinged Swing Drag

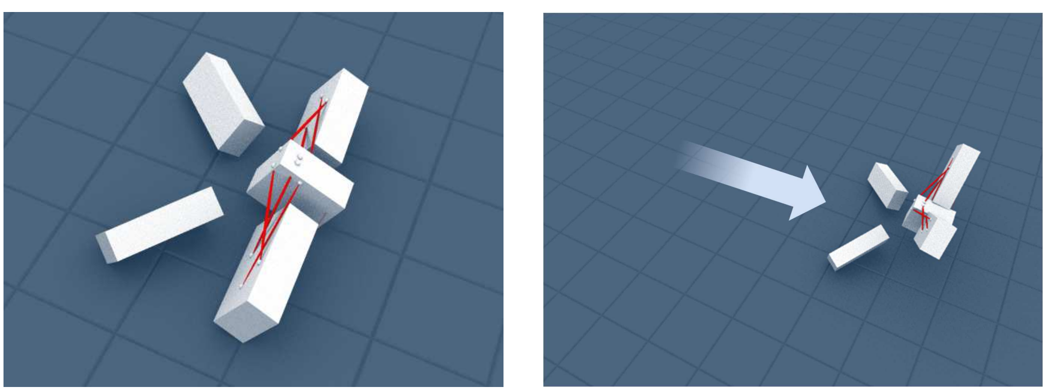

Figure 6.16: Front-hinged swing drag, from generation 2000 of Run 18 of the locomotion task.

In Figure 6.16, 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. This result is from generation 2000 of Run 18 (Figure 6.21) and achieved a fitness of 0.020.

6.5.7 Locomotion Result 6: Square Wheelbarrow

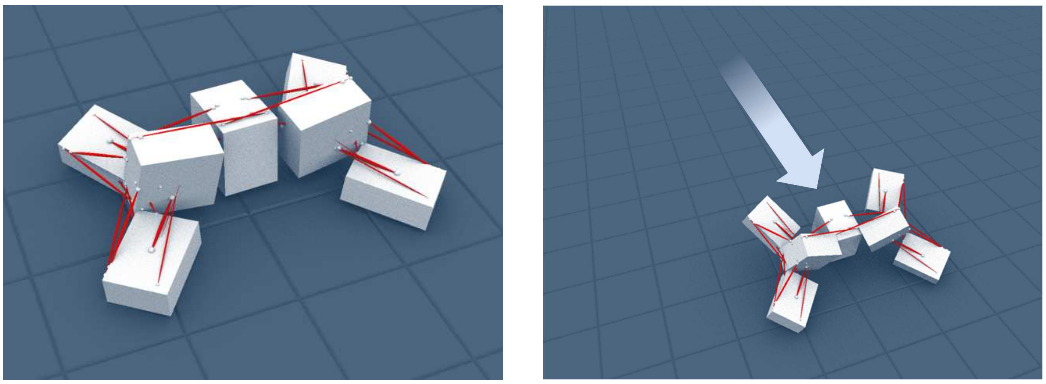

Figure 6.17: Square wheelbarrow, from generation 1250 of Run 20 of the locomotion task.

In Figure 6.17, a different morphology employs the same basic technique as Locomotion Result 4 to again produce a sliding wheelbarrow-like forward movement. This result is from generation 1250 of Run 20 (Figure 6.21) and achieved a fitness of 0.014.

6.5.8 Locomotion Result 7: Complex Swing Step

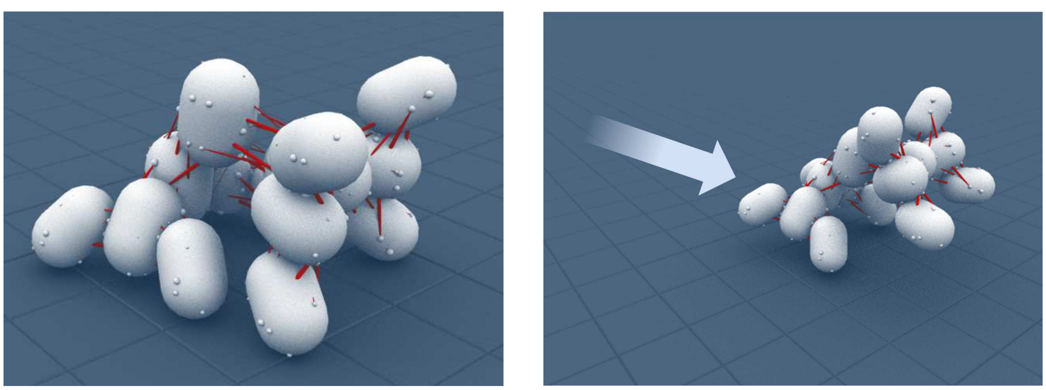

Figure 6.18: Complex swing step, from generation 1000 of Run 23 of the locomotion task.

In Figure 6.18, 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. This result is from generation 1000 of Run 23 (Figure 6.22) and achieved a fitness of 0.027.

6.5.9 Locomotion Result 8: High Hop

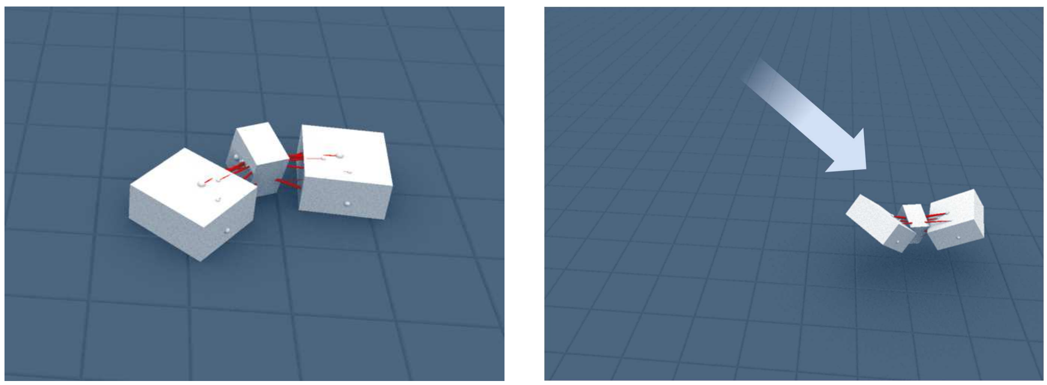

Figure 6.19: High hop, from generation 1000 of Run 31 of the locomotion task.

In one of the simplest yet most effective locomotion results (Figure 6.19), the creature uses clusters of strong muscles to swing up heavy limbs, lifting its comparatively small root segment in a high-jumping locomotive technique. This result is from generation 1000 of Run 31 (Figure 6.23) and achieved a fitness of 0.045.

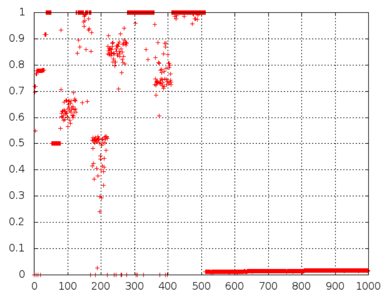

(a) Run 1

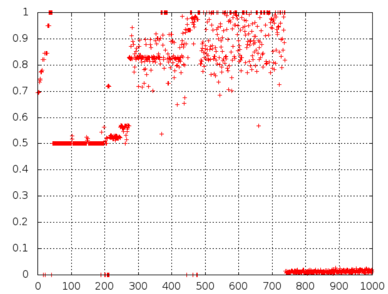

(b) Run 2

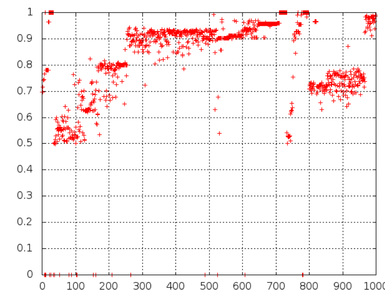

(c) Run 3

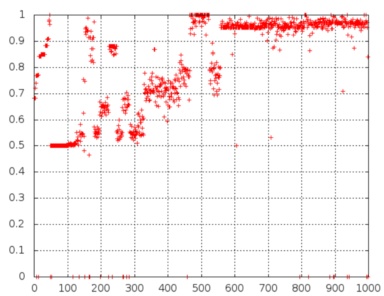

(d) Run 4

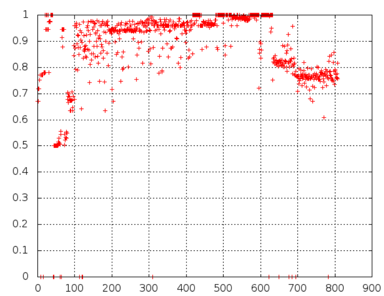

(e) Run 5

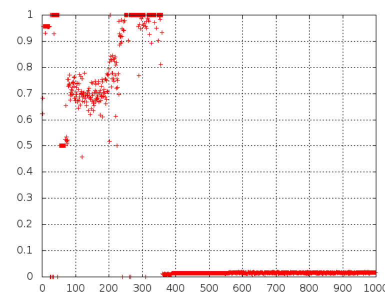

(f) Run 6

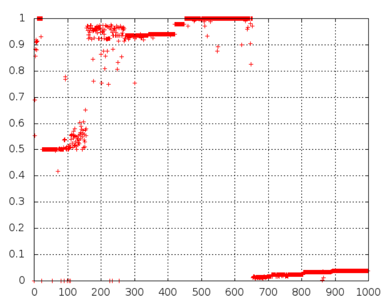

(g) Run 7

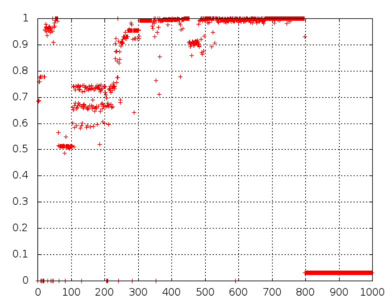

(h) Run 8

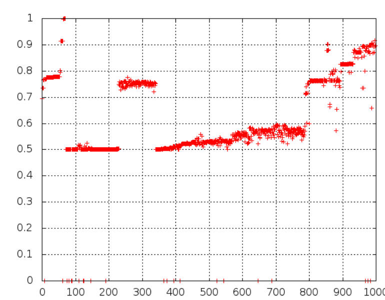

(i) Run 9

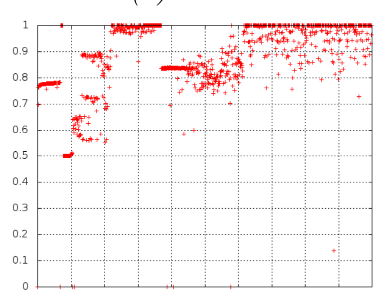

(j) Run 10

Figure 6.20: Fitness graphs for Run 1 through 10 of the locomotion task. This set of runs contained the double front-armed swing hop of Figure 6.12 and the single front-armed swing hop of Figure 6.13.

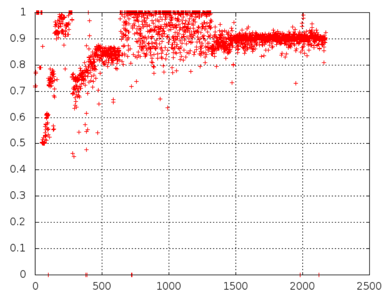

(a) *Run 11*

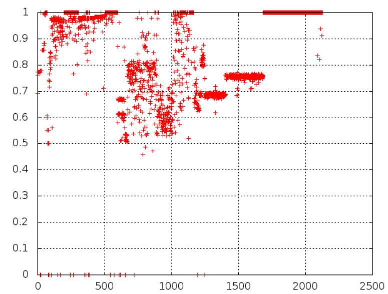

(b) *Run 12*

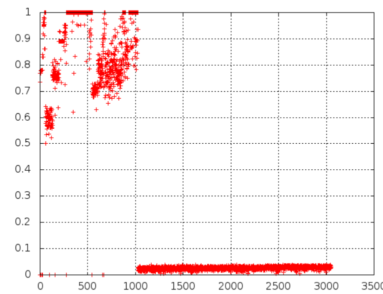

(c) *Run 13*

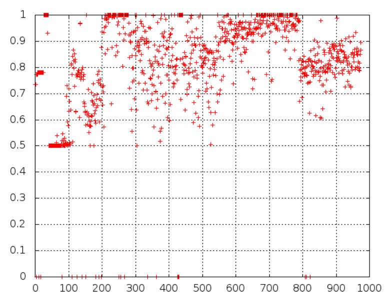

(d) *Run 14*

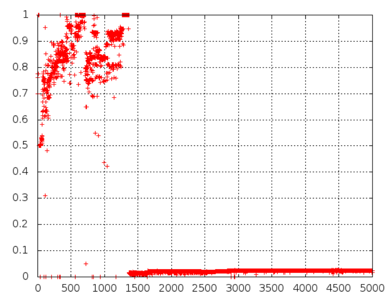

(e) *Run 15*

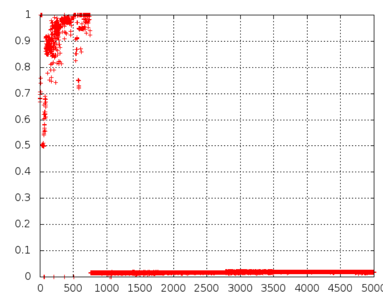

(f) *Run 16*

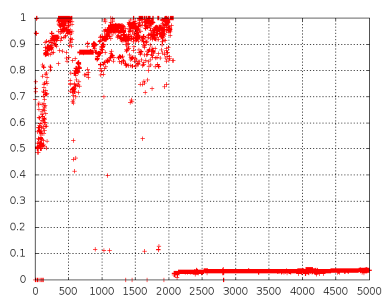

(g) *Run 17*

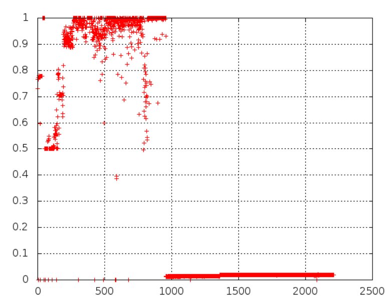

(h) *Run 18*

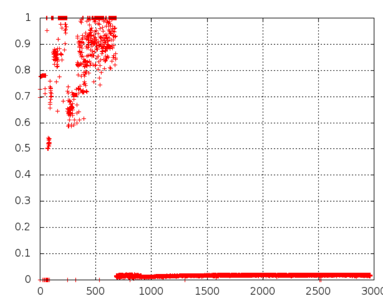

(i) *Run 19*

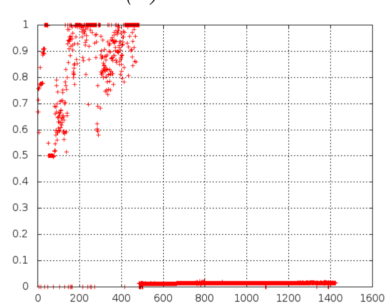

(j) *Run 20*

Figure 6.21: Fitness graphs for Run 11 through 20 of the locomotion task. This set of runs contained the front-armed swing step of Figure 6.14, the delta wheelbarrow of Figure 6.15, the front-hinged swing drag of Figure 6.16, and the square wheelbarrow of Figure 6.17.

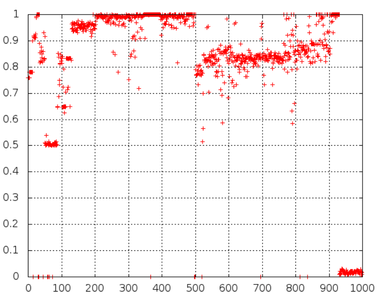

(a) Run 21

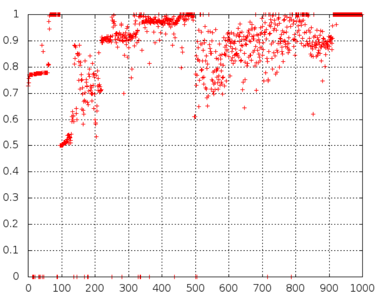

(b) Run 22

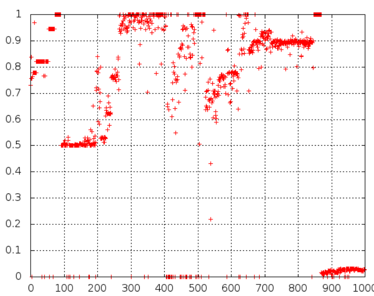

(c) Run 23

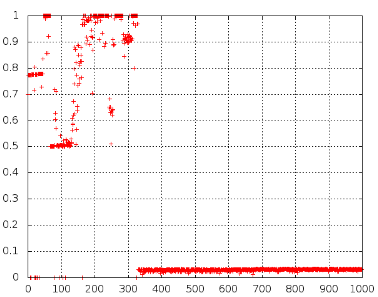

(d) Run 24

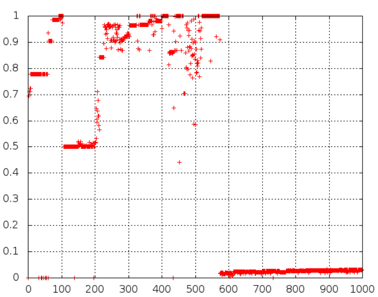

(e) Run 25

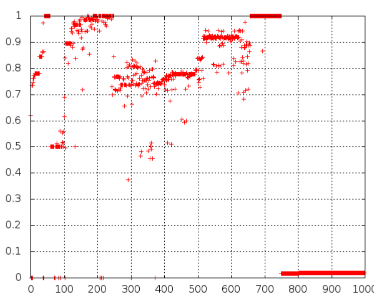

(f) Run 26

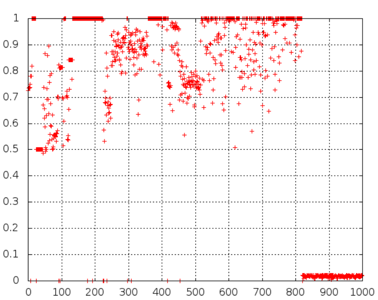

(g) Run 27

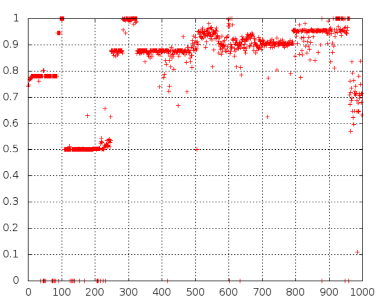

(h) Run 28

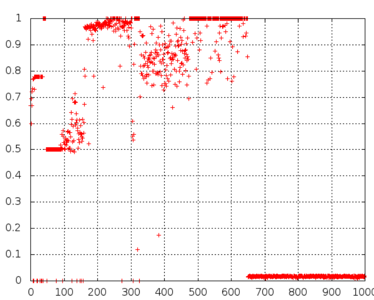

(i) Run 29

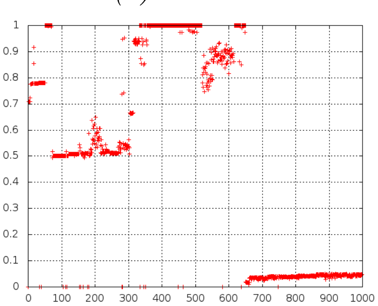

(j) Run 30

Figure 6.22: Fitness graphs for Run 21 through 30 of the locomotion task. This set of runs contained the complex swing step of Figure 6.18.

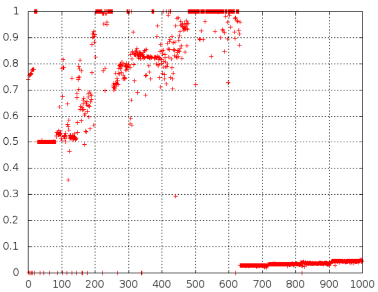

(a) Run 31

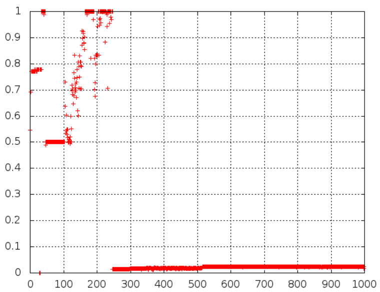

(b) Run 32

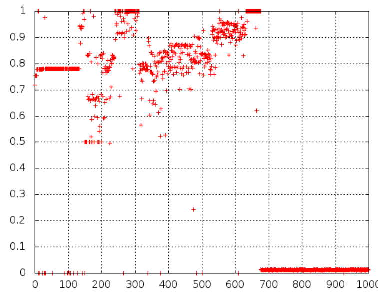

(c) Run 33

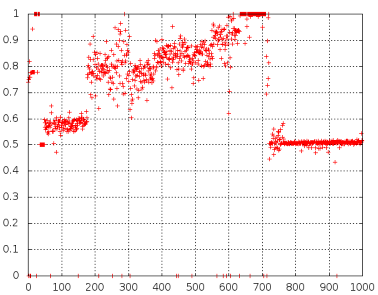

(d) Run 34

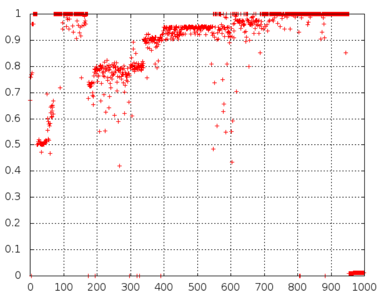

(e) Run 35

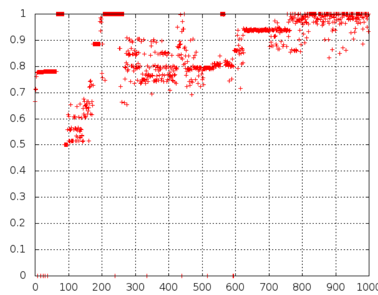

(f) Run 36

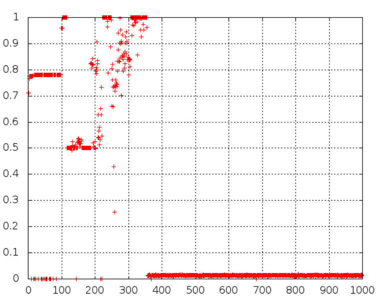

(g) Run 37

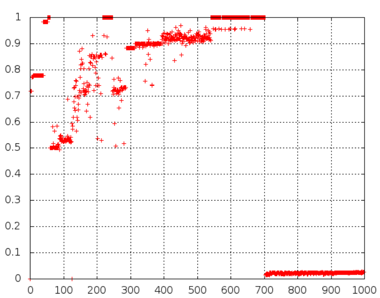

(h) Run 38

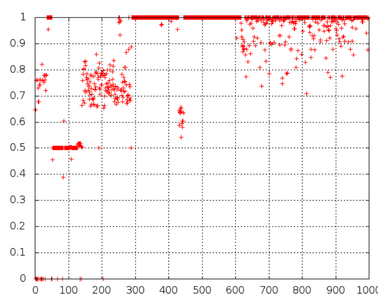

(i) Run 39

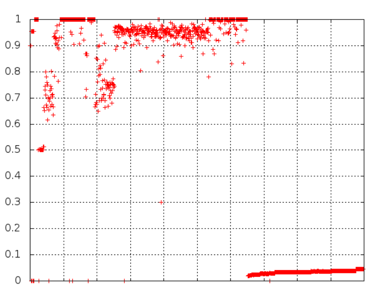

(j) Run 40

Figure 6.23: Fitness graphs for Run 31 through 40 of the locomotion task. This set of runs contained the high hop of Figure 6.19.

6.6 Discussion

The results in this chapter demonstrate the benefits of the evolved muscle drives: removing a measure of the burden from control intelligence and embodying that intelligence as functional morphological complexity. These benefits 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 in this chapter is intended to establish that 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, as will be discussed in Chapter 7.

6.7 Conclusion

This chapter described this dissertation's third contribution: A novel actuation system based on biologically inspired simulated muscles. This system of artificial-muscle drives produces both meaningful morphological complexity and reduces the required control intelligence, potentially making it possible to create more complex EVCs than before.

With the dissertation's three primary contributions now established, the following chapter describes potential directions for future work that build upon these accomplishments.

Chapter 7

Discussion and Future Work

This dissertation advances both behavioral and morphological complexity in evolved virtual creatures. In this chapter, remaining issues with these advances are discussed, namely a clarification of the differences between Original and Extended ESP, as well as the quantification of physical intelligence. Also, promising directions for future progress are presented in both of these areas, specifically: increasing morphological complexity, refining the contributed algorithms, and long-term applications.

7.1 Original ESP vs. Extended ESP

The original ESP method for evolving virtual creatures was presented in Chapter 4, and extended to continued morphology evolution in Chapter 5. It is natural to ask whether Extended ESP subsumes and obviates the original version. It does not: These two versions were presented individually not only because it was natural to describe them in succession, but also because they solve different problems. Extended ESP increases ESP's ability to adapt morphology to multiple tasks, but it does so at the cost of increasing evaluation times. When such morphological changes are not required, and the focus is on behavioral complexity, the original system would be preferred. Only when the morphological goals justify the increased evaluation costs should the extended system be employed.

7.2 Quantifying Physical Intelligence

Chapter 6 described the potential for muscle drives to embody a useful degree of control complexity that would otherwise be required in the brain. This shift in complexity from control to musculature was clear in the examples presented—the brains were essentially removed, but the behavior remained. The amount of control complexity transferred to the muscles was significant, i.e. it was sufficient for useful benchmark behaviors, such as jumping and locomotion. However, it is important to move beyond such qualitative descriptions to a quantitative evaluation of the degree of complexity transferred.

One way to obtain quantitative estimates might be with an experiment that compares creatures with muscle drives and no brains directly against creatures with traditional joint motors and traditional brains. First, the muscled, brainless creature would be evolved to some optimal level of ability at a skill, e.g. best jump height or fastest locomotion. Second, a traditional creature would be evolved for the same level of ability. Then, the traditional creature could be rewarded for minimizing brain complexity (perhaps measured by number of control nodes and wires in some way), while still maintaining its performance level. The resulting traditional creature's brain would thus demonstrate the minimal brain complexity required to match the muscle drives' embodied intelligence, i.e. the amount of brain capacity liberated by moving from traditional drives to muscle drives.

7.3 Biologically Inspired Morphological Complexity

While the evolvable musculature of Chapter 6 provided a significant increase in morphological complexity for EVCs, it was only the first step on a potentially long path of biologically inspired advances. This section suggests a sequence of such steps, taking EVC bodies from their current state to something closer to the fascinating complexity displayed by creatures in the natural world.

Muscles with Volume. One obvious next step would be to give the muscles volume by replacing the current system's simulated linear springs with simulated soft bodies or pressurized cloth. Previous work with simulated muscles [14] already demonstrated that such an approach is feasible, and specifically within PhysX. Allowing muscles to help define the distribution of the body's mass, as they do in real creatures, would advance biologically inspired purposeful 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.

Skin. Another biologically inspired refinement of the rigid-segment EVC model would be to simulate skin, as anticipated by Sims 20 years ago [46]. With powerful cloth simulation widely available (including in PhysX), this extension has become a conceivable next step in life-like morphological complexification. In particular, combining simulated-cloth skin with massed muscles could produce a particularly rich simulation, with a skin stretching and sliding over muscles as they extend and contract.

Emergent Joints. Further in 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 [2], 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.

Exoskeletons. Similarly, 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 would be expected to emerge.

7.4 Algorithmic Refinements

Although visually obvious improvements to morphology are appealing, the complex new algorithms introduced in this work (ESP in Chapter 4 and Extended ESP in Chapter 5) bring their own opportunities for further development, through more subtle recombination of skills, and a new path to morphological complexity.

Whole-Syllabus Adaptation. Although Extended ESP allows morphology evolution to continue through multiple skills, it still only applies to those skills that are *leaves* in the syllabus graph (Section 5.2.1). In principle, it would be desirable to continue morphological evolution throughout all skill adaptations. In this manner, it might be possible to develop morphologies that make transitions between skills easier—for example by making the creature more stable or more agile. To accomplish this goal in the current Extended ESP system would require considering all skills to be leaves, leading to $\mathcal{O}(n^2)$ growth in testing requirements applied to all n skills in the syllabus. To make such evolution possible without increasing testing excessively, a rolling horizon of leaf skills that travels through the syllabus hierarchy might be implemented. The idea is that once a lower ability will no longer be explicitly required for any subsequent skills, it need not be retested or maintained at all. With a well-chosen sequence of skill learning, an approximately constant-sized wave of leaves might result, sweeping gradually through the hierarchy.

Morphological Complexity from Multiple-Task Adaptation. While Extended ESP was employed in this dissertation to produce a variety of results for particular tasks (Chapter 5), it might also prove useful in the pursuit of morphological complexity as an end in itself (Section 6.1). Rather than merely allowing morphology to adapt to improve performance and variety in specific tasks (which may or may not require a more complicated body), a set of goal behaviors could be selected specifically for their ability to increase complexity. For example, the human hand exhibits great physical complexity; most likely this complexity emerged from the need to perform such a variety of tasks with the hand. The ability to embody multiple types of physical intelligence simultaneously through Extended ESP might similarly lead to a greater degree of physical complexity in EVCs.

7.5 Long-Term Applications

Beyond the near-term explorations described above, there are compelling long-term goals for this dissertation's research.

Combat. While Miconi already produced one limited form of combat for EVCs [30], there is a great deal more that can be done. The ESP method, in combination with a number of 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. This goal also presents some of the greatest challenges for increased complexity of morphology and behavior in EVCs.

Fauna on Demand. 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 diversity-promoting techniques such as those seen in [26]. 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.

Physical Intelligence for Nanotechnology. Physical intelligence through inhomogeneous drives (such as the muscle drives of Chapter 6) can significantly reduce the need for control logic, in some cases rendering it almost completely unnecessary. This reduction could ultimately prove useful 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.

Interactive ESP. Perhaps the highest expression of the ESP system would be an interactive version, in which direct human input completely replaces the explicit definition of fitness functions and the predefined ESP syllabus. This goal may present the greatest challenges, as well as offering some of the greatest rewards for the long-term future development of evolved virtual creatures.

7.6 Conclusion

This chapter discussed unresolved issues in this dissertation, namely the differences between Original and Extended ESP, and the quantification of physical intelligence. It also presented a number of promising avenues for continued development in morphological complexity, algorithmic development, and long-term applications. The contributions and main impact are reviewed in the next chapter.

Chapter 8

Conclusion

Evolved virtual creatures to be used as content in movies, video games, and virtual environments can benefit greatly from increased behavioral and morphological complexity. This dissertation presented three primary contributions toward those goals. This chapter summarizes those contributions and evaluates their potential future impact.

8.1 Contributions

The first contribution was the original version of the ESP system. As described in Chapter 4, it allowed evolved virtual creatures to achieve a level of behavioral complexity (as defined in the introduction) which approximately doubles the previous state of the art. The behavioral complexity of evolved virtual creatures has thus not yet been exhausted, and in fact it may continue to increase so as to one day approach the behavioral complexity of creatures from the real world—with all of the potential for content creation that this achievement might bring.

The second contribution was an extension of the original ESP system that makes it possible to continue adapting the morphology beyond the initial skill (Chapter 5), while still producing high-complexity behaviors incrementally. The benefits of this continued adaptation were demonstrated through experiments in which the extended ESP system generated a greater variety of solutions, and solutions with higher fitness. As discussed in Section 7.1, these are separate contributions because the original system is more efficient when focusing on complex behavior alone; the extended version does not obviate the original.

Third, the dissertation described a version of evolved virtual creatures in which traditional joint-motor drives are replaced by a simple yet powerful evolvable musculature (Chapter 6). 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 the body makes it visible, which constitutes progress toward meaningful morphological complexity. In addition, for these basic tasks, the typical EVC brain is made essentially superfluous. This result illustrates that these muscle drives can embody much of the normal control burden, liberating the brain's computational resources for other, more complex work.

8.2 Conclusion

This work has made clear contributions to increasing both the behavioral and morphological complexity of evolved virtual creatures, advancing the state of the art for EVCs as content. In the process, this dissertation has also opened a number of new avenues for rewarding future research in this domain. As a result, in both action and form, evolved virtual creatures are now a significant step closer to having the entertainment value of the real-world creatures that we all know and love.

Bibliography

[1] MIT Press Journals - Artificial Life, 2013.

[2] Joshua E. Auerbach and Josh 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.

[3] Joshua E Auerbach and Josh C Bongard. On the relationship between environmental and mechanical complexity in evolved robots. In *Artificial Life*, volume 13, pages 309–316, 2012.

[4] Joshua E. Auerbach and Joshua 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.

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

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

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

[8] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod 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.

[9] Jeff Clune and Hod Lipson. Evolving 3d objects with a generative encoding inspired by developmental biology. *ACM SIGEVOlution*, 5(4):2–12, 2011.

[10] Richard Dawkins. *The blind watchmaker: Why the evidence of evolution reveals a universe without design.* WW Norton & Company, 1996.

[11] J.A. Doucette, P. Lichodzijewski, and M.I. 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.

[12] Scott Draves. The electric sheep screen-saver: A case study in aesthetic evolution. In *Applications of Evolutionary Computing*, pages 458–467. Springer, 2005.

[13] Emden R. Gansner and Stephen C. North. An open graph visualization system and its applications to software engineering. *SOFTWARE - PRACTICE AND EXPERIENCE*, 30(11):1203–1233, 2000.

[14] 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.

[15] Faustino Gomez and Risto Miikkulainen. Incremental evolution of complex general behavior. *Adaptive Behavior*, 5:317–342, 1997.

[16] Javier Serrano González, Angel G Gonzalez Rodriguez, José Castro Mora, Jesús Riquelme Santos, and Manuel Burgos Payan. Optimization of wind farm turbines layout using an evolutive algorithm. *Renewable Energy*, 35(8):1671–1681, 2010.

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

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

[19] Gregory Hornby, J Lohn, and D Linden. Computer-automated evolution of an x-band antenna for nasa's space technology 5 mission. *Evolutionary computation*, 19(1):1–23, 2011.

[20] G.S. Hornby, J.B. 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.

[21] 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.

[22] MichałJoachimczak and Borys 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.

[23] Maciej 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.

[24] Peter Krcah. Evolving virtual creatures revisited. In *GECCO*, volume 7, page 341, 2007.

[25] 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.

[26] J. Lehman and K.O. 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.

[27] Dan Lessin, Don Fussell, and Risto 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.

[28] Dan Lessin, Don Fussell, and Risto Miikkulainen. Trading control intelligence for physical intelligence: Muscle drives in evolved virtual creatures. In *Proceeding of the Sixteenth Annual Conference on Genetic and Evolutionary Computation Conference*, GECCO '14, New York, NY, USA, 2014. ACM.

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

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

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

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

[33] Prasanth B Nair and Andrew J Keane. Passive vibration suppression of flexible space structures via optimal geometric redesign. *AIAA journal*, 39(7):1338–1346, 2001.

[34] Shigeru Obayashi, Daisuke Sasaki, Yukishiro Takeguchi, and Naoki Hirose. Multiobjective evolutionary computation for supersonic wing-shape optimization. *Evolutionary Computation, IEEE Transactions on*, 4(2):182–187, 2000.

[35] G.B. Peterson. A day of great illumination: B.F. Skinner's discovery of shaping. *Journal of the Experimental Analysis of Behavior*, 82(3):317, 2004.

[36] Rolf Pfeifer and Josh C. Bongard. *How the Body Shapes the Way We Think: A New View of Intelligence (Bradford Books)*. The MIT Press, 2006.

[37] Marcin L. Pilat and Christian 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.

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

[39] J. Schrum and R. Miikkulainen. Evolving multi-modal behavior in NPCs. In *Computational Intelligence and Games, 2009. CIG 2009. IEEE Symposium On*, pages 325–332. IEEE, 2009.

[40] Jimmy Secretan, Nicholas Beato, David B D Ambrosio, Adelein Rodriguez, Adam Campbell, and Kenneth O Stanley. Picbreeder: evolving pictures collaboratively online. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1759–1768. ACM, 2008.

[41] 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, November 1958. HMSO.

[42] Y.S. Shim and C.H. 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.

[43] Y.S. Shim, S.J. Kim, and C.H. 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.

[44] Karl Sims. *Artificial evolution for computer graphics*, volume 25. ACM, 1991.

[45] Karl Sims. Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372, 1994.

[46] Karl 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.

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

[48] Kenneth O. Stanley. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines*, 8(2):131–162, June 2007.

[49] 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.

[50] Joseba Urzelai, Dario Floreano, Marco Dorigo, and Marco Colombetti. Incremental robot shaping. *Connection Science*, 10:341–360, 1998.

[51] Shimon Whiteson, Nate Kohl, Risto Miikkulainen, and Peter Stone. Evolving keepaway soccer players through task decomposition. *Machine Learning*, 59(1):5–30, 2005.

[52] Ka-Chun Wong, Kwong-Sak Leung, and Man-Hon Wong. Protein structure prediction on a lattice model via multimodal optimization techniques. In *Proceedings of the 12th Annual Conference on Genetic and Evolutionary Computation*, GECCO '10, pages 155–162, New York, NY, USA, 2010. ACM.