Evolving a Diversity of Creatures through Novelty Search and Local Competition

Joel Lehman
Department of EECS
University of Central Florida
Orlando, FL 32816-2362 USA
jlehman@eecs.ucf.edu

Kenneth O. Stanley
Department of EECS
University of Central Florida
Orlando, Ohio 32816-2362 USA
kstanley@eecs.ucf.edu

ABSTRACT

An ambitious challenge in artificial life is to craft an evolutionary process that discovers a wide diversity of well-adapted virtual creatures within a single run. Unlike in nature, evolving creatures in virtual worlds tend to *converge* to a single morphology because selection therein greedily rewards the morphology that is easiest to exploit. However, *novelty search*, a technique that explicitly rewards *diverging*, can potentially mitigate such convergence. Thus in this paper an existing creature evolution platform is extended with multi-objective search that balances drives for both novelty and performance. However, there are different ways to combine performance-driven search and novelty search. The suggested approach is to provide evolution with both a novelty objective that encourages diverse morphologies and a local competition objective that rewards individuals outperforming those most similar in morphology. The results in an experiment evolving locomoting virtual creatures show that novelty search with local competition discovers more functional morphological diversity within a single run than models with global competition, which are more predisposed to converge. The conclusions are that novelty search with local competition may complement recent advances in evolving virtual creatures and may in general be a principled approach to combining novelty search with pressure to achieve.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning

General Terms

Algorithms

Keywords

Virtual creatures, Natural evolution, Artificial life, Novelty search

1. INTRODUCTION

Aspiring to understand and imitate nature, researchers have developed artificial environments that evolve virtual creatures [1, 2, 7, 9, 11, 18]. However, all such existing systems fall short of the impressive benchmark set by nature. An important question raised by this gap is what research directions will facilitate developing environments of virtual creatures with dynamics more like that of natural evolution.

A key factor in this problem is representation. Though the DNA encoding is ubiquitous in nature, virtual worlds crafted by researchers employ coarser, more abstract encodings that are computationally tractable. There are many approaches to representing a virtual creature, including L-systems [9], gene regulatory networks [2], hierarchical nested graphs [11, 18], and compositional pattern producing networks [1, 7].

The motivation for such research is that improved creature representations may be more evolvable and thereby pave a more forgiving path for evolution to follow towards more complex functional creatures. However, the full potential of even the best representation may remain unachieved if evolution is handicapped by a deficient system of rewards. Therefore, an important goal for creature evolution systems is to improve the evolutionary process that is at their core.

One shortcoming of traditional evolutionary algorithms (EAs) in this context is their general tendency, irrespective of representation, to *converge*. EAs do not naturally maintain diversity because most treat evolution as an optimizer; they generally march towards an optimum and then converge upon reaching it, stifling future innovation. Thus because creature evolution experiments are built upon traditional EAs, innovation in such investigations tends eventually to stagnate through convergence [2, 9, 15, 18].

In contrast to traditional EAs, a technique called *novelty search* explicitly rewards *diverging* from past designs [14]. Thus by applying such novelty search to discover novel morphologies it may be possible to combat the morphological convergence common in creature evolution experiments. While in previous novelty search experiments the measure of novelty was designed to correlate at least weakly with performance [13, 14, 17], morphology can be explored independently of locomotion ability, which is a typical measure of performance in virtual creature worlds [2, 9, 18]. However, a problem with exploring only morphological novelty is that arbitrary morphologies will almost always be non-functional. Thus in such situations it is necessary to combine novelty search with a complementary drive towards functionality.

Therefore, the experiment in this paper extends an existing creature evolution platform with a multi-objective search algorithm able to balance drives for novelty [14] and performance. However, the simple combination of novelty and performance objectives ignores the fact that different niches

GECCO'11, July 12–16, 2011, Dublin, Ireland.

have different *capacities* for fitness. For example, the morphology of a cheetah is predisposed to faster travel than that of an ant, although both locomote in interesting ways.

Thus the hypothesis in this paper is that discovering diverse solutions can be impeded by global competition between species; that is, in most creature evolution experiments, all organisms, no matter how varied, compete directly against one another on the same benchmark (e.g. how far the creature can travel) [1, 2, 7, 9, 11, 15].

However, in nature, grass does not preclude grass-hoppers, nor do bacteria compete directly with bears. Instead, competition is restricted and often localized within a niche or between connected niches. Such restricted competition implies an implicit reward for discovering new niches, whereby an organism can *escape* competition. Because of this implicit reward, nature is prone to *diverge* and continuously uncover new ways of living, unlike most EAs.

Thus the suggested approach is to provide evolution with both a novelty objective that encourages diverging from previous morphologies and a *local competition* objective that restricts competition locally among individuals with similar morphologies. This local competition contrasts with the *global competition* objective standard in traditional EAs in which each individual competes with all others on the same measure of progress. The result from such local competition is more functional diversity within a population; interesting behaviors are discovered by local competition in niches that are effectively ignored when competition is global.

However, there is naturally a trade-off: The most fit individuals found by local competition are less fit than those found by global competition. Unlike in natural evolution, the population size in virtual worlds is nearly always fixed, which necessarily implies that exploring new niches reallocates resources away from exploiting existing niches. In this paper, this key trade-off is demonstrated through an experiment evolving creatures that are able to locomote effectively in a three-dimensional physically realistic simulation.

The trade-off illuminates a subtle facet of what is interesting about artificial life. We argue that finding the singular most optimally fit organism is not as interesting as uncovering a wide variety of creatures well-adapted to their niches. Though ultimately such interestingness is subjective, it is important at least to be aware that global optimality is not the only desirable goal of an evolutionary system.

The conclusion is that novelty search with local competition, inspired by the limited interactions between niches in nature, can help create more natural creature evolution through encouraging both morphological diversity and performance localized within morphological niches. This advance may complement recent advances in creature representation as well as provide a good general approach to combining the search for novelty with the pressure to achieve.

2. BACKGROUND

This section reviews previous attempts to evolve virtual creatures including the Evolving Robotic Organisms (ERO; [11]) system, and novelty search, an evolutionary technique with which ERO is extended in the experiment in this paper.

2.1 Evolving Virtual Creatures

The most well-known example of evolving virtual creatures is the seminal work of Sims [18]. In this work, an EA evolves the morphology and control policy of virtual creatures specified by an encoding able to represent hierarchical, symmetric, and repeating structures. Creatures are simulated in a realistic three-dimensional environment and learn to locomote in different ways. The results have since been replicated [11, 15]. A problem with Sims' system is that a particular run *converges* to a single morphology [18], contrasting with nature's accumulation of diverse organisms.

Of the systems that have been designed since, most focus on alternate ways of representing organisms, e.g. L-systems [9] or genetic regulatory networks [2]. A recent more granular approach to representation is applying compositional pattern producing networks (CPPNs [20]) to specify growth patterns of three-dimensional structures [1] or density patterns in three-dimensional space for soft robotics [7]. Such systems can in theory scale indefinitely to represent creatures with a large quantity of morphological parts. Though equipped with more advanced representations, these systems also tend to converge to single morphologies [2, 9]. That is, convergence results more from selection than representation.

The particular system extended for the experiments in this paper is called ERO [11], a modern re-implementation of Sims [18] with extensions. Although newer encodings may provide greater evolvability or potential, ERO is extensively tested and therefore provides a good experimental platform for the local competition approach, which can in principle be generally applied to any creature representation to mitigate morphological convergence. Thus the next section briefly reviews ERO (for a full description see Krčah [11]).

2.2 Evolution of Robotic Organisms (ERO)

Krčah [11] describes the ERO system, applies ERO to recreate the experiments of Sims [18], and tests extensions designed to encourage diversity and enhance crossover. ERO applies an EA to a complex encoding derived from that in Sims [18] with extensions inspired by NEAT [19]; the encoding specifies both the morphology and control policy of a virtual creature. Fitness evaluation takes place in a realistic three-dimensional simulation where creatures are rewarded for their ability to effectively locomote.

ERO's encoding unrolls an evolved genotypic graph structure into a coupled body plan and control policy; this unrolling process yields hierarchical and symmetric repeating structures. In ERO, an organism's genome is represented as a nested graph structure that has outer nodes and inner nodes (figure 1). The outer nodes and the connections among them represent different morphological parts comprising a creature's body, while the inner nodes and the connections among them represent the artificial neural network (ANN) that will control the creature.

Each outer node represents a body part of a creature, and each directed edge between two outer nodes represents a physical joint. Connections can be recursively applied to generate hierarchical structures, and can also be mirrored over different axes to generate symmetrical structures. The idea is to provide evolution with a set of useful regularities similar to those seen in natural evolution, which also exploits symmetry and hierarchical structures.

ERO also encodes the control policy of a creature in the nested graph. Each inner node represents a node in an ANN, and each weighted directed edge between two inner nodes is a connection in the ANN. As the genotypic graph is unrolled into the creature, the ANN represented by the inner nodes and connections is unrolled as well, generating modu-

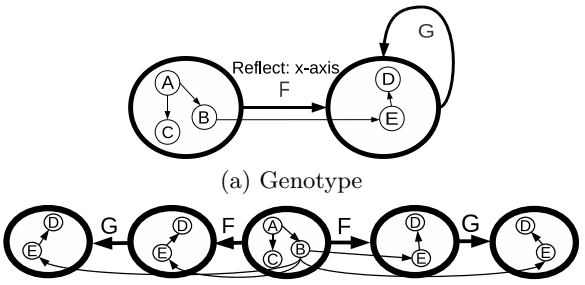

(a) Genotype

(b) Phenotype Blueprint

Figure 1: **ERO Encoding.** A nested graph genotype is shown in (a) that consists of outer nodes (i.e. the larger circles) that represent morphological parts, connections among outer nodes (e.g. F and G) that represent joints, inner nodes (i.e. the smaller circles) that represent neurons in an ANN, and neural connections among inner nodes (e.g. the connection between B and E). The phenotype shown in (b) is unrolled from the genotype (a) according to flags associated with the connections between outer nodes that can encode symmetry (e.g. the joint F and its children are reflected over the x-axis) and repetition (e.g. joints can be repeated to create articulated appendages like fingers), yielding a hierarchical repeating structure.

lar ANNs that repeat in the same way as the morphological nodes they are nested within. ANN nodes can represent sensory organs of the creature, motors driving the joints of the creature, or intermediate functions of sensory signals.

By applying an EA to this representation, the experiments by Krčah [11] evolve creatures that locomote in various ways (walking, swimming, and jumping) and that can reactively follow a light. Even though ERO extends Sims [18] with genotypic speciation (possible in such a graph encoding by exploiting historical markings as in Stanley and Miikkulainen [19]), which does increase performance, preliminary experimentation revealed that such diversity maintenance is typically ineffective at preventing convergence to a single morphology by the end of an evolutionary run, like other related systems [2, 9, 15, 18].

Because it derives from a well-tested and validated model of virtual creatures, ERO is an ideal system to test the hypothesis advanced in this paper: Morphologically local competition combined with a drive for novelty will more effectively explore and exploit a diversity of morphological niches than more traditional approaches.

The next section reviews novelty search, an evolutionary method that accumulates diversity, which is combined with local competition in the experiment in this paper.

2.3 Novelty Search

EAs are often applied as black-box optimization algorithms designed to converge to a globally optimal fitness. In contrast, natural evolution *diverges*, creating and maintaining a wide variety of solutions to the problems of life. *Novelty search* is a *divergent* evolutionary technique, inspired by natural evolution's drive to novelty, that directly rewards novel behavior *instead* of progress towards a fixed objective [12, 14]. Thus it matches well with artificial life domains that are not motivated by a defined set of objectives.

Tracking novelty requires little change to any evolutionary algorithm aside from replacing the fitness function with a *novelty metric*, which measures how different an individual is from other individuals, thereby creating a constant pressure to do something new. The key idea is that instead of rewarding performance on an objective, novelty search rewards diverging from prior behaviors. Therefore, novelty needs to be *measured*.

The novelty of a newly generated individual is computed with respect to the *behaviors* of an *archive* of past individuals and the current population. The archive is initially empty, and individuals' behaviors are added to it probabilistically to penalize future individuals that exhibit previously explored behaviors [13].

The novelty metric is designed to characterize how far away the new individual is from the rest of the population and its predecessors in *behavior space*, i.e. the space of unique behaviors. A good metric should thus compute the *sparseness* at any point in the behavior space. Areas with denser clusters of visited points are less novel and therefore rewarded less.

A simple measure of sparseness at a point is the average distance to the k-nearest neighbors of that point, where k is a fixed parameter that is determined experimentally. Intuitively, if the average distance to a given point's nearest neighbors is large then it is in a sparse area; it is in a dense region if the average distance is small. The sparseness ρ at point x is given by

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} \text{dist}(x, \mu_i), \quad (1)$$

where μ_i is the ith-nearest neighbor of x with respect to the distance metric *dist*, which is a domain-dependent measure of behavioral difference between two individuals in the search space. The nearest neighbors calculation must take into consideration individuals from the current population and from the permanent archive of novel individuals. Candidates from more sparse regions of this behavioral search space then receive higher novelty scores. It is important to note that this novelty space cannot be explored purposefully, that is, it is not known *a priori* how to enter areas of low density just as it is not known a priori how to construct a solution close to the objective. Thus, moving through the space of novel behaviors requires exploration.

The current generation plus the archive give a comprehensive sample of where the search has been and where it currently is; that way, by attempting to maximize the novelty metric, the gradient of search is simply towards what is *new*, with no other explicit objective.

Once objective-based fitness is replaced with novelty, the underlying evolutionary algorithm operates as normal, selecting the most novel individuals to reproduce. Over generations, the population spreads out across the space of possible behaviors, sometimes encountering an individual that solves a given problem even though progress towards the solution is not directly rewarded. In fact, there have been several successful applications of novelty search in neuroevolution and genetic programming [6, 13, 14, 16, 17].

Instead of rewarding novel *behaviors* as in prior novelty search experiments, in this paper novelty search explores a space of *morphologies*, in effect rewarding novel morphologies that exhibit characteristics different from those previously encountered.

3. APPROACH

One reason virtual creature experiments are interesting is the morphological and functional variety that evolution might discover in such experiments. However, in most existing systems, uncovering such diversity becomes almost pro-

hibitively computationally expensive; though *different* runs may yield varied solutions, a particular run will usually converge to a single morphology [2, 9, 15, 18]. This dynamic also departs strikingly from natural evolution, which concurrently explores an increasing diversity of morphologies.

To remedy such morphological convergence, the idea in this paper is to apply novelty search to reward *novel morphologies* by defining a space of morphological niches and a novelty metric that quantifies how different a new morphology is from previous morphologies in this space. For example, a two-dimensional morphology space can be constructed by considering the height and mass of a virtual creature. A search for novelty within this space will effectively explore the space of morphologies spanning those that are short and light to tall and heavy.

However, a raw search for morphological novelty does not guarantee that any of the discovered morphologies are *functional*. That is, rewarding pure morphological novelty does not address the control policy of the virtual creature; all of the creatures discovered might not move at all. Such disconnect between morphology and functionality differs from prior experiments with novelty search in which the space it explores also distinguishes *behaviors* of interest [14, 17]. Thus what is needed is an extension to novelty search that encourages functionality when exploring dimensions like morphology that themselves do not distinguish functionality.

The next section describes one approach to balancing competing drives for diversity and functionality through a Pareto-based Multi-objective Evolutionary Algorithm (MOEA).

3.1 Balancing Achievement with Novelty

Multi-objective optimization is a popular paradigm within EC that addresses how to optimize more than one objective at the same time in a principled way [3]. Such multi-objective search suggests a simple way to combine the drive to optimize performance with the search for novelty: Reward both performance and novelty at the same time by making them separate objectives in a multi-objective search [16]. Yet while it is tempting simply to combine these different objectives and novelty through an MOEA, it is important to understand what this approach might inadvertently imply.

The concept of non-dominance is central to Pareto-based multi-objective search; the key insight is that when comparing the performance of two individuals over multiple objectives, if both individuals are better on different subsets of the objectives then there is no meaningful way to directly rank such individuals because neither entirely *dominates* the other. That is, ranking such mutually non-dominating individuals would require placing priority or weight on one objective at the cost of another; traditionally one individual dominates another only if it is no worse than the other over all objectives and better than the other individual on at least one objective.

Thus the best individuals in a population are those that are not dominated by any others. Such best individuals form the *non-dominated front*, which defines a series of trade-offs in the objective space. That is, the non-dominated front contains individuals that specialize in various combinations of optimizing the set of all objectives; some will maximize one at the expense of all the rest, while some may focus equally on all of the objectives.

Therefore if novelty and fitness are simply two separate objectives, then at one extreme of the resulting non-dominated front, maximizing novelty at the expense of performance will be rewarded. At the other extreme of the front, maximizing performance at the expense of novelty will be rewarded. In between these two extremes will be various trade-offs between performance and novelty. In effect, this type of search is a mixture of traditional objective-based search and novelty search; this simple combination can be seen as an objective-based search that encourages novelty to maintain behavioral diversity.

While this option is viable for search and may sometimes be more effective at achieving higher performance than searching only for novelty or only for functionality, it fails to exploit the fact that some niches may naturally support different levels of fitness than others. For example, bacteria reproduce dramatically faster than humans, and thus technically have higher fitness, though both are viable ways of life in nature. Sometimes an interesting goal is not narrowly to achieve the *highest* fitness, but rather to discover a wide diversity of individuals well-adapted to *their own niches*, like the diversity seen in natural evolution.

In nature, the success of one particular niche does not affect the success of another unrelated niche; that is, bears do not generally suffer from the reproductive success of bacteria. In contrast, in EC a global trade-off between fitness and novelty implies that the success of one niche may render other niches unappealing. For example, in the context of a virtual world, a medium-size, medium-mass organism may be more optimal for traveling quickly than a tall, low-mass organism; a global trade-off between fitness (i.e. locomotion speed) and novelty means that most resources will be spent on the medium morphology even though the way that the tall morphology would eventually perform at its best relative to nearby niches may ultimately be just as interesting. Thus a Pareto front trading off fitness and novelty does not really capture the spirit of the diverse set of locally-optimal designs we might hope to emerge in an artificial world.

3.2 Local Competition

The suggested response is to limit competition among organisms *locally* within a niche-space (i.e. not *physical* locality); that is, individuals compete only with other individuals nearby in niche-space. For example, in an artificial world a space of niches can be constructed by considering the space of possible morphologies. Local competition within such a space implies that individuals compete for fitness only with other individuals in similar morphological niches.

The key change is that local competition within a niche space transforms the fitness objective from being a global measure to one relative to an individual's neighbors in niche space. In effect, the search pressure changes from managing a global trade-off between fitness and novelty that is biased towards particular niches to encouraging both novelty and better performance relative to an individual's niche.

The idea is thus to explore the merits of *each niche* rather than to exploit greedily only the best niches. This strategy reflects that what is most interesting about some domains may not always be discovering the most optimal behavior. That is, among the various strategies for achieving the highest performance score in a given domain, the factors that decide which is ultimately optimal may only be an incidental side effect of how the domain is constructed. For example, one particular artificial world may have a physics engine in which quadruped locomotion is optimal, and another may

have an engine in which worm-like locomotion is optimal. Thus the particular optimality of one strategy or another may be uninteresting; instead, it may be more interesting to observe the creative force of evolution exploiting a wide variety of niches than it is to see only the optimal product of evolution narrowly converging.

It is important to note the differences between local competition and existing diversity maintenance techniques in EC [5, 8, 10]. Some such techniques do restrict competition, e.g. by level of fitness [10] or by age of genotype [8], but not in a way that necessarily encourages finding and maintaining *functionally* distinct solutions, i.e. genotypic age and fitness level have no inherent link to an individual's *means* of functionality (e.g. its particular strategy or morphology). Other methods modulate fitness based on genotypic diversity [5] but generally do not change the *scale* of fitness itself; though rare genotypes may have a better chance of surviving, such methods do not acknowledge that different niches may have widely differing distributions of fitness values that they are capable of supporting (e.g. bacterial niches support absolute fitness values of several orders of magnitude higher than mammalian niches).

In practice, transforming a global competition score (e.g. the fitness function) into a local competition score requires a comparison of an individual's performance to that of its nearest neighbors in niche space. The more neighbors it outperforms, the higher its local competition score.

Therefore, it is straightforward to implement local competition as an extension of novelty search, which already calculates nearest neighbors in morphology space. When determining the novelty of an individual by measuring the average distance to its nearest neighbors in morphology space, the number of such nearest neighbors with lower fitness than that of the individual is also counted. This number is assigned as the *local competition objective* for that individual, which measures an individual's performance *relative* to its niche. Thus two objectives in the MOEA become novelty and local competitiveness. The expectation is that balancing achievement and novelty in this way will lead to more natural evolutionary dynamics through a gradual accumulation of functionally-diverse well-adapted individuals.

4. EXPERIMENT

The experiment in this paper extends ERO with Pareto multi-objective search, novelty search, and local competition. It is important to note that although this composite system has several components, the central underlying intuition motivating the system is simple: Encouraging both morphological novelty and local competition may lead to a greater diversity of functional creatures.

The MOEA with which ERO is extended is based on NSGA-II, a well-known MOEA [4]. One difference is that while NSGA-II has a mechanism to encourage diversity along the non-dominated front, the implementation described here *replaces* that mechanism with a separate objective explicitly rewarding *genotypic* diversity. This departure is logical because both novelty and local competition are relative measures; two individuals with exactly the same novelty or local competition scores that would be conflated by encouraging diversity on the Pareto front may actually be quite different morphologically or with respect to performance.

The characterization of morphology space has important consequences for evolution and strongly biases the types of creatures evolution will discover. In this paper the morphological space that novelty search explores and within which competition is localized has three dimensions: height, mass, and number of active joints. This space encourages creatures of different sizes, with different distributions of mass, and that exploit different means of locomotion.

In all experimental setups, the fitness function rewards virtual creatures according to how far they are able to travel within an evaluation. The novelty metric is the squared Euclidean distance separating two individuals in morphology space, and thus the novelty of a creature is proportional to how different its morphology is from that of other creatures currently in the population.

The genotypic diversity objective encourages elaborating upon existing genotypes by assigning higher scores to more novel genotypes. A rough characterization of a genotype in ERO is the number of outer genotypic nodes it has; a genome starts with a small number of outer genotypic nodes, more of which can be added through mutation as evolution progresses, and the amount of which limit the maximal morphological complexity a genotype can represent. Because adding a new outer genotypic node often initially disrupts existing functionality, it is beneficial if such innovation is temporarily protected. Thus genotypic diversity is rewarded inversely proportionally to how many other genotypes in the population have the same amount of outer genotypic nodes.

Four experimental setups are tested: (1) **fitness alone**, (2) **novelty alone**, (3) **novelty search with global fitness competition**, and (4) **novelty search with local fitness competition**. The only difference between these setups is the choice of objectives given to the MOEA. Genotypic diversity is an objective in all variants.

The first setup, fitness alone, has a global fitness objective and a genotypic diversity objective. This setup is a control, roughly corresponding to the original ERO experiment, which employed fitness-based search coupled with speciation to encourage genotypic diversity.

The second setup, novelty alone, has a novelty objective and a genotypic diversity objective. This setup is a second control designed to show that without additional pressure to adapt, a search for morphological novelty alone will not generate creatures able to effectively locomote.

The third setup, novelty with global competition, has a novelty objective, a global fitness objective, and a genotypic diversity objective. This setup is much like the combination of novelty and fitness in Mouret [16]. The idea is to extend standard objective-based search with an additional drive for novelty that may encourage diversity, thereby reaping some of the benefit of novelty search while mostly maintaining the familiar objective-based search paradigm.

The final setup, novelty with local competition, abstracts competition between niches in nature; instead of a global fitness objective between all morphologies, evolution is provided a local competition objective that restricts competition within the morphology space. That is, an individual receives a local competition score relative to how many of its nearest neighbors in morphological space it outperforms.

Additionally, in the final setup the genotypic diversity objective is also localized within the morphology space; similar in motivation to that of local competition, local diversity ensures that genotypic diversity is not only exploited in those morphological niches in which such diversity is incidentally most easily expressed.

In all setups, the distribution of individuals in morphology space is recorded, as well as the final generation of the individuals themselves and the progression of highest fitness individuals throughout the run. The idea is to quantify how well each morphological niche is being exploited and how much morphological diversity is discovered and *maintained* throughout a run.

4.1 Experimental Parameters

The population size for all experiments is $1,000$, and a run consists of 500 generations. The nearest-neighbor size for novelty search and local competition is 15. The three morphology dimensions are rescaled so that their values fill the range between zero and three; the height dimension originally spans between zero and 1.5 meters, the mass dimension originally spans between zero and 3.0 kilograms, and the active joint dimension originally spans between one and four active joints. The crossover rate was 75% and mutated parameters would be perturbed with power 0.25. The selection method was changed to tournament selection (with tournament size two) to agree with NSGA-II [4]. Other parameters are the same as in Krčah [11].

5. RESULTS

Three metrics are described next that respectively quantify the ability of each setup to discover high fitness individuals, to maintain morphological diversity, and to exploit morphological niches. For each of the metrics, the Kruskal-Wallis test was first applied across the four experimental setups to demonstrate that the distributions are significantly different ($p < 0.001$). The remainder of the statistical tests below are Mann-Whitney U tests measuring the significance of pair-wise differences between experimental setups.

The first metric is the maximum fitness discovered by a particular setup. These results are shown in figure 2. The worst-performing setup for this metric is morphological novelty alone because it enforces no selection pressure towards functionality (of course, another kind of novelty, such as behavioral, might have produced better fitness, as Lehman and Stanley [14] have shown, but would not produce the desired morphological diversity). The global competition setup, which heavily favors morphological niches that most easily facilitate locomotion, significantly outperforms the other setups ($p < 0.001$). Fitness alone and local competition do not perform significantly differently from each other, but perform better than novelty alone ($p < 0.001$).

The second metric is niche coverage, which quantifies how well a setup has encouraged and *maintained* morphological diversity until the end of a run. The morphology space is overlaid with a regularly-spaced grid. From each point in the grid, the distance in morphology space to the nearest individual in the final population is calculated. The more well-covered the niche space is, the less the sum of all such nearest distances will be; if the population has spread throughout niche space, then on average there will be an individual near to each point in the grid. The results are shown in figure 3. Note that the larger the niche coverage metric is, the *worse* the morphological space is covered. Thus the worst-performing setup is fitness alone, which has no direct selection pressure towards discovering morphological diversity beyond that useful in maximizing fitness ($p < 0.001$). The best performing setups are novelty alone and local competition ($p < 0.001$), which do not perform significantly differently from each other. Local competition outperforms global competition on this metric because it allocates resources more fairly among morphological niches.

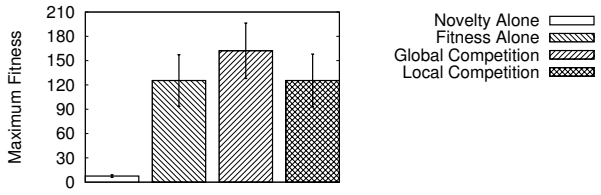

Figure 2: **Absolute Performance Comparison (larger is better).** For each setup, the maximum fitness discovered in a particular run is shown (averaged over 15 runs). The main result is that novelty search with global competition discovers the most fit individuals ($p < 0.001$).

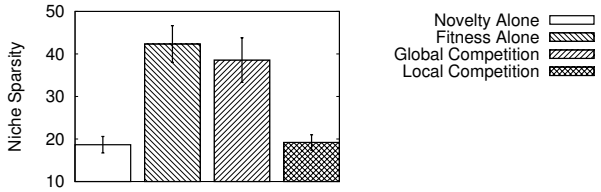

Figure 3: **Niche Sparsity (larger is worse).** For each setup, the average coverage of morphology space of the final population of a particular run is shown (lower is better; averaged over 15 runs). The main result is that novelty with local competition and novelty search alone cover the niche space the best ($p < 0.001$).

The last metric is niche exploitation, which quantifies how well on average a setup has accumulated and maintained diverse individuals that exploit niches. The morphological space is first divided into equal-sized bins. Next, each individual in the final population is placed into the bin that matches its morphology. Then, for each bin the best fitness among all individuals in it is recorded; these represent the best solutions for each niche that evolution was able to find. If all niches supported the same level of fitness, then taking the average of these best fitnesses would be a reasonable measure of niche exploitation. However, figure 4, which illustrates niche capacities, demonstrates the uneven distribution of solutions discovered across morphology space. Intuitively, some morphological niches are simply better suited to locomotion than others; e.g. it is difficult to construct an effective, very tall, very light creature.

Because of the uneven distribution of maximum fitnesses across niches, which span several orders of magnitude, fitness values for each niche are normalized first logarithmically, then are divided by the *best* fitness values for each niche (also normalized logarithmically) over *all* runs; i.e. each niche is normalized by its capacity for fitness. Finally, these normalized fitness values are then averaged over all of the niches, yielding the average niche exploitation; this number indicates for a particular run the average percentage of a niche's capacity evolution was able to exploit.

The results are shown in figure 5. This figure illustrates the trade-off between localized and globalized competition: While global competition is better able to find the maximum fitness, local competition is better able to exploit all of the morphological niches.

6. TYPICAL RUN

Typical runs of novelty alone accumulate diverse morphologies, but none are functional (figure 6a). In contrast,

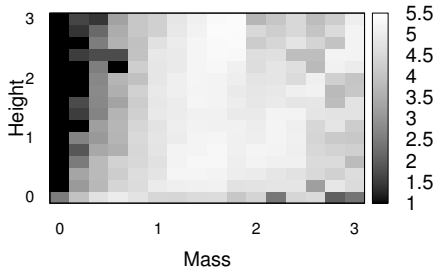

Figure 4: **Niche Capacity.** The capacity of evolution to exploit different morphological niches is illustrated above. Each square represents a segment of morphology space (only two out of three morphological dimensions are visualized), and its darkness is proportional to the logarithm of the highest fitness value found within that segment of morphology space (i.e. darker means more fit), over all runs of all variants; thus it is an estimate of the upper bound of fitness that the niche supports.

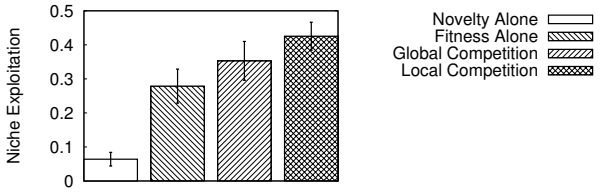

Figure 5: **Niche Exploitation (larger is better).** For each setup, the average niche exploitation is shown (averaged over 15 runs). The main result is that novelty with local competition exploits niches significantly better on average than the other setups ($p < 0.001$).

typical runs of fitness alone (figure 6b) and novelty with global competition (figure 6c) generally converge to a single body plan. However, global competition *stretches* that body plan further through morphology space than fitness alone because of novelty pressure. Finally, local competition (figure 6d) typically finds diverse functional morphologies, often maintaining worm-like, unipedal, and quadrupedal locomotion throughout a run, all at once; such diversity would otherwise require many independent runs [2, 9, 11, 15, 18].

Images of typical such creatures discovered within a single run are shown in figure 7, and videos of typical ones discovered within a single run are available at:
`http://eplex.cs.ucf.edu/ns/vc.html`

7. DISCUSSION

The results support the hypothesis in this paper that local competition combined with novelty search can better maintain and exploit a diversity of individuals than can global competition between all individuals combined with novelty, albeit at a cost in absolute performance. This result is potentially important given the ubiquity of global competition objectives throughout EC.

While this paper combines drives towards novelty and achievement, prior successes with searching for novelty alone [14, 17] illustrate that such a combination is not always necessary. The poor relative performance of searching for novel morphologies alone in this paper is not surprising nor does it contradict such prior successes. That is, of course morphological novelty does not distinguish competent creatures as another type of novelty might (e.g. behavioral novelty). Rather, the idea in this paper is that sometimes it is most natural to express and encourage the desired outcome of evolution (e.g. a morphologically diverse collection of competent creatures) as a combination of loosely connected drives.

In the context of artificial life, it is interesting to further consider the relationship between local competition and niches in nature. Morphologically localized competition is an abstraction of interacting niches in nature. However, the abstraction is inexact: In nature, niches are not aligned in fixed-dimensional spaces, and interactions between niches, although *restricted*, may not always be *localized*. For example, a large predator may impinge on the fitness of small prey, or a small parasite may make a larger animal ill.

That is, natural niches are not fundamentally functions of morphology but of ways of life, which are defined implicitly and tend to grow through a process of bootstrapping. One way of life may provide the substrate for many others. For example, the stomach of an animal may be a novel niche, and the same animal's waste products can found yet another niche. An interesting future research direction is to construct a similar system wherein niches and interactions are implicitly defined and thereby lead to indefinite open-ended discovery of new niches.

It is important to note that lack of exact fidelity to nature does not imply that local competition may not often be a useful tool in encouraging functional diversity whenever a niche space can defined. In fact, this formulation may in general be a good way of merging novelty search with the drive to optimize an objective. It may be particularly important in the virtual worlds of artificial life wherein the hope is to encourage an explosion of diverse and competent (though not necessarily globally optimal) creatures.

The idea of local competition maintaining a diversity of interesting functional individuals aligns well with the dynamics of natural evolution; additionally, the results presented here may provide anecdotal evidence of a problem in how evolutionary algorithms are nearly always applied: Evolution may ultimately prove better suited to progressively uncovering a widening interconnected web of interesting, complex artifacts than it is to the constrained pursuit of a single optimal individual.

8. CONCLUSIONS

This paper proposed a new approach to evolving virtual creatures that is independent of and potentially complementary to different ways of representing creatures. The experiments demonstrated that evolutionary algorithms may often create and maintain a wider diversity of functional organisms through a process of local competition combined with novelty search. The conclusion is that such local competition may be a viable way in general of combining pressure to achieve with the drive to uncover and preserve novelty.

Acknowledgements

This research was supported by DARPA grant HR0011-09-1-0045.

References

[1] J.E. Auerbach and J.C. Bongard. Dynamic Resolution in the Co-Evolution of Morphology and Control. In *Proceedings of the Twelfth International Conference on Artificial Life (ALIFE XII)*, 2010.

[2] J.C. Bongard and R. Pfeifer. Evolving complete agents using artificial ontogeny. *Morpho-functional machines: The new species (designing embodied intelligence)*, pages 237–258, 2003.

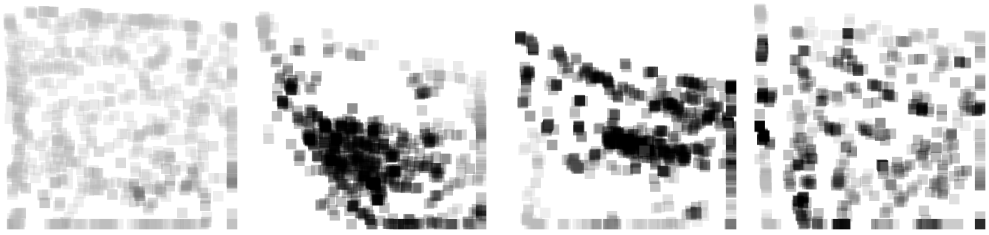

(a) Novelty Only (b) Fitness Only (c) Global Competition (d) Local Competition

Figure 6: **Comparing exploration and exploitation of morphological niches.** The final distribution of virtual creatures from a single run in two morphological dimensions (height and mass) is shown for typical runs of (a) novelty only, (b) fitness only, (c) novelty search with global competition, and (d) novelty search with local competition. The darkness of each point indicates how fit each virtual creature is. The main result is that local competition effectively exploits more of the morphological niche space than the other methods.

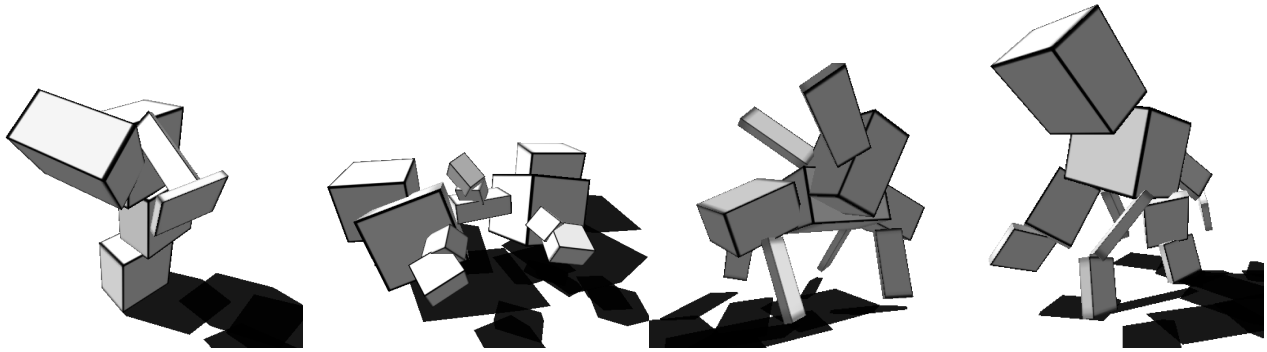

(a) Hopper (b) Crab (c) Quadruped (d) Tailed Quadruped

Figure 7: **Diverse competent morphologies discovered within a typical *single* run of local competition.** Various creatures are shown that have specialized to effectively exploit particular niches of morphology space. These creatures were all found in the final population of a typical run of local competition. The hopper (a) is a unipedal hopper that is very tall, (b) is a heavy short crab-like creature, and (c) and (d) are distinct quadrupeds. Creature (c) drives a large protrusion on its back to generate momentum, and (d) has a tail for balance.

[3] C.A.C. Coello. A comprehensive survey of evolutionary-based multiobjective optimization techniques. *Knowledge and Information systems*, 1(3):129–156, 1999.

[4] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

[5] David E. Goldberg and Jon Richardson. Genetic algorithms with sharing for multimodal function optimization. In *Proceedings of the Second International Conference on Genetic Algorithms*, pages 41–49, Hillsdale, NJ, USA, 1987. L. Erlbaum Associates Inc.

[6] H.J. Goldsby and B.H.C. Cheng. Automatically Discovering Properties that Specify the Latent Behavior of UML Models. In *Proceedings of MODELS 2010*, 2010.

[7] J.D. Hiller and H. Lipson. Evolving Amorphous Robots. In *Proceedings of the Twelfth International Conference on Artificial Life (ALIFE XII)*, 2010.

[8] Gregory S. Hornby. ALPS: the age-layered population structure for reducing the problem of premature convergence. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2006)*, pages 815–822, New York, NY, USA, 2006. ACM.

[9] G.S. Hornby, H. Lipson, and J.B. Pollack. Evolution of generative design systems for modular physical robots. In *IEEE Intl. Conf. on Robotics and Automation*, volume 4, pages 4146–4151. IEEE; 1999, 2001.

[10] Jianjun Hu, Erik Goodman, Kisung Seo, Zhun Fan, and Rondal Rosenberg. The hierarchical fair competition (HFC) framework for sustainable evolutionary algorithms. *Evol. Comp.*, 13(2):241–277, 2005.

[11] P. Křcah. Evolutionary development of robotic organisms. *Master's thesis, Charles University in Prague*, 2007.

[12] Joel Lehman and Kenneth O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *Pro. of the Eleventh Intl. Conf. on Art. Life (ALIFE XI)*, Cambridge, MA, 2008. MIT Press.

[13] Joel Lehman and Kenneth O. Stanley. Efficiently evolving programs through the search for novelty. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2010)*. ACM, 2010.

[14] Joel Lehman and Kenneth O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 2011. To appear.

[15] Thomas Miconi. *The Road To Everywhere: Evolution, Complexity and Progress in Nature and in Computers*. PhD thesis, Univ. of Birmingham, 2007.

[16] Jean-Baptiste Mouret. Novelty-based multiobjectivization. In *Proc. of the Workshop on Exploring New Horizons in Evol. Design of Robots,2009 IEEE/RSJ Intl. Conf. on Intell. Robots and Systems*, 2009.

[17] S. Risi, C.E. Hughes, and K.O. Stanley. Evolving plastic neural networks with novelty search. *Adaptive Behavior*, 2010.

[18] K. Sims. Evolving virtual creatures. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22. ACM, 1994.

[19] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.

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