Effective Diversity Maintenance in Deceptive Domains

Joel Lehman
University of Texas at Austin
1 Inner Campus Drive
Austin, TX 78712 USA
joel@cs.utexas.edu

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

Risto Miikkulainen
University of Texas at Austin
1 Inner Campus Drive
Austin, TX 78712 USA
risto@cs.utexas.edu

ABSTRACT

Diversity maintenance techniques in evolutionary computation are designed to mitigate the problem of deceptive local optima by encouraging exploration. However, as problems become more difficult, the heuristic of fitness may become increasingly uninformative. Thus, simply encouraging genotypic diversity may fail to much increase the likelihood of evolving a solution. In such cases, diversity needs to be directed towards potentially useful structures. A representative example of such a search process is novelty search, which builds diversity by rewarding behavioral *novelty*. In this paper the effectiveness of fitness, novelty, and diversity maintenance objectives are compared in two evolutionary robotics domains. In a biped locomotion domain, genotypic diversity maintenance helps evolve biped control policies that travel farther before falling. However, the best method is to optimize a fitness objective and a behavioral novelty objective together. In the more deceptive maze navigation domain, diversity maintenance is ineffective while a novelty objective still increases performance. The conclusion is that while genotypic diversity maintenance works in well-posed domains, a method more directed by phenotypic information, like novelty search, is necessary for highly deceptive ones.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning—*connectionism and neural nets, concept learning*

General Terms

Algorithms

Keywords

NEAT, Diversity Maintenance, Novelty Search, Deception

1. INTRODUCTION

Deception is a fundamental problem in evolutionary computation (EC) wherein optimizing a fitness function often

leads an evolutionary algorithm (EA) to converge to a suboptimal solution [7, 16]. Thus, many researchers have proposed mechanisms for combating deception, often by encouraging or maintaining genotypic diversity during evolutionary search [7, 11, 12, 18]. The motivation is that by exploring a wider set of possible solutions it should be less likely for search to prematurely converge to degenerate locally optimal solutions.

However, recent research in evolutionary robotics suggests that deception may be a more serious problem than previously thought [14, 17]. While fitness heuristics (and objective-based optimization approaches in general) are founded on the assumption that steps to the solution will increasingly resemble it, in complex problems steps that ultimately lead towards the objective may yet poorly resemble it [17]. Thus with increasing problem difficulty, fitness functions may sometimes become *perversely deceptive*, i.e. nearly all of the steps in the sequence of evolutionary innovations necessary to solve a particular problem may actively be *punished* by optimizing resemblance to the desired outcome. For example, at nearly all points along the shortest path through a circuitous maze, traveling directly towards the goal will actually oppose progress.

Problematically, in such perversely deceptive problems the fitness function becomes a liability, and genotypic diversity maintenance alone may not significantly increase the effectiveness of search because the gradient of increasing fitness remains search's primary compass. Indeed, it has been shown in many deceptive problems that driving search *only* by phenotypic diversity can be more effective than guiding the search for the objective *by* the objective (i.e. optimizing fitness) [13, 16, 17, 22]. Although there exist a variety of approaches for encouraging phenotypic diversity in evolutionary robotics [9, 16, 20, 27], novelty search is different in that it can be used alone to solve problems, i.e. without being combined with pressure to optimize objective fitness [13, 14, 22].

Motivated by such results, Lehman and Stanley [17] proposed that search processes driven primarily by phenotypic diversity (e.g. novelty search) may scale to solve difficult problems more readily than search processes guided primarily by the objective (e.g. traditional objective-based fitness functions). The reason is that such search does not depend upon the fitness function conveniently identifying the order of stepping stones that lead to the objective. Instead, it relies on the assumption that previously discovered innovations are stepping stones to further evolutionary innovations, which may potentially hold more broadly. In other words,

while genotypic diversity maintenance creates random raw material from which fitness selects the most promising ones, novelty search discovers structure in the domain, thereby making progress possible even when fitness is uninformative.

This paper evaluates this hypothesis further by testing various combinations of common evolutionary search objectives (e.g. novelty, fitness, and diversity maintenance objectives) empirically on problems of increasing difficulty. In this way, in addition to exploring how fitness-based search and novelty search scale with respect to problem difficulty, insight may be gained into what combinations of objectives might be most useful for practitioners to apply to different types of problems.

To facilitate such a study, two well-studied evolutionary robotics domains are adapted to include an axis of increasing difficulty: biped locomotion [16, 21, 28] and maze navigation [16, 17, 20]. Across both domains and over problem instances of varying difficulty, various combinations of fitness, novelty, and diversity maintenance objectives are employed to guide evolutionary search.

In the biped domain, diversity maintenance techniques in general increase performance, although the overall most effective approach is optimizing a fitness objective and a novelty objective simultaneously. The biped domain is an example of a challenging but not perversely deceptive problem, wherein diversity can allow search to cross locally optimal fitness peaks to improve solution quality.

In contrast, in maze navigation there is little benefit from any of the genotypic diversity maintenance objectives. However, search still benefits from the novelty objective. The reason is that many of the intermediate products that ultimately lead to successfully navigating a particular maze may be actively punished by the fitness function. Because mazes can be seen as an abstract model for search problems in general, this result suggests that increasingly ambitious problems may in general become perversely deceptive.

The conclusion for practitioners is that for difficult real-world problems in which deception is not perverse, combining fitness-based search with novelty search is an effective approach, as has indeed been demonstrated before [20]. However, for researchers scaling evolutionary algorithms to more deceptive problems, the message is different: The objective-based search paradigm may be poorly-suited, and instead, open-ended evolutionary approaches such as novelty search may be needed.

2. BACKGROUND

This section reviews the NEAT neuroevolution approach, multi-objective optimization, diversity maintenance techniques in EC, and novelty search.

2.1 NEAT

In experiments described in this paper, behaviors are evolved for robots that are controlled by artificial neural networks (ANNs). Thus a neuroevolution (NE) method is needed to underpin these experiments. The NEAT method is appropriate because it is widely applied [1, 3, 16, 22, 24–26] and well understood.

The NEAT method was originally developed to evolve ANNs to solve difficult control and sequential decision tasks [24–26]. Evolved ANNs control agents that select actions based on their sensory inputs. Like the SAGA method [10] introduced before it, NEAT begins evolution with a population of small, simple networks and *complexifies* the network topology into diverse species over generations, leading to increasingly sophisticated behavior. A similar process of gradually adding new genes has been confirmed in natural evolution [19]. This section briefly reviews the NEAT method; for comprehensive introductions see e.g. [25, 26].

To keep track of which gene is which while new genes are added, a historical marking is uniquely assigned to each new structural component. During crossover, genes with the same historical markings are aligned, producing meaningful offspring efficiently. Speciation in NEAT protects new structural innovations by reducing competition among differing structures and network complexities, thereby giving newer, more complex structures room to adjust. Networks are assigned to species based on the extent to which they share historical markings. Complexification, which resembles how genes are added over the course of natural evolution [19], is thus supported by both historical markings and speciation, allowing NEAT to establish high-level features early in evolution and then later elaborate on them. Further, NEAT's ability to evolve increasingly complex ANNs fits well with this paper's motivation of solving increasingly difficult problems.

Note that as originally described, NEAT speciates the population to encourage genotypic diversity. Because this paper's experiments explore the effects of diversity maintenance on search, NEAT is extended here to run without speciation, and speciation is replaced by other diversity maintenance techniques. Additionally, NEAT is also extended to simultaneously optimize two objectives through multi-objective optimization, which is reviewed next.

2.2 Multi-objective Optimization

Multi-objective optimization is a popular paradigm within EC that addresses how more than one objective can be optimized at the same time in a principled way [4]. It has previously been applied to neuroevolution with promising results [2, 20]. Some of the experiments in this paper apply an implementation of NSGA-II [5], a well-established Pareto-based multi-objective search algorithm, to optimize two objectives concurrently.

The concept of 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.

In this way, 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. In this way, various trade-offs of competing objectives such as novelty, fitness, and measures of diversity can be explored during a single

evolutionary run. The idea is that such trade-offs allow for principled exploration of both objectives in cases where there is no clear and principled mechanism for combining two objectives together into one.

In the experiments in this paper, one of such simultaneously explored objectives maintains diversity in the population, which is reviewed next.

2.3 Diversity Maintenance in EC

Because naively optimizing an objective function often results in converging to a point that is optimal only locally [7, 16], practical EC approaches often employ diversity maintenance techniques that incentivize evolving individuals for reasons beyond simply increasing performance [7, 11, 12, 18]. A central motivating idea is that diversity maintenance may allow an EA to avoid being entirely *deceived* by overcomitting to a singular gradient of increasing fitness that actually leads *away* from the solution. That is, diversity maintenance encourages exploring multiple promising pathways through the search space simultaneously in hopes of thereby reaching higher fitness peaks.

In general, what differentiates most diversity maintenance techniques is the choice of by what dimension to characterize diversity. For example, diversity can be encouraged in the space of genotypes [8, 18], the age of genotypes [11, 23], the fitness of genotypes [12], or the space of phenotypic behaviors [9, 16, 20, 27]. Specifically, this paper applies genotypic speciation, an age objective, and a behavioral novelty objective as representative diversity maintenance techniques in evolutionary robotics.

Speciation, as applied in the canonical NEAT approach [25], separates individuals into different species through applying a genotypic distance metric. The idea is to encourage exploring ANNs with different topologies and also to ensure that crossover is restricted to recombining sufficiently similar ANNs. In particular, NEAT's speciation is based on explicit fitness sharing [8].

In contrast, the age-layered population structure (ALPS; [11]) characterizes genotypes more coarsely, by their *age*, i.e. how many generations have elapsed since they were randomly generated. In ALPS, on a regular basis new random genetic material is injected into the system, and genomes are segmented into layers based on their age such that older genomes do not compete with younger genomes. The idea is that similar to how fresh PhDs do not compete with senior professors in the job market, younger genomes may benefit from reduced competition to potentially converge to a different and perhaps more promising optimum than their older peers.

As originally described, ALPS consists of many separated layers with specific rules for transmission between them [11]; however, Schmidt and Lipson [23] showed that simply having an *age objective* in the context of multi-objective search often achieves similar performance. The idea is that if all else is equal then a search should prefer a younger genome over an older one to increase diversity. Thus for reasons of simplicity the experiments in this paper apply this concept of an age objective instead of the full ALPS setup.

It is important to note that speciation and ALPS encourage diversity without taking into account any aspect of phenotypic behavior. However, because many genotypes may map to the same phenotype and many genes may serve no functional purpose, distance metrics based only on geno-

types may not always encourage sufficient *functional* diversity to escape deceptive local optima [16]. Likewise, simply protecting younger genomes may not help escaping from deceptive fitness gradients if they are pervasive. In other words, most diversity maintenance techniques do not drive towards inherently useful information and would not make sense as the sole guidance for search [14]. In contrast, searching for novel behaviors encourages different functionality and can solve problems even when it is not combined with optimizing objective-based fitness [16]. Thus it may sometimes be more effective to encourage behavioral novelty or diversity [9, 16, 20, 27] instead of diversity in the genotypic space.

This motivates the next section, which reviews the idea of searching for behavioral novelty. Although such novelty search can be applied to complement the search for higher fitness, in its original conception it acts as the *primary* guidance for evolutionary search.

2.4 Novelty Search

In contrast to most EAs, which tend to converge, novelty search is a *divergent* evolutionary technique. It is inspired by natural evolution's drive to novelty, and directly rewards novel behavior *instead* of progress towards a fixed objective [15, 16]. This paper applies novelty search alone and in combination with other objectives to examine how the performance of such differing combinations of objectives scales to solve problems of increasing difficulty.

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 metric characterizes 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. 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

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

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. Candidates from more sparse regions of the behavior space then receive higher novelty scores.

If novelty is sufficiently high at the location of a new individual, i.e. above some minimal threshold ρ_{min}, then the individual is entered into the permanent archive that characterizes the distribution of prior solutions in behavior space. 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. However, even without an
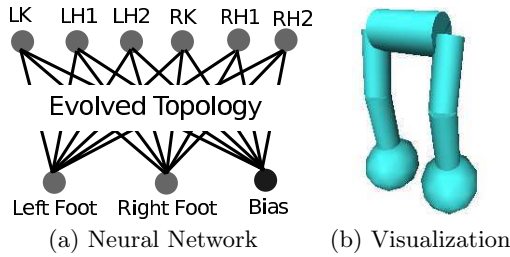
(a) Neural Network (b) Visualization

Figure 1: **Biped Robot.** In the biped locomotion domain, the ANN in (a) controls the biped robot that is visualized in (b). The robot has motors that apply forces to achieve the joint angles that are output by the ANN. In particular, the biped has motors for each of its six degrees of freedom: One in both its left and right knees (LK and RK), and two in each hip (LH1, LH2, RH1, and RH2). Additionally, the robot's ANN receives input from sensors in its feet that activate when they touch the ground. The goal is for the robot to locomote as far as possible, which is deceptive from the local optimum induced by lunging forward violently. Difficulty in the domain can be scaled by adjusting the intensity of forces that help to keep the robot upright.

explicit objective, novelty search is still driven by meaningful information; that is, behaving in a novel way often requires learning the structure of the domain.

Once objective-based fitness is replaced with novelty (or alternatively is *augmented* by 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.

3. APPROACH

The approach in this paper is to study empirically how the effectiveness of different combinations of common objectives for EAs scale with problems of increasing difficulty. Thus a means of generating problems of variable difficulty is needed. For this purpose, two well-explored evolutionary robotics domains are adapted: biped locomotion [16, 21, 28] and maze navigation [16, 17, 20].

3.1 Variable Difficulty in Biped Locomotion

The biped locomotion domain from Lehman and Stanley [16] is adapted for these experiments. In this domain, the goal is to evolve an ANN controller for a three-dimensional simulated biped robot (figure 1) that can locomote as far as possible within the given time limit, which is challenging because of the need for balance and oscillatory gaits. As in Lehman and Stanley [16], the fitness of a biped controller is evaluated as the squared distance the robot walks before it falls, while its behavior (for calculating novelty) is derived from sampling its center of gravity over time. That is, fitness-based search looks for walkers that travel farther while novelty search rewards *different* walking trajectories. Exact details of the simulated biped robot can be found in Lehman and Stanley [16]. Note that in all the biped experiments in this paper, evolution was run for 2,000 generations with a population size of 500; other parameters are the same as those used by Lehman and Stanley [16].

In its original formulation, the biped problem is deceptive because an attractive local optimum is to lunge forward, which maximizes distance traveled without having any relation to a successful oscillatory gait [16, 28]. To create *easier*

variations of this problem with *less* deception, supportive forces are added that act to keep the robot upright, like the approach that van de Panne and Lamouret [28] called "the hand of God." In particular, assuming the angular velocity of the robot's torso is v, the ideal up vector representing the desired orientation of the robot's torso is u, the current imperfect up vector is p, and a difficulty parameter ranging from 0.0 to 1.0 is d, then a lifting force is added with intensity $20(1.0 - d)$ newtons, a damping torque is added equal to $-0.3(1.0 - d)v$ newton-meters, and a corrective torque is added equal to $2.0(1.0 - d)p \times u$ newton-meters. The idea of these forces is to make balancing easier, in order to reduce the difficulty in evolving an effective cyclical gait. Thus the deceptiveness of the problem can be manipulated by varying the difficulty parameter d, which controls the intensity of the corrective forces (e.g. decreasing d decreases the supportive forces and makes the problem more deceptive). For example, when d is set to its maximum (1.0), then no supportive forces are applied and the domain is at its most deceptive and is equivalent to how it is described in Lehman and Stanley [16].

The next section describes how a similar axis of difficulty is created in the maze navigation domain.

3.2 Variable Difficulty in Maze Navigation

The maze navigation domain is adapted from Lehman and Stanley [16], where simulated robots controlled by evolved ANNs must navigate in a maze from a starting point to a goal point within a fixed amount of time. Traversal of the maze is complicated by walls that require circumnavigation. The simulated robot is shown in figure 2. Note that the simulation is ego-centric and continuous (i.e. the robot has incomplete local information about its environment, and its sensors, effectors, position, and orientation are represented by continuous variables), which makes the problem more difficult. As in previous studies [16, 17, 20], the fitness function measures how close a maze navigator ends to the goal location, which is an intuitive measure of progress in navigation tasks; likewise, for novelty search the behavior of a navigator is defined as the Cartesian coordinate of the navigator at the end of an evaluation. In this way, fitness-based search seeks individuals that minimize distance to the goal while novelty search seeks individuals that end in different locations from those previously encountered during search.

To generate a wide variety of problems, a recursive backtracking algorithm was applied to generate random mazes [6]. The initial and goal points of the maze were also randomly chosen, which, when considered with the walls, specifies a complete problem instance. The difficulty of such a generated problem was estimated by the length of the optimal path from the starting point to the goal, guided by the intuition that longer paths will in general require more complex policies to navigate (this intuition is supported empirically by the results).

In particular, 15 random maze problems on a $4x4$ grid were generated for each path length ranging from four to 14, yielding a total of 165 problem instances. Examples of an easy and a difficult maze problem are shown in figure 3. Evolution was run for a maximum of 1,000 generations with a population size of 250; other parameters are the same as those used by Lehman and Stanley [16].
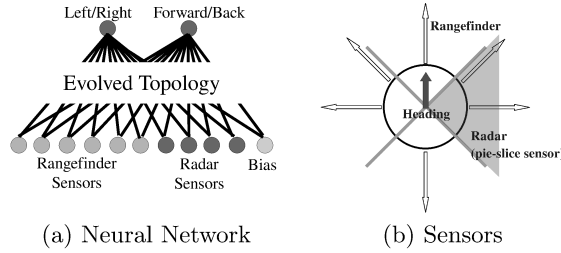
(a) Neural Network   (b) Sensors

Figure 2: **Maze Navigating Robot.** In the maze navigation domain, the ANN in (a) controls the maze navigating robot. The layout of the robot's sensors is shown in (b). Each arrow outside of the robot's body in (b) is a rangefinder sensor that indicates the distance to the closest obstacle in that direction. The robot has four pie-slice sensors that act as a compass towards the goal, activating when a line from the goal to the center of the robot falls within the pie-slice. The solid arrow indicates the robot's heading. The task is for the robot to navigate through a maze of parametrized difficulty. The mazes become perversely deceptive as the shortest path through them becomes increasingly circuitous.
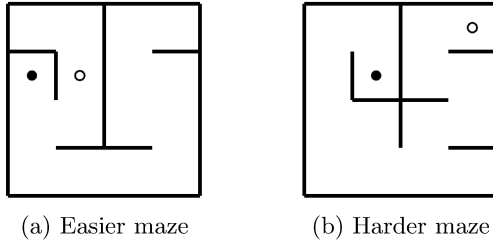

(a) Easier maze   (b) Harder maze

Figure 3: **Examples of Randomly Generated Mazes.** In both mazes, the unfilled circle represents the starting position of the robot and the filled circle represents the goal. A short solution path exists in (a), which is why it classified as an easy problem instance. The optimal path from the starting position to the goal in (b) is longer and this is instance is thereby classified as being more difficult.

3.3 Combinations of Evolutionary Search Objectives

This section describes the particular combinations of evolutionary search objectives that are tested in the experiments. First, the controls are *fitness alone* and *novelty alone*, which are driven only by a fitness objective and a novelty objective, respectively, i.e. they are not combined with a traditional diversity maintenance technique. In practice these methods were implemented by disabling the speciation mechanism that is specified by the standard NEAT approach. In this way it is possible to establish a baseline level of performance and to explore in what situations *lacking* any kind of standard diversity maintenance is problematic.

Next there are methods that combine both fitness and novelty objectives with common genotypic diversity maintenance techniques. First, genotypic speciation is *enabled* to test both fitness and novelty search with the standard NEAT approach as in prior novelty-search NEAT studies [16, 22]. These are the *fitness with speciation* and *novelty with speciation* methods. While such speciation is not an objective in a Pareto multi-objective sense, in effect it serves to combine optimization pressure with a drive towards genotypic diversity. Similarly, the addition of an ALPS-like age objective [11, 23] is tested by applying a multi-objective version

of NEAT to simultaneously optimize the age objective and either a fitness or novelty objective, which are, respectively, the *fitness with age* and *novelty with age* methods. Note that this multi-objective version of NEAT operates without speciation, thereby isolating the effects of the age objective to encourage genotypic diversity.

Finally, the last method combines a fitness objective and a novelty objective together, similarly to the multi-objectivization approach by Mouret and Doncieux [20]. The same multi-objective version of NEAT as in the age-objective method is applied to optimize both a novelty and fitness objective simultaneously, resulting in the *fitness and novelty* method. The idea is that novelty search may drive towards a diverse set of behaviors and that when it uncovers an individual in the basin of attraction of a solution, optimizing the fitness objective may then efficiently encourage that solution to be evolved.

The next section describes the results of applying these various methods to problems of increasing difficulty in the two domains.

4. RESULTS

Figure 4 compares the average performance of various methods across instances of the biped problem of increasing difficulty. Note that performance is measured as the average distance traveled by run champions; such performance of methods generally declines as problem difficulty increases, supporting the way problem difficulty was implemented in this domain. A coarser view of the results is given by figure 5, which aggregates a method's average performance across all instances of the biped problem.

Note that the overall best-performing method is the fitness and novelty method, reflected by its significantly higher aggregate performance when compared pairwise to the other methods (Student's t-test; $p < 0.05$). At a finer granularity, it also generally performs better than the other methods in the 11 biped domain difficulty variants (in 54 out of 66 such pairwise comparisons; Student's t-test; $p < 0.05$). Additionally, figures 4a and 4b show that combining either fitness or novelty with speciation provides a general benefit; for both fitness and novelty, the differences in performance between including and not including speciation are significant for each of the 11 biped difficulty levels (Student's t-test; $p < 0.05$). In contrast, augmenting novelty or fitness with an age objective is significantly deleterious in this domain. It is possible that using the complete ALPS setup may provide different results; in such a case these results would suggest that the results of Schmidt and Lipson [23] do not completely generalize.

Figure 4c compares the most representative techniques, showing that novelty search with speciation tends to outperform fitness-based search with speciation, although the advantage decreases with increasing problem difficulty. In particular, of the 11 possible pairwise comparisons, novelty with speciation is significantly better in the seven intermediate difficulty settings from 0.3 to 0.9, is significantly worse when difficulty is set to 0.2, and does not differ significantly in the three remaining difficulties (0.0,0.1,and 1.0). An interesting result is the lack of significant difference between novelty search with speciation and fitness-based search with speciation when the difficulty is set to 1.0; this particular setting (where there are no supportive forces) is equivalent to the original problem from Lehman and Stanley [16]. How-

(a) Fitness-based variants   (b) Novelty search variants   (c) Direct comparison
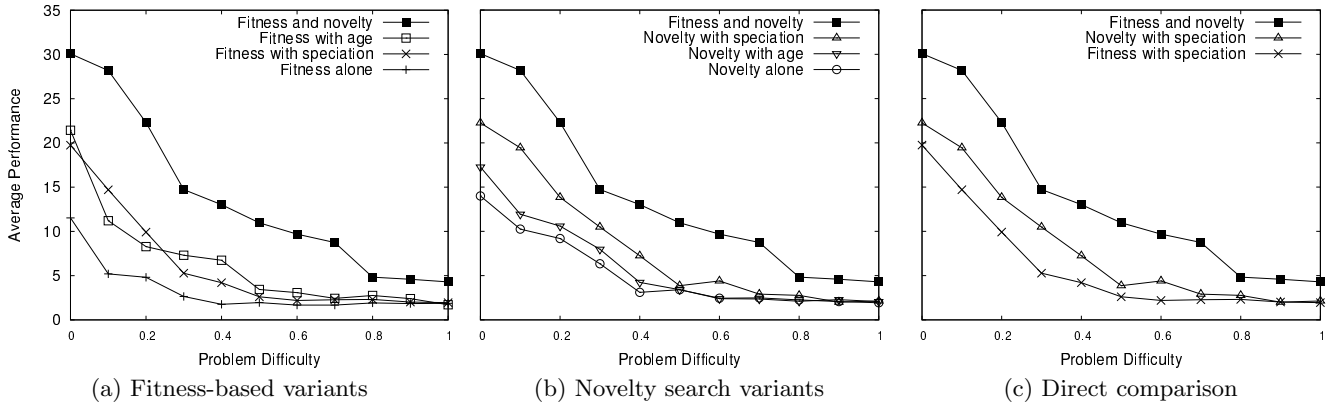
Figure 4: **The effectiveness of guiding search by various combinations of evolutionary objectives as the difficulty of the biped domain increases.** Diversity maintenance techniques are shown combined with (a) fitness-based search and (b) novelty search. A direct comparison of representative approaches is shown in (c). The conclusion is that the most effective approach in this domain is optimizing fitness and novelty together.

ever, in that paper novelty search did significantly outperform fitness-based search. This discrepancy suggests that the *generational* novelty search applied in the experiments in this paper may provide a less informative gradient of novelty than the *steady-state* novelty search of Lehman and Stanley [16]. That is, gradients of novelty are more stable when replacing only one individual at a time instead of replacing the entire population. Generational search was used (as in some prior novelty search studies [17, 20]) because the NSGA-II search algorithm chosen for optimizing multiple objectives is generational, and standardizing on generational search allows for uniformity and easier interpretation of results.

Next, the results for the maze navigation experiment are shown in figures 6 and 7. As in the biped domain, performance of all methods tends to decrease as problem difficulty increases. Interestingly however, the maze results exhibit little benefit from either adding diversity maintenance to fitness-based search (figure 6a) or novelty search (figure 6b), or from optimizing fitness and novelty simultaneously (fig 6c). Instead, the main differentiator is simply whether or not a novelty objective is included: With a novelty objective, more of the difficult problem instances are solved.

Supporting this statement, figure 7 presents the total proportion of successful evolutionary runs from each of the methods. This quantity is significantly higher for methods that include a novelty objective compared to those that do not (Fischer's exact test; $p < 0.001$). There is no significant difference *among* the methods including a novelty objective.

Furthermore, when looking more granularly at the performance of methods across the 11 individual problem difficulty levels shown in figure 6, there are no significant differences among methods including a novelty objective nor among methods not including it. However, in 73 out of 132 possible pairwise comparisons of performance on problems of a specific difficulty *between* methods including novelty and those not including it, those including novelty performed significantly better. Also, such significant differences only occurred when difficulty was seven or higher. In the remaining 59 pairwise comparisons (mainly comprising the easiest problem instances) there was no significant difference in performance between the novelty methods and those without novelty.
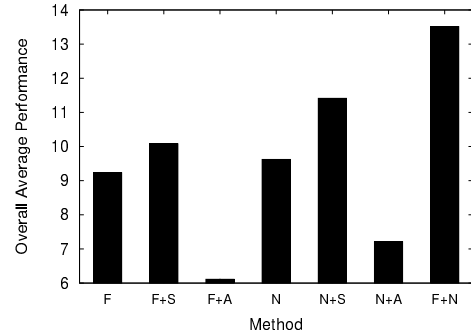

Figure 5: **Aggregate comparison of the biped experimental methods.** The average performance of the champions from the 220 total runs for each method are shown. The methods are abbreviated according to the driving forces of its search: N indicates novelty, F is for fitness, S is for speciation, and A indicates an age objective (e.g. N+S represents the novelty with speciation method and F+N represents the fitness and novelty method). The main result is that the method combining novelty and fitness is the most effective.

5. DISCUSSION

The results in this paper differ significantly over the two explored domains. In the biped domain, combining fitness and novelty is more effective than optimizing either one separately, and there is some benefit achieved from actively maintaining genotypic diversity. In contrast, in the maze navigation domain the most important factor is whether a behavioral novelty objective is included or not, and there is little benefit from the other objectives, whether fitness or genotypic diversity maintenance.

Importantly, these two different problems can be seen as representative of the distinct challenges facing practitioners and researchers. For practitioners seeking to exploit EC to solve challenging problems, the biped problem provides a good model; it is deceptive and difficult, but the stepping stones to a solution are known (oscillatory leg movements and balance) and are generally recognized by intuitive fitness functions (e.g. distance traveled before falling). This is a problem that objective-based search can reasonably solve,

(a) Fitness-based variants   (b) Novelty search variants   (c) Direct comparison
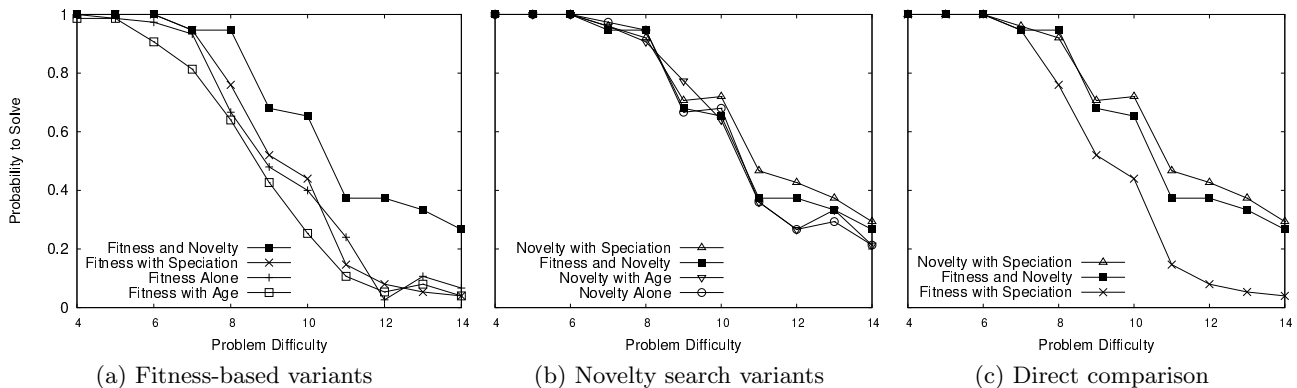
Figure 6: **Maze Navigation Results.** This figure shows how effective guiding search by various combinations of evolutionary objectives is as the difficulty of randomly generated mazes increases (which is estimated by the length of the optimal path through the maze). A variety of diversity objectives are shown combined with (a) fitness-based search and (b) novelty search. A direct comparison of representative approaches is shown in (c). The main result is that the methods including a novelty objective perform better than methods without such an objective in more difficult maze problems.

and more powerful diversity maintenance techniques can increase performance. In particular, the biped results here augment existing evidence that phenotypic diversity may be a more meaningful class of diversity to encourage than diversity measures based on genotypes alone [9, 16, 20, 27].

While the biped problem is interesting and difficult, for researchers seeking to understand the ultimate limitations of both diversity maintenance and the objective-based search paradigm that dominates EC, the maze navigation problem is a more appropriate model. It illustrates that no amount of diversity maintenance may help if a problem is perversely deceptive. Most of the stepping stones needed to solve the problem are not recognized by the fitness function, and may be actively *penalized* by it. Such perverse deception arises in this domain as the mazes become more difficult and the path from the start to the goal becomes increasingly circuitous, rendering the heuristic of Euclidean distance to the goal nearly meaningless. In such cases, encouraging genotypic diversity does not help because the gradients of the fitness function overwhelmingly direct search actively away from the solution.

However, although the maze navigation domain is perversely deceptive, it is not without structure that can be exploited by search. There are fundamental stepping stones that make sense even without an explicit fitness objective, such as traversing corridors, avoiding dead ends, and exploring an area systematically. Discovering such stepping stones makes it possible to develop more novel and more complex behaviors, thereby facilitating progress in search even with extreme deception.

Of course, the maze domain is well-suited for novelty search and thus the message is *not* that novelty search itself will always outperform objective-based search, or that it will scale to solve the most ambitious problems. Rather, the idea is that novelty search is representative of a different class of evolutionary heuristics that may have better scalability properties than traditional objective-based fitness.

That is, the maze navigation domain and the application of fitness and novelty search can be interpreted abstractly as a model representing problem solving in general [16, 17]; the goal point can be seen abstractly as the solution for a problem and the walls as domain constraints that compli-

cate the solution's form. The hypothesis is that the more complicated the problem is, the less likely it is that the necessary stepping stones in the evolutionary lineage leading to a solution will increasingly resemble that same final solution.

In other words, solving increasingly complex problems with a fitness-based approach may be unproductive even with diversity maintenance unless increasing amounts of information about the stepping stones are also injected into the fitness function. Because such injection runs counter to the idea of *artificial* intelligence, it may be important to further consider open-ended evolutionary search processes able to automatically *accumulate* information from the domain instead of having it specified by a domain expert. In this way, novelty search is an interesting example of such a search process that may point towards promising new research directions for EC.

Finally, note that all methods decline in performance as problem complexity increases. At some very high level of problem complexity, no method may solve a particular problem reliably. Yet even then novelty may help discover highly creative and complex artifacts that do not necessarily solve a particular given problem, but that are important discoveries nonetheless. That is, search need not only be about solving specific problems, but can also be about discovering unanticipated artifacts in the search space that are interesting in their own right. This possibility is interesting and deserves further study.

6. CONCLUSION

This paper evaluated how genotypic diversity maintenance techniques combined with fitness-based search and novelty search scale to solve problems of increasing complexity. Traditional diversity maintenance techniques facilitate greater exploration around high-fitness areas by generating random diversity on a genetic level. In contrast, searching for novelty can drive towards useful stepping stones that reflect underlying structure in a domain; that is, incorporating information about the domain into the genome often facilitates generating novelty. Thus in difficult but only superficially deceptive problems, genotypic diversity maintenance is often effective: Random diversity can aid escaping from shallow local optima. However, with increased deception the gradients of
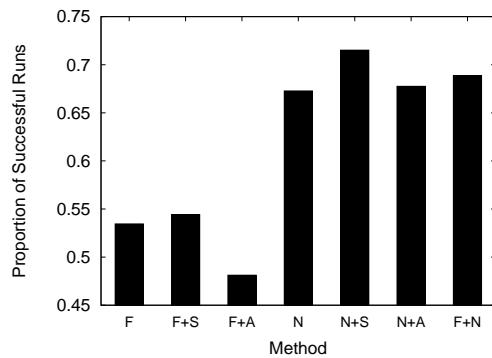
Figure 7: **Aggregate comparison of maze navigation experimental methods.** The percentage of the 825 total runs for each method that are successful are shown. The methods are abbreviated by the search objectives that they reward: N indicates novelty, F is for fitness, S is for speciation, and A indicates an age objective (e.g. N+S represents the novelty with speciation method and F+N represents the fitness and novelty method). The main result is that the methods including novelty have a significantly higher percentage of successful runs than those methods not including novelty.

the fitness function become overwhelmingly misleading and genotypic diversity maintenance may not provide much benefit. Thus, abstractions of natural evolution that are based on phenotypic diversity, and especially those that discover fundamental structure in the domain, may become increasingly important to EC researchers hoping to scale EC to evolve biological levels of complexity.

References

[1] T. Aaltonen et al. Measurement of the top quark mass with dilepton events selected using neuroevolution at CDF. *Physical Review Letters*, 2009.

[2] H. A. Abbass. A memetic pareto evolutionary approach to artificial neural networks. In *Proc. of the 14th Australian Joint Conf. on Artificial Intelligence*, 2001.

[3] Brian Allen and Petros Faloutsos. Complex networks of simple neurons for bipedal locomotion. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009.

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

[5] 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.

[6] M. Foltin. Automated maze generation and human interaction. 2008.

[7] David E. Goldberg. Simple genetic algorithms and the minimal deceptive problem. In L. D. Davis, editor, *Genetic Algorithms and Simulated Annealing, Research Notes in Art. Intell.*. Morgan Kaufmann, 1987.

[8] David E. Goldberg and J. Richardson. Genetic algorithms with sharing for multimodal function optimization. In J. J. Grefenstette, editor, *Proceedings of the Second International Conference on Genetic Algorithms*, pages 148–154. San Francisco: Kaufmann, 1987.

[9] F.J. Gomez. Sustaining diversity using behavioral information distance. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 113–120. ACM, 2009.

[10] Inman Harvey. *The Artificial Evolution of Adaptive Behavior*. PhD thesis, School of Cognitive and Computing Sciences, University of Sussex, Sussex, 1993.

[11] 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.

[12] Marcus Hutter and Shane Legg. Fitness uniform optimization. *IEEE Transactions on Evolutionary Computation*, 10:568–589, 2006.

[13] Peter Krcah. Solving deceptive tasks in robot body-brain co-evolution by searching for behavioral novelty. In *ISDA*, pages 284–289. IEEE, 2010.

[14] Joel Lehman. *Evolution Through the Search for Novelty*. PhD thesis, University of Central Florida, Orlando, Florida, 2012.

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

[16] Joel Lehman and Kenneth O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evol. Comp.*, 19(2):189–223, 2011.

[17] Joel Lehman and Kenneth O. Stanley. Novelty seach and the problem with objectives. In *Genetic Programming in Theory and Practice IX (GPTP 2011)*, chapter 3, pages 37–56. Springer, 2011.

[18] Samir W. Mahfoud. *Niching methods for genetic algorithms*. PhD thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA, 1995.

[19] Andrew P. Martin. Increasing genomic complexity by gene duplication and the origin of vertebrates. *The American Naturalist*, 154(2):111–128, 1999.

[20] J.-B. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: an empirical study. *Evolutionary Computation*, 20(1):91–133, 2012.

[21] Torsten Reil and Phil Husbands. Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168, April 2002.

[22] S. Risi, C.E. Hughes, and K.O. Stanley. Evolving plastic neural networks with novelty search. *Adaptive Behavior*, 2011.

[23] M. Schmidt and H. Lipson. Age-fitness pareto optimization. *Genetic Programming Theory and Practice VIII*, pages 129–146, 2011.

[24] Kenneth O. Stanley, Bobby D. Bryant, and Risto Miikkulainen. Real-time neuroevolution in the NERO video game. *IEEE Transactions on Evolutionary Computation Special Issue on Evolutionary Computation and Games*, 9(6):653–668, 2005.

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

[26] Kenneth O. Stanley and Risto Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Art. Intell. Research*, 21:63–100, 2004.

[27] Leonardo Trujillo, Gustavo Olague, Evelyne Lutton, Francisco Fernandez De Vega, León Dozal, and Eddie Clemente. Speciation in behavioral space for evolutionary robotics. *Journal of Intelligent & Robotic Systems*, 64(3-4):323–351, 2011.

[28] Michiel van de Panne and Alexis Lamouret. Guided optimization for balanced locomotion. In D. Terzopoulos and D. Thalmann, editors, *Sixth Eurographics Workshop on Animation and Simulation*, volume 95, pages 165–177. Springer-Verlag, 1995.