Contents

1
Novelty Search and the Problem with Objectives 1
Joel Lehman and Kenneth O. Stanley

Index 21

Chapter 1

NOVELTY SEARCH AND THE PROBLEM WITH OBJECTIVES

TO APPEAR IN: *GENETIC PROGRAMMING THEORY AND PRACTICE IX (GPTP 2011).*

NEW YORK, NY: SPRINGER

Joel Lehman[1] and Kenneth O. Stanley[1]

[1]*Department of EECS, University of Central Florida, Orlando, Florida, USA;*

Abstract By synthesizing a growing body of work in search processes that are not driven by explicit objectives, this paper advances the hypothesis that there is a fundamental problem with the dominant paradigm of objective-based search in evolutionary computation and genetic programming: Most *ambitious* objectives do not illuminate a path to themselves. That is, the gradient of improvement induced by ambitious objectives tends to lead not to the objective itself but instead to dead-end local optima. Indirectly supporting this hypothesis, great discoveries often are not the result of objective-driven search. For example, the major inspiration for both evolutionary computation and genetic programming, natural evolution, innovates through an open-ended process that lacks a final objective. Similarly, large-scale cultural evolutionary processes, such as the evolution of technology, mathematics, and art, lack a unified fixed goal. In addition, direct evidence for this hypothesis is presented from a recently-introduced search algorithm called novelty search. Though ignorant of the ultimate objective of search, in many instances novelty search has counter-intuitively outperformed searching directly for the objective, including a wide variety of randomly-generated problems introduced in an experiment in this chapter. Thus a new understanding is beginning to emerge that suggests that searching for a fixed objective, which is the reigning paradigm in evolutionary computation and even machine learning as a whole, may ultimately *limit* what can be achieved. Yet the liberating implication of this hypothesis argued in this paper is that by embracing search processes that are *not* driven by explicit objectives, the breadth and depth of what is reachable through evolutionary methods such as genetic programming may be greatly expanded.

Keywords: Novelty search, objective-based search, non-objective search, deception, evolutionary computation

1. Introduction

Evolutionary computation (EC; De Jong, 2006; Holland, 1975) and genetic programming (GP; Koza, 1992) are algorithmic abstractions of natural evolution, inspired by nature's prolific creativity and the astronomical complexity of its products. Supporting such abstractions, evolutionary algorithms (EAs) have achieved impressive results, sometimes exceeding the capabilities of human design (Koza et al., 2003; Spector et al., 1999). Yet the ambitious goal of evolving artifacts with complexity comparable to those crafted by natural evolution remains daunting.

An interesting question is what prevents EAs from evolving artifacts with a functional complexity of the magnitude seen in biological organisms. There are many possible answers, each pointing to potential faults in current EAs. For example, representation, selection, or the design of problem domains could each possibly be the paramount issue preventing higher achievement in EC, and there are researchers who investigate ways of improving each of these components (Pelikan et al., 2001; Stewart, 2001). This paper focuses on *selection* and argues that the currently dominant objective-based selection paradigm significantly limits the potential of EAs.

This handicap results from a well-known problem facing EAs called deception (Goldberg, 1987): Sometimes a mutation increases fitness but actually leads *further* from the objective. That is, the fitness function in EC is a heuristic and thus there is no guarantee that increasing fitness actually decreases the distance to the objective of the search. The fundamental problem is that the stepping stones that lead to the objective may not resemble the objective itself. For example, humans bear little resemblance to their flatworm ancestors. In the words of John Stuart Mill, it is a fallacy to assume that the "conditions of a phenomenon must, or at least probably will, resemble the phenomenon itself." (Mill, 1846, p. 470). Yet this fallacy is the very *foundation* of typical fitness functions and thus ultimately limits their effectiveness.

In practice, researchers become accustomed to the fragility of fitness functions and learn certain rules of thumb that guide their efforts. One prominent such inductive rule is that the more ambitious the objective is, the less likely evolution will be able to solve it. This intuition, supported by experiments in this paper with increasingly difficult problems, highlights the critical problem undermining objective-based search: While we want to harness evolution to solve ambitious problems, the more ambitious the objective is, the less informative the gradient of the induced objective function will be. A provocative question is whether the quest for the objective itself sometimes precludes search from achieving anything remarkable. In other words, could ignoring the ultimate objective of search, or even searching entirely without an explicit objective, sometimes be a more viable approach to discovery?

A recent technique in EC called novelty search (Lehman and Stanley, 2011; Lehman and Stanley, 2008) shows that indeed new opportunities for discovery arise once explicit objectives are abandoned: Contrary to intuition, searching without regard to the objective can often outperform searching explicitly for the objective. Instead of searching for the objective, novelty search only rewards individuals with functionality *different* from those that precede them in the search; inspired by nature's drive towards diversity, novelty search directly incentivizes novelty in lieu of any notion of progress. In a growing number of experiments, novelty search has successfully been applied to solving problems, often solving them more effectively than an algorithm searching directly for the objective (Lehman and Stanley, 2011; Risi et al., 2010; Lehman and Stanley, 2010a; Doucette, 2010; Mouret, 2009; Lehman and Stanley, 2010b). For example, an experiment introduced in this chapter demonstrates the advantages of novelty search in a wide variety of randomly-generated maze navigation problems.

However, novelty search provides but one example of a non-objective search, i.e. a search without a final explicit objective. A more prominent example is natural evolution itself, the process from which both EC and GP are inspired. While some might view reproduction as the goal of natural evolution, complex organisms such as ourselves are less efficient and slower to reproduce than simple single-celled creatures. Thus, what we might be tempted to characterize as progress in natural evolution is in fact quantitatively detrimental to the supposed objective of reproduction. That is, most innovation seen in natural evolution may result more from finding new ways of meeting life's challenges (i.e. founding new niches) than from simply optimizing reproductive fitness. Furthermore, nature does not aim at any single point or set of points in the space of organisms. In contrast, the objective in EC or GP is usually just such a point or set of points (e.g. the optimum or one of a set of optima).

Similarly, the evolution of mathematics, art, and technology are facilitated by exploration around recent discoveries, serendipity, and a plethora of diverse and conflicting individual objectives. That is, these human-driven processes of search also do not aim at any unified society-wide singular objective. Thus the types of search processes that continually innovate to produce radical advancements often lack a final predefined goal. This observation makes sense because a single fixed goal would either (1) be deceptive and therefore bring search to a point at which progress would effectively halt, or (2) if the goal is not so deceptive then innovation would cease once the goal is met.

The most awe-inspiring forms of search, which continually discover complex and interesting novelty, tend to build exploratively and incrementally upon prior innovations while lacking final objectives. When search is framed in this way, it is natural to ask, why is the typical approach in EC and GP to start from a random initial population and then to search narrowly towards a fixed goal? While it does indeed succeed in some cases, such objective-based search does not scale to the most ambitious objectives, e.g. the ones that natural evolution is able to reach, because the objective-based search paradigm constrains evolution in a particularly restrictive way. That is, natural evolution succeeds because it divergently explores many ways of life while optimizing a behavior (i.e. reproduction) largely *orthogonal* to what is interesting about its discoveries, while objective-based search directly follows the gradient of improvement until it either succeeds or is too far deceived.

The main implication of the hypothesis advanced in this paper is that to reach truly ambitious goals, EAs may need to be modified to exploit richer gradients of information than estimated distance to a fixed objective. Behavioral novelty is one such gradient, yet although novelty search does outperform objective-based search in many deceptive problems, it too pales in comparison to the creative achievement of natural evolution. That is, there still remains much work to be done in developing powerful non-objective search algorithms. Thus, while illustrating the limitations of objective-based search may be a negative result, at the same time it also illuminates an exciting and potentially profound challenge for researchers in GP and EC: Through exploring the mostly untamed wilderness of non-objective search algorithms we may be able to finally devise truly creative algorithms that continually yield innovative complex artifacts. This paper reviews a spectrum of recent work that supports this view, ultimately building an argument in favor of a wider perspective for GP and EC.

2. Deception

In this section we argue that deception is a deleterious fundamental property of ambitious objectives that paradoxically prevents such objectives from being reached when searching directly for them.

Investigating Deception

The motivation behind characterizing deception and problem difficulty is to understand what properties of problems may cause EAs to fail, so that such properties can potentially be remedied or avoided.

The original definition of deception (Goldberg, 1987) is based on the building blocks hypothesis, in which small genetic building blocks are integrated to form larger blocks (Holland, 1975). In the original conception, a problem is deceptive if lower-order building blocks, when combined, do not lead to a global optimum. Thus, in deceptive problems the fitness function may actively steer search away from exactly what is necessary to solve the problem.

Some alternative measures of problem difficulty attempt to quantify the ruggedness of the fitness landscape, motivated by the intuition that optimizing more rugged landscapes is more difficult (Weinberger, 1990). Importantly, because the fitness landscape is induced by the objective function, the problem of ruggedness, presupposing reasonable settings for the EA, can be attributed to the objective function itself.

Interestingly, other researchers suggest that ruggedness is overemphasized and that neutral fitness plateaus (i.e. neutral networks) are key influences on evolutionary dynamics (Barnett, 2001; Stewart, 2001). However, even neutral networks suggest a deficiency in the objective function: By definition a neutral part of the search space contains no gradient information with respect to the objective function. That is, in a neutral network the compass of the objective function is ambiguous with respect to which way search should proceed.

In summary, there are many ways to consider, measure, and model the difficulty of problems for EAs. While in general the exact properties of a problem that make it difficult for EAs are still a subject of research, in this paper the term deception will refer to an intuitive definition of problem hardness: A deceptive problem is one in which a reasonable EA (with a reasonable representation, parameters, and search operators) will not reach the desired objective in a reasonable amount of time. That is, a deceptive problem is simply a problem in which following the gradient of the objective function leads to local optima.

It is important to note that this definition of deception is different from the traditional definition (Goldberg, 1987). This intuitive approach helps to isolate the general problem with particular objective functions because the word "deception" itself reflects a fault in the *objective function* (as opposed to in the algorithm itself): An objective function with the pathology of deceptiveness will *deceive* search by actively pointing the wrong way.

Mitigating Deception

Ideally, there would exist a silver bullet method immune to the problem of deception such that any objective would be reachable in a reasonable amount of time. Although it is impossible that any such general silver bullet method exists (Wolpert and Macready, 1995), researchers strive to create methods that can overcome deception in practice.

Common approaches in EC to mitigating deception are diversity maintenance techniques (Mahfoud, 1995), building models that derive additional information from an imperfect fitness function (Pelikan et al., 2001), or accelerating search through neutral networks (Stewart, 2001). However, all of these techniques remain vulnerable to sufficiently uninformative objective functions.

In direct response to the problem of local optima when evolving towards sophisticated behaviors, some researchers incrementally evolve solutions by sequentially applying carefully crafted objective functions (Gomez and Miikkulainen, 1997). However, with ambitious objectives crafting an appropriate sequence of objectives may be difficult or impossible to achieve a priori. Additionally, the requirement of such intimate domain knowledge conflicts with the aspiration of *machine* learning.

In addition to single-objective optimization, there also exist evolutionary methods that optimize several objectives at once: Multi-Objective Evolutionary Algorithms (MOEAs) (Veldhuizen and Lamont, 2000). However, these MOEAs are not immune to the problem of deception (Deb, 1999), and adding objectives does not always make a problem easier (Brockhoff et al., 2007).

Another approach in EC related to deception is coevolution, wherein interactions between individuals contribute towards fitness. The hope is that continual competition between individuals will spark an evolutionary *arms race* in which the interactions between individuals continually creates a smooth gradient for better performance (Cliff and Miller, 1995). However, in practice such arm races often converge to situations analogous to local optima in standard objective-based search, e.g. mediocre stable-states, cycling between behaviors without further progress, or unbalanced adaptation where one species significantly out-adapts other competing species (Ficici and Pollack, 1998).

In summary, because deception is a significant problem in EC, there are *many* methods that have been designed to mitigate deception. However, while they may sometimes work, ultimately such methods do not cure the underlying pathology of the objective function that causes deception: The gradient of the objective function may be misleading or uninformative to begin with. Given a sufficiently uninformative objective function, it is an open question whether *any* method relying solely on the objective function will be effective. Thus an interesting yet sobering conclusion is that some objectives may be unreachable by direct objective-based search alone. Furthermore, as task complexity increases it is more difficult to successfully craft an appropriate objective function (Ficici and Pollack, 1998; Zaera et al., 1996). These insights match many EC practitioners' experience that the difficulty in ambitious experiments is often in crafting a sufficient fitness function. Thus the ultimate conclusion is that the more ambitious the experiment, the more likely it is that objective-based search will lead to mediocre local optima as opposed to the desired goal.

3. Non-objective Search

If one accepts that there can be no general solution to the fundamental problem of deception in objective-based search (Wolpert and Macready, 1995), it becomes important to consider alternative paradigms such as searches in which there is no a priori fixed objective.

Interestingly, natural evolution, which inspires both EC and GP, is an example of such a non-objective search. That is, there is no final organism for which natural evolution searches. While competition between organisms may increase reproductive fitness, the complexity of biological organisms that we are tempted to attribute to selection is instead nearly always quantitatively *detrimental* to fitness. That is, large complex organisms reproduce slower and less efficiently than simple single-celled organisms. Indeed, some biologists have argued that selection pressure may not explain innovation (Gould, 1996; Lynch, 2007). The conclusion is that innovation may result more from accumulating novel ways of life (i.e. new niches) than from optimizing fitness.

Similarly, cultural evolutionary processes such as the evolution of mathematics, art, and technology also lack a single unified objective. That is, there is no fixed final theory of mathematics, no final transcendent pinnacle of art, and no final culminating technology that these systems singularly strive towards. Innovation in such systems branches from existing innovations by local search empowered by a diversity of differing individual goals and serendipitous discovery (Drexler and Minsky, 1986; Kelly, 2010, pp. 165–166).

Finally, an interesting, well-studied microcosm of open-ended innovation is provided by an online system called Picbreeder (Secretan et al., 2011) wherein users interactively evolve pictures that are represented as compositions of mathematical functions. During evolution, a user can publish a particular image to the Picbreeder website, where other users can see and rate it. Users can evolve images starting from a random initial population or they can start instead from any one of the images already published. Most evolution in this system happens through users branching from already-evolved pictures because they are more complex and visually appealing than random images. Thus, branching in Picbreeder fosters a collaborative system that leads to an accumulation of diverse, complex pictures. It is important to note that there is no overall drive to the system besides the wide-ranging individual preferences of the users.

Though there is no system-wide goal and no bias in the encoding towards particular classes of images, surprisingly, many pictures resembling real-world phenomena such as faces, cars, and butterflies have been evolved. That is, through collaborative interactive evolution users have discovered mathematical representations of recognizable images. Creating an interactive evolution system that can discover recognizable images is difficult, and is rare among such systems. Because the evolutionary history of all images is preserved by Picbreeder, one can trace the ancestors of such complex images all the way to their random origins in an initial population. Whereas one might guess that users discovered these images by intentionally selecting images that resemble them, interestingly, that is not the case. In fact, for most cases of complex images, the nearly-immediate predecessors to what look like particular real-world objects do *not* resemble that same object (figure 1-1). That is, the precursors to an image resembling a car were not chosen because they were car-like, but for some other aesthetic merit, mirroring biological exaptation. In fact, users are often frustrated when they try and fail to evolve towards a specific image class (Secretan et al., 2011), yet those image classes are still discovered – but

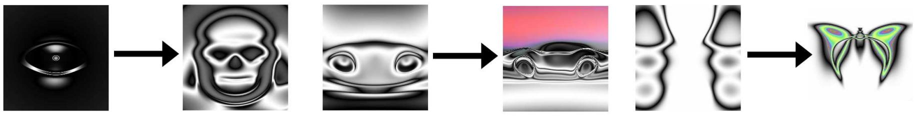

Figure 1-1. **Deceptive Precursors.** Three pairs of related images evolved by users on Picbreeder are shown above. Each left image is a close evolutionary ancestor of the image to its right. The important insight is that the precursor pictures that are stepping stones to a particular image often do not resemble that final image.

only when discovering them is *not* the goal. In other words, the success of Picbreeder at finding so many recognizable images results from of its lack of an overarching goal. Furthermore, (Woolley and Stanley, 2011) shows that pictures evolved on Picbreeder cannot be re-evolved effectively by the same algorithm and encoding inside Picbreeder if those pictures are set as objectives for evolution.

The implication from reviewing these examples of non-objective searches is that the types of systems that foster open-ended innovation and lead to the accumulation of complex interesting artifacts tend to lack singular fixed objectives. While the evidence presented thus far has been inductive, the next section reviews novelty search, a non-objective search algorithm that can be quantitatively compared in various domains to more traditional objective-based search. Thus novelty search provides an opportunity to test directly whether abandoning the single-minded search for the objective is ever beneficial.

4. Novelty Search

Recall that the problem with the objective-based search paradigm that is common in EC models is that the objective function (e.g. the fitness function) does not necessarily reward the intermediate stepping stones that lead to the objective. These stepping stones often do not resemble the objective itself, especially as objectives become more ambitious, which makes it increasingly difficult to identify these stepping stones *a priori*.

The approach in novelty search (Lehman and Stanley, 2008; Lehman and Stanley, 2011), is to identify novelty as a *proxy* for stepping stones. That is, instead of searching for a final objective, the learning method rewards instances with functionality significantly different from what has been discovered before. Thus, instead of a traditional objective function, evolution employs a *novelty metric*. That way, no attempt is made to measure overall progress. In effect, such a process performs explicitly what natural evolution does passively, i.e. gradually accumulating novel forms that ascend the complexity ladder.

For example, in a biped locomotion domain, initial attempts might simply fall down. An objective function may explicitly *reward* falling the farthest, which is unrelated to the ultimate objective of walking and thus exemplifies a deceptive local optimum. In contrast, the novelty metric would reward simply falling down in a different way, regardless of whether it is closer to the objective behavior or not. After a few ways to fall are discovered, the only way to be rewarded is to find a behavior that does *not* fall right away. In this way, behavioral complexity rises from the bottom up. Eventually, to do something new, the biped would have to successfully walk for some distance even though it is not an objective.

Novelty search succeeds where objective-based search fails by rewarding the stepping stones. That is, anything that is genuinely different is rewarded and promoted as a jumping-off point for further evolution. While we cannot know which stepping stones are the right ones, if we accept that the primary pathology in objective-based search is that it cannot detect the stepping stones at all, then that pathology is remedied. This idea is also related to research in *curiosity seeking* in reinforcement learning (Schmidhuber, 2006).

Evolutionary algorithms like GP or neuroevolution (Yao, 1999) are well-suited to novelty search because the population of genomes that is central to such algorithms naturally covers a wide range of expanding behaviors. In fact, tracking novelty requires little change to any evolutionary algorithm aside from replacing the fitness function with a *novelty metric*.

The novelty metric 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, the novelty search rewards diverging from prior behaviors. Therefore, novelty needs to be *measured*. There are many potential ways to measure novelty by analyzing and quantifying behaviors to characterize their differences. Importantly, like the fitness function, this measure must be fitted to the domain.

The novelty of a newly-generated individual is computed with respect to the *behaviors* (i.e. *not* the genotypes) of an *archive* of past individuals and the current population. The aim is 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

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} \mathrm{dist}(x, \mu_i), \tag{1.1}$$

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. Note that this novelty space cannot be explored purposefully; 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 explicitly-specified objective within the search space.

Novelty search in general allows any behavior characterization and any novelty metric. Although generally applicable, novelty search is particularly suited to domains with deceptive fitness landscapes, intuitive behavioral characterization, and domain constraints on possible expressible behaviors. More generally, novelty search can be applied even when the experimenter has no clear objective in mind at all. For example, in some domains, rather than optimality, the aim may be to collect all the interesting behaviors in the space.

Once objective-based fitness is replaced with novelty, the underlying EA operates as normal, selecting the highest-scoring individuals to reproduce. Over generations, the population spreads out across the space of possible behaviors, continually ascending to new levels of complexity to create novel behaviors as the simpler variants are exhausted.

5. Experiments with Novelty Search

There have been many successful applications of novelty search in EC (Lehman and Stanley, 2011; Risi et al., 2010; Goldsby and Cheng, 2010; Mouret, 2009; Lehman and Stanley, 2008; Lehman and Stanley, 2010a; Lehman and Stanley, 2010b; Doucette, 2010), both with GP (Doucette, 2010; Lehman and Stanley, 2010a; Goldsby and Cheng, 2010) and neuroevolution (Lehman and Stanley, 2011; Risi et al., 2010; Mouret, 2009; Lehman and Stanley, 2010b). This section reviews some such results to provide evidence that search can indeed function effectively without an explicit objective.

Novelty search was first introduced in a conference paper in 2008 (Lehman and Stanley, 2008) in which it was combined with the NEAT neuroevolution method (Stanley and Miikkulainen, 2002) and tested in a deceptive maze-navigation domain. In the harder of the two tested mazes, novelty search solved the maze in 39 out of 40 attempts (even though solving the maze was *not* the objective), while objective-based search nearly always failed (succeeding only three times out of 40 even though solving the maze *was* the objective). These results were also reproduced in combination with a multi-objective EA (Mouret, 2009). Novelty-related methods have also been shown beneficial in evolving plastic neural networks that learn from experience (Risi et al., 2010).

Novelty search was further applied to biped locomotion (Lehman and Stanley, 2011), a difficult control task that is popular within machine learning (Reil and Husbands, 2002). Though it was not looking directly for stable gaits, novelty search evolved controllers that traveled farther (4.04 meters, $sd = 2.57$) than solutions evolved by objective-based search (2.88 meters, $sd = 1.04$) on average over 50 runs of both methods. More dramatically, the *best* gait discovered by novelty search traveled 13.7 meters, while the best gait discovered by objective-based search traveled only 6.8 meters.

In GP, novelty search has worked successfully in the artificial ant benchmark (Lehman and Stanley, 2010a; Doucette, 2010), maze navigation (Lehman and Stanley, 2010a; Doucette, 2010), and in finding latent bugs in software models (Goldsby and Cheng, 2010). Novelty search with GP has outperformed standard objective-based search (Lehman and Stanley, 2010a; Doucette, 2010), proven less prone to program bloat (Lehman and Stanley, 2010a), and found more general solutions than objective-based search (Doucette, 2010).

Building on prior results in maze navigation with GP (Lehman and Stanley, 2010a; Doucette, 2010), the next section describes an experiment that investigates how the performance of novelty search and traditional objective-based search degrade with increasing problem complexity.

6. Scaling Problem Complexity in Maze Navigation

A hypothesis advanced by this chapter is that as problems grow more difficult, the gradient defined by measuring distance to the objective becomes increasingly deceptive and thereby less informative. Thus as deceptiveness increases, non-objective search methods like novelty search may outperform more traditional objective-based search methods. However, while not susceptible to traditional deception, novelty search also is not guaranteed to consistently find *specific* objectives as problems become more complex.

Therefore, an interesting experiment is to compare how the relationship between problem complexity and performance varies in both traditional objective-based search and novelty search, which serves as an example of a non-objective search algorithm. Maze navigation is a natural choice of domain for such an investigation because it is a good model for search problems in general (Lehman and Stanley, 2011), because it is the basis for previous comparisons between novelty search and objective-based search (Lehman and Stan-

Objective:	Find a robot that navigates the maze
Terminal set:	Left (turn left), Right (turn right), Move (move forward one square)
Functions set:	IfWallAhead (execute one of two child instructions based on whether there is a wall directly ahead), IfGoalAhead (execute one of two child instructions based on whether the goal is within a 90 degree cone projected outwards from where the robot is facing), Prog2 (sequentially execute the two child instructions)
Fitness cases:	One of 360 randomly-generated mazes
Wrapper:	Program repeatedly executed for 200 time steps
Population Size:	500
Termination:	Maximum number of generations = 200, 400 and 600

Table 1-1. Parameters for the Maze Problem

ley, 2010a; Doucette, 2010; Lehman and Stanley, 2011; Mouret, 2009), and because it is easy to generate mazes of parameterized complexity.

Experiment Description

The GP maze domain works as follows. A robot controlled by a genetic program must navigate from a starting point to an end point within a fixed amount of time. The task is complicated by occlusions and cul-de-sacs that prevent a direct route and create local optima in the fitness landscape. The robot can move forward, turn, and act conditionally based on whether there is a wall directly in front of it or not, or whether it is facing the general direction of the goal or not. The robot is successful in the task if it reaches the goal location. This setup is similar to previous GP maze navigation experiments (Lehman and Stanley, 2010a; Doucette, 2010). Table 1-1 describes the parameters of the experiment.

Objective fitness-based GP, which will be compared to novelty search, requires a fitness function to reward maze-navigating robots. Because the objective is to reach the goal, the fitness f is defined as the distance from the robot to the goal at the end of an evaluation: $f = b_f - d_g$, where b_f is the maximum distance possible and d_g is the distance from the robot to the goal. Given a maze with no deceptive obstacles, this fitness function defines a monotonic gradient for search. The constant b_f ensures that all individuals will have positive fitness.

GP with novelty search, on the other hand, requires a novelty metric to distinguish between maze-navigating robots. Defining the novelty metric requires careful consideration because it biases the search in a fundamentally different way than the fitness function. The novelty metric determines the behavior-space through which search will proceed. It is important that the type of behaviors that one hopes to distinguish are recognized by the metric.

As in prior maze navigation experiments (Lehman and Stanley, 2011; Lehman and Stanley, 2010a), the behavior of a navigator is defined as its ending position. The novelty metric is then the Euclidean distance between the ending positions of two individuals. For example, two robots stuck in the same corner appear similar, while one robot that simply sits at the start position looks very different from one that reaches the goal, though they are both equally viable to the novelty metric.

To compare how effectively fitness-based search and novelty search evolve navigational policies for increasingly complex maze problems, both search methods were tested on 360 randomly-generated mazes. These mazes were created by a recursive division algorithm (Reynolds, 2010), which divides an initally empty maze (i.e. without any interior walls) into two subareas by randomly adding a horizonal or vertical wall with a single randomly-located hole in it (which makes all open points reachable from any other open point in the maze.) This process continues recursively within each subarea until no areas can be further subdivided without making the maze untraversable, or until the limit for subdivisions (chosen randomly between 2 and 50 for each maze in this experiment) is exceeded. The starting position of the maze navigating robot and the goal position it is trying to reach are also chosen randomly. Examples of mazes generated by such recursive division are shown in figure 1-2.

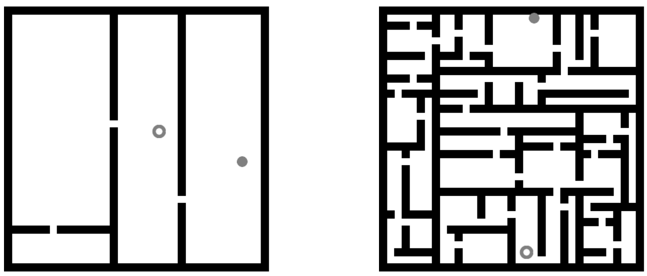

(a) Simple Maze Problem

(b) More Complex Maze Problem

Figure 1-2. **Two randomly-generated maze problems created by the recursive division algorithm.** In both mazes, the filled circle represents the starting location and the unfilled circle represents the goal location. The maze shown in (a) has fewer subdivisions and a shorter optimal path to the goal than the maze shown in (b).

The length of the shortest possible path between the start and goal position was found to be a good heuristic for problem complexity. Intuitively, longer paths potentially require more complex navigational policies. In addition, increasing path length was highly correlated with decreasing performance for all of the search methods (adjusted $R^2 > 0.75$ for each method). Thus mazes were sampled such that 4 maze problems were chosen for each shortest-path length between 10 and 100. For each of the 360 mazes, 10 independent runs were conducted for both fitness-based search, novelty search, and GP with random selection. Random selection was considered as a control to differentiate novelty search from random exploration of the search space. Experiments were conducted with limits of 200, 400, and 600 generations. A given run is considered successful if a navigator was evolved that reaches the goal within the time limit of 200 steps.

Results

The main result, as illustrated by figures 1-3a and 1-3b, is that novelty search solves significantly more instances of the generated maze problems ($p < 0.001$, Fischer's exact test) and that it scales to solving more complex instances significantly better than objective fitness-based search or random search ($p < 0.001$, the intercept values of the linear regression models are significantly different according to an ANCOVA test.) In addition, figure 1-3b shows that novelty search better exploits additional evaluations than fitness-based search or random search. While random search may waste many evaluations with policies that are the same and fitness-based search may waste many evaluations attempting to escape from deceptive local optima, novelty search constantly incentivizes discovering new behaviors.

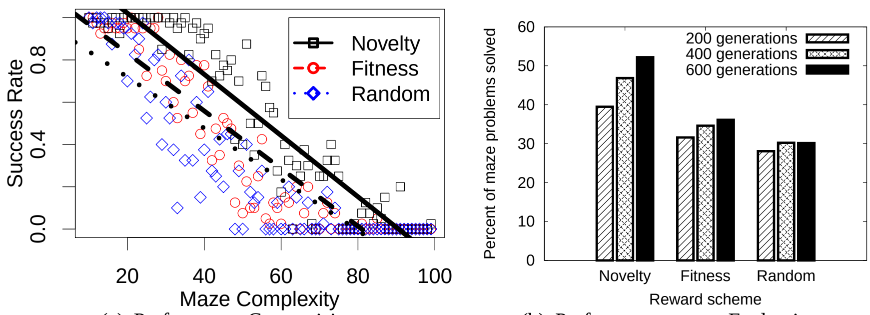

(a) Performance Comparision (b) Performance versus Evaluations

Figure 1-3. **Reward scheme comparisons.** The effectiveness of novelty search, fitness-based search, and random search at solving problems of increasing complexity is plotted along with linear regressions in (a). Novelty search is the most effective although its performance also degrades with increasing problem complexity. Each plotted point is from ten runs, each lasting 600 generations. The effect on performance of varying the amount of generations for novelty search, fitness-based search, and random search is shown in (b). Novelty search exploits additional evaluations more effectively than fitness-based search or random search.

It is important to note that the performance of each of the three compared methods decreases with increasingly complex maze instances. Few instances are reliably solved by any of the methods with optimal path length greater than 80. Thus while novelty search may outperform the other methods in this domain, it too struggles to discover *specific* ambitious objectives from first principles; this result tentatively supports the hypothesis that in some cases attempting to achieve specific objectives from a random starting point may ultimately be futile beyond a certain level of problem complexity.

This experiment and the previous section (which enumerated prior results with novelty search) offer strong empirical evidence that novelty search may often be a viable approach. Though such evidence of its effectiveness continues to accumulate, because it challenges the intuitive assumption that search should be driven by an objective, some skepticism of novelty search is a natural response, as addressed next.

7. Common Objections to Novelty Search

Common objections to novelty search include: (1) that it is not general, (2) that it is ineffective if the space of behaviors to be explored becomes vast, (3) that maintaining a growing archive of past behaviors is too computationally expensive, and (4) that novelty itself is an objective and thus that novelty search is still an objective-based search.

7.1 Generality

Because maze-navigation is easily understood and makes a good model for search in general, it has often been used to illustrate properties of novelty search (Lehman and Stanley, 2010a; Lehman and Stanley, 2008; Lehman and Stanley, 2011; Mouret, 2009; Lehman and Stanley, 2010b). However, novelty search has also been applied to diverse domains such as biped walking (Lehman and Stanley, 2011), discovering latent bugs in software models (Goldsby and Cheng, 2010), T-mazes common in learning experiments with rats (Risi et al., 2010), and the artificial ant GP benchmark (Lehman and Stanley, 2010a). It has been combined with multi-objective search (Mouret, 2009) and applied with both neuroevolution and GP. Thus, while of course it will not always succeed in every domain, evidence of its generality as a viable tool in the toolbox of EC continues to accumulate.

7.2 Dimensionality

Though smaller behavior spaces can be more thoroughly explored, novelty search incentivizes the discovery of maximally different behaviors in any size of behavior space. Thus even in vast behavior spaces novelty search will expand to cover the space loosely, uncovering a wide variety of representative behaviors, some of which may be solutions to problems. For example, even when exploring a 400-dimensional behavior space constructed from considering nearly every point in the trajectory of a maze-navigating robot, novelty search still consistently discovers neural networks that solve mazes (Lehman and Stanley, 2011). Of course, it is possible to construct scenarios in which vast sections of the behavior space are uninteresting. However, search in such cases can be accelerated by restricting which behaviors are considered novel (Lehman and Stanley, 2010b).

7.3 Costs of Expanding the Archive

The purpose of the archive of previously novel behaviors in novelty search is to discourage backtracking by keeping track of where search has previously been in the behavior space. Though the archive in general grows slowly, because it grows continually over the course of evolution, it may eventually become computationally expensive to calculate the novelty of a new individual. However, kd-trees or other specialized data structures can reduce the computational complexity of such calculations, and experiments have shown that in practice it may often be possible to limit the size of the archive without harming novelty search's performance (Lehman and Stanley, 2011).

7.4 Novelty as an Objective

While some might say that rewarding novelty effectively means that novelty is just a special kind of objective, novelty is not an objective in the usual spirit of the word in EC. That is, for many years objectives in EC have been descriptions of *areas of the search space towards which evolution should be directed*. In other words, *objectives* describe where we want search to go. Yet novelty does *not* favor any particular part of the search space; instead, it is *relative* to what behaviors have been previously encountered, is constantly changing, and is largely orthogonal to the actual objective of the experiment. Thus while the semantics of the word *objective* are likely to continue to invite debate, drawing the distinction between novelty on the one hand and traditional objectives on the other is important because the purpose of scientific language is to *facilitate* drawing such distinctions rather than to obfuscate them. In fact, it could be argued that one reason that non-objective search has received so little attention until now is that a number of different incentives for search have been conflated as being one and the same when they are in fact fundamentally different. Thus language can help to extricate us from such misleading conflations.

The next section examines the larger implications of non-objective search.

8. Implications

The success of novelty search combined with the general tendency of systems that continually innovate to lack fixed objectives is that sometimes it is beneficial to *ignore* the objective rather than to seek it. Although this insight may at first appear strange or unsettling, the evidence suggests that the underlying assumption causing this discomfort, i.e. that searching for something is always the best way to find it, does not hold. That is, novelty search provides direct evidence that searching without knowledge of the objective can sometimes provide an advantage.

It is important to note that while novelty search is a viable alternative to objective-based search and has been shown to outperform it in non-trivial domains, it is no panacea and also so far falls short of natural evolution. However, as a new tool for EC and a working example of what can be gained from considering non-objective search, novelty search may inspire future research into more powerful EAs that can achieve more by not directly trying to do so.

That is, acknowledging that the prevailing paradigm of objective-based search is not the only way nor always the best way to guide search opens our minds to the possibilities of non-objective search, of which novelty search is only one example. Other examples of non-objective search include natural evolution and the cultural evolutions of math, art, and technology. These systems continually innovate and yield artifacts of increasing complexity though they are not guided by progress towards a fixed objective. New innovations in such objective-less systems typically branch exploratively outwards from prior innovations. This dynamic contrasts starkly with the objective paradigm in EC wherein evolution in any experiment almost always starts from an unstructured, random population and then evolves narrowly to a particular goal artifact.

The problem is that we often attribute the creativity and complexity of nature to *optimization* of reproductive fitness by natural selection. Thus the optimization process has become the key abstraction driving most EAs. However, natural evolution's characteristic prolific creativity and accumulation of complexity may be natural byproducts of diversity and open-ended innovation instead of the struggle to optimize a particular metric (Gould, 1996; Lynch, 2007). That is, the driving abstraction behind EC may rest upon the wrong bedrock principle. However, alternative abstractions can easily be investigated while still preserving the main components of EC and GP algorithms by simply changing the selection criteria to be driven by something other than explicit objectives.

A narrow interpretation of this argument is that we might sometimes more effectively achieve our objectives by searching for something other than themselves. However, the implication of this body of work is actually more fundamental: There may be a trade-off in search that has to date received little attention yet that nevertheless explains why computational approaches seem so qualitatively different from processes in nature. In particular, it may be that search can either be prodded toward a specific yet not-too-ambitious objective (e.g. as with the traditional fitness function) *or* it can discover a multitude of interesting artifacts, none of which are anticipated a priori or even necessarily desired at all (e.g. as in nature). Thus, on the downside, perhaps it is not possible to *purposefully* drive search to our most ambitious objectives. However, on the upside, perhaps artificial processes *can* discover artifacts of unprecedented scope and complexity, yet only if we relinquish the insistence that we must define a priori what those discoveries should be. In effect, we might *search without objectives*. Who knows what we will find?

9. Conclusion

In conclusion, the reliance on objectives that pervades EC and GP may not capture the essential dynamics of natural evolution. Indeed, this prevalent paradigm may be preventing us from realizing computationally the profound creativity of natural evolution. That is, although natural evolution does not narrowly search for something in particular, that is how we as practitioners typically constrain evolution. Far from bearing a negative message, this paper highlights the opportunity to explore the yet untamed wilderness of non-objective search algorithms to create open-ended systems that yield a ratcheting proliferation of complex novelty.

10. Acknowledgements

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

References

Barnett, Lionel (2001). Netcrawling - optimal evolutionary search with neutral networks. In *Proc. of the 2001 IEEE Intl. Conf. on Evol. Comp.*, pages 30–37. IEEE Press.

Brockhoff, Dimo, Friedrich, Tobias, Hebbinghaus, Nils, Klein, Christian, Neumann, Frank, and Zitzler, Eckart (2007). Do additional objectives make a problem harder? In *GECCO '07: Proc. of the 9th Annual Conf. on Genetic and Evol. Comp.*, pages 765–772, New York, NY, USA. ACM.

Cliff, Dave and Miller, Geoffrey (1995). Tracking the red queen: Measurements of adaptive progress in co-evolutionary simulations. pages 200–218.

De Jong, Kenneth A. (2006). *Evolutionary Computation: A Unified Approach*. MIT Press.

Deb, Kalyanmoy (1999). Multi-objective genetic algorithms: Problem difficulties and construction of test problems. *Evol. Comp.*, 7:205–230.

Doucette, John (2010). Novelty-based fitness measures in genetic programming. Master of science in computer science, Dalhouise University.

Drexler, K.E. and Minsky, M. (1986). *Engines of creation*. Anchor Press.

Ficici, Sevan and Pollack, Jordan B. (1998). Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In *Proc. of the Sixth Intl. Conf. on Art. Life*, pages 238–247. MIT Press.

Goldberg, David E. (1987). Simple genetic algorithms and the minimal deceptive problem. In Davis, L. D., editor, *Genetic Algorithms and Simulated Annealing, Re- search Notes in Artificial Intelligence*. Morgan Kaufmann.

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

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

Gould, Steven Jay (1996). *Full House: The Spread of Excellence from Plato to Darwin*. Harmony Books.

Holland, John H. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. University of Michigan Press, Ann Arbor, MI.

Kelly, K. (2010). *What technology wants*. Viking Press.

Koza, John R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, USA.

Koza, John R., Keane, Martin A., Streeter, Matthew J., Mydlowec, William, Yu, Jessen, and Lanza, Guido (2003). *Genetic Programming IV: Routine Human-Competitive Machine Intelligence*. Kluwer Academic Publishers.

Lehman, Joel and Stanley, Kenneth O. (2008). 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. MIT Press.

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

Lehman, Joel and Stanley, Kenneth O. (2010b). Revising the evolutionary computation abstraction: Minimal criteria novelty search. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2010)*. ACM.

Lehman, Joel and Stanley, Kenneth O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evol. Comp.* To appear.

Lynch, Michael (2007). The frailty of adaptive hypotheses for the origins of organismal complexity. In *Proc Natl Acad Sci USA*, volume 104, pages 8597–8604.

Mahfoud, Samir W. (1995). *Niching methods for genetic algorithms*. PhD thesis, Champaign, IL, USA.

Mill, John Stuart (1846). *A System of Logic, Ratiocinative and Inductive*.

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

Pelikan, Martin, Pelikan, Martin, Goldberg, David E., and Goldberg, David E. (2001). Escaping hierarchical traps with competent genetic algorithms. In

Proc. of the Genetic and Evolutionary Computation Conference (GECCO-2001), pages 511–518. Morgan Kaufmann.

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

Reynolds, AM (2010). Maze-solving by chemotaxis. *Physical Review E*, 81(6).

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

Schmidhuber, J. (2006). Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18(2):173–187.

Secretan, J., Beato, N., D'Ambrosio, D.B., Rodriguez, A., Campbell, A., Folsom-Kovarik, J.T., and Stanley, K.O. (2011). Picbreeder: A Case Study in Collaborative Evolutionary Exploration of Design Space. *Evol. Comp.*

Spector, Lee, Barnum, Howard, Bernstein, Herbert J., and Swamy, Nikhil (1999). Quantum computing applications of genetic programming. In Spector, Lee, Langdon, William B., O'Reilly, Una-May, and Angeline, Peter J., editors, *Advances in Genetic Programming 3*, chapter 7, pages 135–160. MIT Press, Cambridge, MA, USA.

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

Stewart, T. C. (2001). Extrema selection: Accelerated evolution on neutral networks. In *Proc. of the 2001 IEEE Intl. Conf. on Evol. Comp.* IEEE Press.

Veldhuizen, David A. Van and Lamont, Gary B. (2000). Multiobjective evolutionary algorithms: Analyzing the state-of-the-art. *Evolutionary Computation*, 8(2):125–147. PMID: 10843518.

Weinberger, Edward (1990). Correlated and uncorrelated fitness landscapes and how to tell the difference. *Biological Cybernetics*, 63(5):325–336.

Wolpert, David H. and Macready, William (1995). No free lunch theorems for search. Technical Report SFI-TR-95-01-010, The Santa Fe Institute, Santa Fe, NM.

Woolley, Brian G. and Stanley, Kenneth O. (2011). On the deleterious effects of a priori objectives on evolution and representation. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2011)*. ACM.

Yao, Xin (1999). Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447.

Zaera, N., Cliff, D., and Bruten, J. (1996). (Not) evolving collective behaviours in synthetic fish. In *From Animals to Animats 4: Proc. of the Fourth Intl. Conf. on Simulation of Adaptive Behavior*. MIT Press Bradford Books.

Index

Ambitious objectives, 4
Artificial ant benchmark, 11
Behavior space, 9
Biped locomotion, 11
Building blocks hypothesis, 4
Coevolution, 6
Complexity, 17
Deception, 2
Diversity maintenance, 5
Exaptation, 7
Fitness landscape, 5
Lehman Joel, 1
Local optima, 5, 12
Maze navigation, 11
Multi-Objective Evolutionary Algorithms, 6
Natural evolution, 3–4, 17
NEAT, 10
Neuroevolution, 9
Neutral network, 5
Non-objective search, 3, 6, 17
Novelty metric, 8
Novelty search, 3, 8, 16
Objective-based search, 4, 17
Picbreeder, 7
Problem complexity, 15
Problem difficulty, 4
Program bloat, 11
Recursive division algorithm, 13
Serendipity, 3
Stanley Kenneth O., 1
Stepping stones, 2, 8