

Exploiting Open-Endedness to Solve Problems Through the Search for Novelty

In: *Proceedings of the Eleventh International Conference on Artificial Life (ALIFE XI)*. Cambridge, MA: MIT Press, 2008

Joel Lehman and Kenneth O. Stanley

School of Electrical Engineering and Computer Science
University of Central Florida, Orlando, FL, 32816
{jlehman,kstanley}@eecs.ucf.edu

Abstract

This paper establishes a link between the challenge of solving highly ambitious problems in machine learning and the goal of reproducing the dynamics of open-ended evolution in artificial life. A major problem with the objective function in machine learning is that through deception it may actually prevent the objective from being reached. In a similar way, selection in evolution may sometimes act to *discourage* increasing complexity. This paper proposes a single idea that both overcomes the obstacle of deception *and* suggests a simple new approach to open-ended evolution: Instead of either explicitly seeking an objective or modeling a domain to capture the open-endedness of natural evolution, the idea is to simply *search for novelty*. Even in an objective-based problem, such *novelty search* ignores the objective and searches for behavioral novelty. Yet because many points in the search space collapse to the same point in behavior space, it turns out that the search for novelty is computationally feasible. Furthermore, because there are only so many simple behaviors, the search for novelty leads to increasing complexity. In fact, on the way up the ladder of complexity, the search is likely to encounter at least one solution. In this way, by decoupling the idea of open-ended search from only artificial life worlds, the raw search for novelty can be applied to real world problems. Counterintuitively, in the deceptive maze navigation task in this paper, novelty search significantly outperforms objective-based search, suggesting a surprising new approach to machine learning.

Introduction

The problem of overcoming deception and local optima to find an objective in machine learning is not often linked to the goal of creating a truly open-ended dynamic in artificial life. Yet this paper argues that the same key idea addresses both challenges.

The concept of the objective function, which rewards getting closer to the goal, is ubiquitous in machine learning [22]. However, objective functions come with the pathology of local optima; landscapes from objective (e.g. fitness) functions are often deceptive [9, 21]. As a rule of thumb, the more ambitious the goal, the more likely it is that search can be deceived by local optima. The problem is that the objective function does not necessarily reward the *stepping stones* in the search space that ultimately lead to the objective.

For example, it is difficult to train a simulated biped without first suspending it from a string because it simply falls down on every attempt, obfuscating to the objective function any improvements in leg oscillation [30].

For these reasons, ambitious objectives are often carefully sculpted through a curriculum of graded tasks, each chosen delicately to build upon the prior [8, 10, 30]. Yet such incremental training is difficult and ad hoc, requiring intimate domain knowledge and careful oversight.

In contrast to the focus on objective optimization in machine learning, researchers in artificial life often study systems without explicit objectives, such as in *open-ended evolution*. An ambitious goal of this research is to reproduce the unbounded innovation of natural evolution. A typical approach is to create a complex artificial world in which there is no final objective other than survival and replication [4, 32]. Such models assume that biologically-inspired evolution supports creating an open-ended dynamic that leads to unbounded increasing complexity [3, 4, 16].

However, a growing yet controversial view in biology is that the drive towards complexity is a passive force, i.e. not driven primarily by selection [15, 18, 19]. In fact, in this view, the path towards complexity in natural evolution can sometimes be *inhibited* by selection pressure. Thus although open-endedness is often framed as an adaptive competition in artificial life worlds [3, 16], this paper decouples the idea of open-endedness from the domain by capitalizing on a simpler perspective: An open-ended evolutionary system is simply one that continually produces novel forms [25].

This perspective leads to a key idea that addresses the problems in both artificial life and machine learning: Instead of modeling natural evolution with the hope that novel individuals will be continually discovered, it is possible to search *directly* for novelty. Thus this paper introduces the *novelty search* algorithm, which searches with no objective other than continually finding novel behaviors in the search space. By defining novelty in this domain-independent way, novelty search can be applied to real world problems as directly as artificial life worlds. In fact, because there are only so many ways to behave, some of which must be more com-

plex than others [6], the passive force in nature that leads to increasing complexity is *accelerated* by searching for behavioral novelty.

To demonstrate the power of novelty search, in this paper it is compared to objective-based search in a deceptive two-dimensional robot maze navigation task. Counterintuitively, novelty search, which ignores the objective, evolves successful maze navigators that reach the objective in significantly fewer evaluations than the objective-based method. For harder mazes, the objective-based method almost always fails, while novelty search is successful in nearly every attempt. These results defy the premise in much of machine learning that the objective is the proper impetus for search.

The conclusion is that by abstracting the process through which natural evolution discovers novelty, it is possible to derive an open-ended search algorithm that applies naturally to *both* real-world machine learning problems and artificial life worlds. Novelty search overcomes the problems of deception and local optima inherent in objective optimization by ignoring the objective, suggesting the surprising conclusion that ignoring the objective in this way may often *benefit* the search for the objective.

Background

This section reviews open-endedness in natural evolution and evolutionary computation, as well as the neuroevolution method used in the experiments.

Open-endedness in Natural Evolution

Natural evolution fascinates practitioners of search because of its profuse creativity, lack of volitional guidance, and perhaps above all its drive towards complexity.

A subject of longstanding debate is the arrow of complexity [3, 19], i.e. the idea that evolutionary lineages sometimes tend towards increasing complexity. What about evolutionary search in nature causes complexity to increase? This question is important because the most difficult problems in search, e.g. an intelligent autonomous robot, may require discovering a prohibitive level of solution complexity.

The topic of complexity in natural evolution is much in contention across biology, artificial life, and evolutionary computation [15, 19, 23, 28]. One important question is whether there is a selective pressure towards complexity in evolution. A potentially heretical view that is gaining attention is that progress towards higher forms is not mainly a direct consequence of selection pressure, but rather an inevitable passive byproduct of random perturbations [15, 19]. Researchers like Miconi [19] in artificial life, and Lynch [14, 15] in biology are arguing that natural selection does not always explain increases in evolutionary complexity. In fact, they argue that to the extent that fitness (i.e. in nature, the ability to survive and reproduce) determines the direction of evolution, it can be deleterious to increasing complexity. In other words, rather than laying a path towards the next

major innovation, fitness (like the objective function in machine learning) in effect prunes that very path away.

In particular, Miconi [19] points out that natural selection restricts the breadth of evolution because only designs with high fitness can be further explored. Lynch [15], a biologist, goes even further, arguing that selection pressure in general does not explain innovation, and that nonadaptive processes are often undeservedly ignored.

These arguments lead to the main idea in this paper that it may be most effective to simply search explicitly for novel behaviors.

Open-Ended Evolutionary Computation

The open-ended evolution community in artificial life aims to produce simulated worlds that allow a similar degree of unconstrained exploration as Earth. Tierra [24], PolyWorld [32] and Geb [4] are typical examples. There is no objective beyond that of survival and reproduction. The motivation behind this approach is that as evolution explores an unbounded range of life forms, complexity will inevitably increase [4, 19].

Bedau and Packard [1] and Bedau et al. [2] have contributed to formalizing the notion of unbounded open-ended dynamics by deriving a test (called *activity statistics*) that classifies evolutionary systems into categories of open-endedness. Geb and others have passed this test [4, 16], but the results nevertheless do not appear to achieve the levels of diversity or complexity seen in natural evolution. This apparent deficiency raises the question of what element is missing from current models [25]. Many suggest that more detailed, lifelike domains must be constructed to facilitate the open-ended dynamic of natural evolution [20, 25, 32].

However, this paper presents a more general approach to open-ended evolution that is motivated well by the following insight from Standish [25]: “The issue of *open-ended evolution* can be summed up by asking under what conditions will an evolutionary system continue to produce novel forms.” Thus, instead of modeling natural selection, the idea in this paper is that it is more efficient to search directly for novel behaviors. While not intended to replace previous approaches to open-ended evolution, the advantage of this approach is that it decouples the concept of open-endedness from the problem domain because novelty can be sought in *any* domain. Therefore, it can apply to real-world tasks as easily as artificial life worlds.

It is important to acknowledge that this view of open-endedness contrasts with the more commonly accepted notion of prolonged production of *adaptive* traits [1, 2]. Nevertheless, the simpler view of open-endedness merits consideration on the chance that a dynamic that *appears* adaptive might be possible to capture in spirit with a simpler process.

The experiment in this paper combines this approach to open-ended evolution with the NEAT method, which is explained in the next section.

NeuroEvolution of Augmenting Topologies (NEAT)

The NEAT method was originally developed to evolve artificial neural networks (ANNs) to solve difficult control and sequential decision tasks [26, 27, 29]. Evolved ANNs control agents that select actions based on their sensory inputs. Like the SAGA method [11] 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 [17, 31], and fits well with the idea of open-ended evolution.

However, a key feature that distinguishes NEAT from prior work in complexification is its unique approach to maintaining a healthy diversity of complexifying structures simultaneously, as this section reviews. Complete descriptions of the NEAT method, including experiments confirming the contributions of its components, are available in Stanley et al. [26], Stanley and Miikkulainen [27], and Stanley and Miikkulainen [29]. Let us review the key ideas on which the basic NEAT method is based.

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 between 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 [17], 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. In effect, then, NEAT searches for a compact, appropriate network topology by incrementally complexifying existing structure.

In the experiment in this paper, NEAT is combined with novelty search, which is explained next.

The Search for Novelty

Recall that the problem identified with the objective function in machine learning is that it does not necessarily reward the intermediate stepping stones that lead to the objective. The more ambitious the objective, the harder it is to identify *a priori* these stepping stones.

The suggested approach is to identify novelty as a *proxy* for stepping stones. That is, instead of searching for a final objective, the learning method is rewarded for finding any instance whose functionality is significantly different from what has been discovered before. Thus, instead of an objective function, search 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 maze navigation domain, initial attempts might run into a wall and stop. In contrast with an objective function, the novelty metric would reward simply running into a *different* wall regardless of whether it is closer to the goal or not. In this kind of search, a set of instances are maintained that represent the most novel discoveries. Further search then jumps off from these representative behaviors. After a few ways to run into walls are discovered, the only way to be rewarded is to find a behavior that does not hit a wall right away. In this way, the complexity bucket fills from the bottom up. Eventually, to do something new, a navigator will have to successfully navigate the maze even though it is not an objective!

At first glance, this approach may seem naive. What confidence can we have that a search process can solve a problem when the objective is not provided whatsoever? Yet its appeal is that it rejects the misleading intuition that objectives are an essential means to discovery. The idea that the objective may be the enemy of progress is a bitter pill to swallow, yet if the proper stepping stones do not lie conveniently along its gradient, we must begin to leave behind its false security.

Still, what hope is there that novelty is any better when it contains no information about the direction of the solution? Is not the space of novel behaviors unboundedly vast, creating the potential for endless meandering? One might compare novelty search to exhaustive search: Of course a search that enumerates all possible solutions will eventually find the solution, but at enormous computational cost.

Yet there are good reasons to believe that novelty search is not like exhaustive search, and that in fact the number of novel behaviors is reasonable and limited in many practical domains. The main reason for optimism is that task domains on their own provide sufficient constraints on the kinds of behaviors that can exist or are meaningful, without the need for further constraint from an objective function.

For example, a robot navigating a maze can only do so many things; the robots in the experiments in this paper have only two effectors. Although the search space is effectively infinite because of NEAT's ability to add new genes, the *behavior space* into which points in the search space collapse is limited. For example, after an evaluation in the maze, a robot finishes at a specific location. Suppose the robot's behavior is characterized only by this ending location. While there are many ways to encode a policy that arrives at a particular point, under this measure of novelty, they *all* collapse to the *same* behavior. In fact, the search space collapses into a manageable number of novelty points, significantly differentiating novelty search from exhaustive enumeration.

Furthermore, novelty search succeeds where objective-based search fails by rewarding the stepping stones. That is, anything that is genuinely different is rewarded and pro-

moted 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.

The next section introduces the novelty search algorithm by replacing the objective function with the novelty metric and formalizing the concept of novelty itself.

The Novelty Search Algorithm

Evolutionary algorithms like NEAT are well-suited to novelty search because the population 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, 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 whose behaviors were highly novel when they originated. In addition, if the evolutionary algorithm is steady state (i.e. one individual is replaced at a time) then the current population can also supplement the archive by representing the most recently visited points. The aim is to characterize how far away the new individual is from the rest of the population and its predecessors in *novelty space*, i.e. the space of unique behaviors. A good metric should thus compute the *sparseness* at any point in the novelty space. Areas with denser clusters of visited points are less novel and therefore rewarded less.

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

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

where μ_i is the i th-nearest neighbor of x with respect to the distance metric *dist*, which is a domain-dependent measure of behavioral difference between two individuals in the search space. The nearest neighbors calculation must take into consideration individuals from the current population

and from the permanent archive of novel individuals. Candidates from more sparse regions of this behavioral search space then receive higher novelty scores. It is important to note that this novelty space cannot be explored purposefully, that is, it is not known *a priori* how to enter areas of low density just as it is not known *a priori* how to construct a solution close to the objective. Thus, moving through the space of novel behaviors requires exploration. In effect, because novelty is measured relative to other individuals in evolution, it is driven by a coevolutionary dynamic.

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 novelty space, similarly to archive-based approaches in coevolution [7]. 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.

It is important to note that novelty search resembles prior diversity maintenance techniques (i.e. speciation) popular in evolutionary computation. The most well known are variants of fitness sharing [5]. These also in effect open up the search by reducing selection pressure. However, in these methods, as in Hutter's fitness uniform selection [13], the heretical step of eschewing the fitness function entirely is not taken. In contrast, novelty search *only* rewards behavioral diversity with no concept of fitness or a final objective.

It is also important to note that novelty search is not a random walk; rather, it explicitly maximizes novelty. Because novelty search includes an archive that accumulates a record of where search has been, backtracking, which can happen in a random walk, is effectively avoided in behavioral spaces of any dimensionality.

The novelty search approach in general allows any behavior characterization and any novelty metric. Although generally applicable, novelty search is best suited to domains with deceptive fitness landscapes, intuitive behavioral characterization, and domain constraints on possible expressible behaviors. Changing the way the behavior space is characterized and the way characterizations are compared will lead to different search dynamics, similarly to how researchers now change the objective function to improve the search. The intent is not to imply that setting up novelty search is *easier* than objective-based search. Rather, once novelty search is set up, the hope is that it can find solutions beyond what even a sophisticated objective-based search can currently discover. Thus, the effort is justified in its returns.

Once objective-based fitness is replaced with novelty, the NEAT algorithm 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 (i.e. by




Figure 1: **Maze Navigating Robot.** The artificial neural network that controls the maze navigating robot is shown in (a). The layout of the 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.

expanding the neural networks in NEAT) to create novel behaviors as the simpler variants are exhausted.

Experiment

A good domain for testing novelty search should have a deceptive fitness landscape. In such a domain, the search algorithm following the fitness gradient may perform worse than an algorithm following novelty gradients because novelty cannot be deceived; it ignores fitness entirely. A compelling, easily-visualized domain with this property is a two-dimensional maze navigation task. A reasonable fitness function for such a domain is how close the maze navigator is to the goal at the end of the evaluation. Thus, dead ends that lead close to the goal are local optima to which an objective-based algorithm may converge, which makes a good model for deceptive problems in general.

This paper’s experiments utilize NEAT, which has been proven in many control tasks [27, 29], including maze navigation [26], the domain of the experiments in this paper.

The maze domain works as follows. A robot controlled by an ANN must navigate from a starting point to an end point in a fixed time. The task is complicated by cul-de-sacs that prevent a direct route and that create local optima in the fitness landscape. The robot (figure 1) has six rangefinders that indicate the distance to the nearest obstacle and four pie-slice radar sensors that fire when the goal is within the pie-slice. The robot’s two effectors result in forces that respectively turn and propel the robot. This setup is similar to the successful maze navigating robots in NERO [26].

Two maps are designed to compare the performance of NEAT with fitness-based search and NEAT with novelty search. The first (figure 2a) has deceptive dead ends that lead the robot close to the goal. To achieve a higher fitness than the local optimum provided by a dead end, the robot must travel part of the way through a more difficult path that requires a weaving motion. The second maze (figure 2b) provides a more deceptive fitness landscape that requires the search algorithm to explore areas of significantly lower fitness before finding the global optimum (which is a network that reaches the goal).




Figure 2: **Maze Navigation Maps.** In both maps, the large circle represents the starting position of the robot and the small circle represents the goal. Cul-de-sacs in both maps that lead toward the goal create the potential for deception.

Fitness-based NEAT, 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 a constant bias 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 to follow. The constant b_f ensures all individuals will have positive fitness.

NEAT 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.

Thus, for the maze domain, 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.

This novelty metric rewards the robot for ending in a place where none have ended before; the method of traversal is ignored. This measure reflects that what is important is reaching a certain location (i.e. the goal) rather than the method of locomotion. Thus, although the novelty metric has no knowledge of the final goal, a solution that reaches the goal will appear novel. Furthermore, the comparison between fitness-based and novelty-based search is fair because both scores are computed only based on the distance of the final position of the robot from other points.

Finally, to confirm that novelty search is indeed not anything like random search, NEAT is also tested with a random fitness assigned to every individual regardless of performance, which means that selection is random. If the maze is solved, the number of evaluations is recorded.

Experimental Parameters

Because NEAT with novelty search differs from original NEAT only in substituting a novelty metric for




Figure 3: Comparing Novelty Search to Fitness-based Search. The change in fitness over time (i.e. number of evaluations) is shown for NEAT with novelty search, fitness-based NEAT, and NEAT with random selection on the medium (a) and hard (b) maps, both averaged over 40 runs of each approach. The horizontal line indicates at what fitness the maze is solved. The main result is that novelty search is significantly more effective. Only the first 75,000 evaluations (out of 250,000) are shown because the dynamics remain stable after that point.

a fitness function, it uses the same parameters. All experiments were run using a modified version of the real-time NEAT (rtNEAT) package (available from <http://www.cs.utexas.edu/users/nn/keyword?rtneat>) with a population size of 250. The steady-state rtNEAT evolutionary algorithm performs equivalently to generational NEAT [26]. Offspring had a 0.5% chance of adding a node, a 10% chance of adding a link, and the weight mutation power is 0.8. Parameter settings are based on standard NEAT defaults and were found to be robust to moderate variation. Runs consisted of 250,000 evaluations, which is equivalent to 1,000 generations of 250 individuals in a generational evolutionary algorithm.

The number of nearest neighbors checked in novelty search, k , was set to 15, and is robust to moderate variation. The minimum threshold of novelty before adding to the permanent archive of points, ρ_{min} , was initialized to 3.0, but changed dynamically: If 2,500 evaluations pass, and no new individuals have been added to the archive, the threshold is lowered by 5%. If over four are added in the same amount of evaluations, it is raised by 20%. In addition, any evaluated point has a 0.1% chance to be added to the archive.

A robot is allotted 400 timesteps to navigate through a maze. This number was chosen experimentally to make navigation more difficult; because time is limited, the robot must make efficient movements to reach the goal. The fitness bias f_b was 300.0, which ensures that a positive fitness is awarded to all individuals.

Results

On both maps, a robot that finishes within five units of the goal counts as a solution. On the medium map, both fitness-based NEAT and NEAT with novelty search were able to evolve solutions in every run (figure 3a). Novelty search took on average 18,274 evaluations ($sd = 20,447$) to reach a solution, while fitness-based NEAT was three times

slower, taking 56,334 evaluations ($sd = 48,705$), averaged over 40 runs. This difference is significant ($p < .0001$). NEAT with random selection performed much worse than the other two methods, finding successful navigators in only 21 out of 40 runs, which confirms the difference between novelty search and random search.

Interestingly, the genomic complexity of solutions evolved by fitness-based NEAT for the medium map (66.74 connections, $sd = 56.7$) was almost three times greater ($p < 0.05$) than those evolved by NEAT with novelty search (24.6 connections, $sd = 4.59$), even though both share the same parameters.

On the hard map, fitness-based NEAT was only able to evolve a successful navigator in three out of 40 runs, while NEAT with random selection fared marginally better, succeeding in four out of 40 runs, showing that deception in this map renders the gradient of fitness no more helpful than random search. However, novelty search was able to solve the same map in 39 out of 40 runs, in 35,109 evaluations ($sd = 30,236$) on average when successful, using 33.46 connections on average ($sd = 9.26$). Figure 3b shows this more dramatic divergence. Remarkably, because the second maze is so deceptive, the same rtNEAT algorithm can almost never solve it when solving the maze is made the explicit objective, yet solves it almost every time when finding novel behavior is the objective!

Typical Behavior

Figure 4 depicts behaviors (represented as the final point visited by an individual) discovered during typical runs of NEAT with novelty search and fitness-based NEAT on each map. Novelty search exhibits a more even distribution of points throughout both mazes. Fitness-based NEAT shows areas of density around local optima in the maze.

The typical behavior of a successful robot on either maze was to directly traverse the maze for both methods.

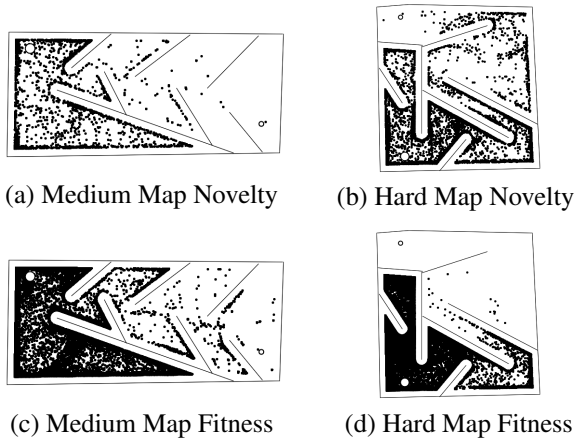


Figure 4: **Final Points Visited Over Typical Runs.** Each maze depicts a typical run, stopping at either 250,000 evaluations or when a solution is found. Each point represents the end location of a robot evaluated during the run. Novelty search is more evenly distributed because it is not deceived.

Discussion and Future Work

Novelty search casts the performance of evolutionary algorithms in a new perspective. Based on the performance of fitness-based NEAT on the second maze, one might conclude that NEAT is no better than random search for solving this particular problem. Yet NEAT-based novelty search, which changes the reward function to *ignore* the objective while preserving the rest of the algorithm, shows that the pathology is *not* in NEAT but rather in the pursuit of the objective itself. In fact, the second maze is consistently solved by NEAT when it is given no specific objective other than to produce individuals that are different functionally from those seen before. It is thus when NEAT is charged simply with continually searching for something new that it effectively solves the problem.

However, novelty search has limitations as well; because it ignores the objective, there is no bias towards optimization once a solution is found. An optimized solution may be produced by novelty search only if an individual can appear novel by reaching such performance. However, it is likely more efficient to take the most promising results from novelty search and further optimize them based on an objective function. This idea exploits the strengths of both approaches: Novelty search effectively finds approximate solutions, while objective optimization is good for *tuning* approximate solutions. Alternatively, novelty search could be applied when a traditional evolutionary algorithm converges, to replenish diversity in the population. These ideas for combining novelty and fitness-based search will be explored in future experimentation.

While glaringly counterintuitive, the idea that the search for novelty can outperform the search for the objective introduces critical insight: Objective fitness by necessity instantiates an imposing landscape of peaks and valleys. For complex problems it may be impossible to define an objective function where these peaks and valleys create a direct

route through the search space. Yet in novelty search, the rugged landscape evaporates into an intricate web of paths leading from one idea to another; the concepts of higher and lower ground are replaced by an agnostic landscape that points only along the gradient of novelty.

This idea further hints at a novel perspective on open-endedness that is fitness-agnostic. Rather than viewing open-ended evolution as an adaptive competition, it can be viewed simply as a passive drift through the lattice of novelty. As Lynch [15] and Miconi [19] suggest, it is often when the reigns of selection pressure are *lifted* that evolution innovates most prolifically. Novelty search is simply an accelerated version of this passive force in natural evolution; unlike in nature it *explicitly* rewards drifting away in the phenotype/behavior space, thereby pushing the innovating process ahead. While this perspective bypasses a longstanding notion of adaptive innovation in open-ended evolution [3, 4, 16], it offers a complementary view that is recommended by its intuitive simplicity: Open-endedness can be defined simply as the continual production of novelty.

The benefit of this view is that it means that we can now endow *any* domain with this kind of open-endedness. No longer are we restricted to complex artificial life worlds in our pursuit of open-ended discovery. As long as novelty can be defined (which will not always be easy), it can be sought explicitly in every domain from simple XOR to the most complex artificial world, putting many practical problems in machine learning within its reach.

For example, it is difficult to evolve a checkers player from scratch against a fixed world-class opponent because early generation individuals are always completely defeated. Yet novelty search abandons the idea that winning is the goal; rather it can simply try to *lose in a different way*. As the approaches to losing are exhausted one by one, eventually it will cross the path to winning, avoiding all deception and providing an entirely new kind of practical search that is nevertheless open-ended.

In addition, in the context of artificial life, it is interesting to consider how novelty search relates to natural evolution. Novelty is preserved in nature as long as a novel individual meets minimal selection criteria. It is also encouraged through niching. Moreover, there is evidence of active novelty search in natural evolution as well: Intersexual selection sometimes biases mate choice towards novelty [12]. Thus it is not unreasonable to view natural evolution as a kind of novelty search in addition to an adaptive competition.

Finally, novelty search provides a new hope for an artificial arrow of complexity. For, as Dawkins has said [6], once all the simple ways to live have been exhausted, the only way to do anything different is to become more complex. In a passive way, this idea explains the arrow of complexity in nature. In novelty search, the principle should also hold true.

In fact, the result that the solutions to the medium maze discovered by NEAT with novelty search contain almost

three times fewer connections than those discovered by fitness-based NEAT suggests that novelty search climbs the ladder of complexity *more efficiently*. While this intriguing result merits further study, a possible explanation is that compact solutions are often missed by objective-based search because they are hidden behind deceptive landscapes. Novelty search is more likely to encounter the most compact solution on the way up the ladder of complexity because it is not susceptible to such deception.

The problem with the objective is that it fails to identify the stepping stones. The more ambitious and complex the problem, the more difficult it is to formalize an objective that rewards the stepping stones along the way. Yet it is exactly those stepping stones that ultimately must be identified and rewarded if search is to find its way up the ladder of complexity [19]. Novelty search is designed to build gradients that lead to stepping stones. By abandoning the objective, all the steps along the way come into greater focus. While the trade-off is a more expansive search, it is better to search far and wide and eventually reach a summit than to search narrowly and single-mindedly yet never come close.

The implications of this approach are far-reaching because it is relevant to all of machine learning. The idea that search is more effective without an objective challenges fundamental assumptions and common intuitions about why search works. It is also the first machine learning approach to take seriously the growing (yet controversial) consensus in biology and artificial life that adaptive selection does not explain the arrow of complexity in nature [14, 15]. Novelty search asks what is left if the pressure to achieve the objective is abandoned. Thus its potential reach is broad. Furthermore, the implication for artificial life is that the adaptive competition is not necessary to promote an open-ended dynamic, suggesting a new approach to modeling evolution in artificial worlds.

In summary, almost like a riddle, novelty search suggests a surprising new perspective on achievement: *To achieve your highest goals, you must be willing to abandon them.*

Conclusions

This paper introduced novelty search, a domain-independent method of open-ended search. Motivated both by the problem of deceptive gradients in objective-based search and the desire for a simple approach to open-ended evolution, novelty search ignores the objective and instead searches only for individuals with novel behaviors. Counterintuitively, experiments in a deceptive navigation task showed that novelty search can significantly outperform objective-based search. Novelty search thus makes it possible to effectively apply the power of open-ended search to real-world problems.

Acknowledgements

This research was supported in part by a gift from Dr. Charles Bailey. Special thanks to Li Yuan for her support in preliminary investigations.

References

- [1] M. A. Bedau and N. H. Packard. Measurement of evolutionary activity, teleology, and life. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Proc. of Art. Life II*, pages 431–461, Redwood City, CA, 1991. Addison-Wesley.
- [2] M. A. Bedau, E. Snyder, and N. H. Packard. A classification of longterm evolutionary dynamics. In C. Adami, R. Belew, H. Kitano, and C. Taylor, editors, *Proc. of Art. Life VI*, pages 228–237, Cambridge, MA, 1998. MIT Press.
- [3] Mark Bedau. Four puzzles about life. *Artificial Life*, 4:125–140, 1998.
- [4] A. Channon. Passing the alife test: Activity statistics classify evolution in geb as unbounded. In *Proceedings of the European Conference on Artificial Life (ECAL-2001)*. Springer, 2001.
- [5] Paul Darwen and Yin Yao. Every niching method has its niche: Fitness sharing and implicit sharing compared. In Hans-Michael Voigt, Werner Ebeling, Ingo Rechenberg, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature – PPSN IV*, pages 398–407, Berlin, 1996. Springer.
- [6] Richard Dawkins. Genetic and evolutionary computation conference (GECCO-2007) Keynote Debate, July 2007.
- [7] E. D. De Jong. The incremental pareto-coevolution archive. In *Proc. of the Genetic and Evol. Comp. Conf. (GECCO-2004)*, Berlin, 2004. Springer Verlag.
- [8] Jeffrey L. Elman. Incremental learning, or the importance of starting small. Technical Report 9101, CRL, La Jolla, CA, 1991.
- [9] David E. Goldberg. Simple genetic algorithms and the minimal deceptive problem. In L. D. Davis, editor, *Genetic Algorithms and Simulated Annealing, Research Notes in Artificial Intelligence*. Morgan Kaufmann, 2007.
- [10] Faustino Gomez and Risto Miikkulainen. Incremental evolution of complex general behavior. *Adaptive Behavior*, 5:317–342, 1997.
- [11] Inman Harvey. *The Artificial Evolution of Adaptive Behavior*. PhD thesis, School of Cognitive and Computing Sciences, U. of Sussex, Sussex, 1993.
- [12] Kimberly A. Hughes, Linh Du, F. Helen Rodd, and David N. Reznick. Familiarity leads to female mate preference for novel males in the guppy, *Poecilia reticulata*. *Animal Behavior*, 58(4):907–916, 1999.
- [13] Marcus Hutter and Shane Legg. Fitness uniform optimization. *IEEE Transactions on Evolutionary Computation*, 10:568–589, 2006.
- [14] Michael Lynch. The evolution of genetic networks by non-adaptive processes. *Nature Reviews Genetics*, 8:803–813, 2007.
- [15] Michael Lynch. The frailty of adaptive hypotheses for the origins of organismal complexity. In *Proc Natl Acad Sci USA*, volume 104, pages 8597–8604, 2007.
- [16] C. C. Maley. Four steps toward open-ended evolution. In *Proc. of the Genetic and Evol. Comp. Conf. (GECCO-1999)*, pages 1336–1343, San Francisco, 1999. Kaufmann.
- [17] Andrew P. Martin. Increasing genomic complexity by gene duplication and the origin of vertebrates. *The American Naturalist*, 154(2):111–128, 1999.
- [18] Daniel W. McShea. Complexity and evolution: What everybody knows. *Biology and Philosophy*, 6(3):303–324, 1991.
- [19] Thomas Miconi. Evolution and complexity: The double-edged sword. *Artificial Life: Special Issue on the Evolution of Complexity*, 2007.
- [20] Thomas Miconi. *THE ROAD TO EVERYWHERE: Evolution, Complexity and Progress in Nature and in Computers*. PhD thesis, U. of Birmingham, 2007.
- [21] Melanie Mitchell, Stephanie Forrest, and John H. Holland. The royal road for genetic algorithms: Fitness landscapes and ga performance. In F. J. Varela and P. Bourguin, editors, *Proceedings of the First European Conference on Artificial Life*, Cambridge, MA, 1992. MIT Press.
- [22] Tom M. Mitchell. *Machine Learning*. McGraw-Hill, New York, 1997.
- [23] C. L. Nehaniv and J. L. Rhodes. On the manner in which biological complexity may grow. In *Math. and Comp. Biology*, volume 26 of *Lectures on Mathematics in the Life Sciences*, pages 93–102. American Mathematical Society, 1999.
- [24] T. Ray. Evolution, ecology and optimization of digital organisms. Technical Report Working paper 92-08-042, Santa Fe Institute, 1992.
- [25] Russell Standish. Open-ended artificial evolution. *International Journal of Computational Intelligence and Applications*, 3(167), 2003.
- [26] Kenneth O. Stanley, Bobby D. Bryant, and Risto Miikkulainen. Real-time neuroevolution in the NERO video game. *IEEE Trans. on Evol. Comp. Special Issue on Evolutionary Computation and Games*, 9(6):653–668, 2005.
- [27] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.
- [28] Kenneth O. Stanley and Risto Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130, 2003.
- [29] Kenneth O. Stanley and Risto Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Art. Int. Research*, 21:63–100, 2004.
- [30] 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*, pages 165–177. Springer Verlag, 1995.
- [31] James D. Watson, Nancy H. Hopkins, Jeffrey W. Roberts, Joan A. Steitz, and Alan M. Weiner. *Molecular Biology of the Gene Fourth Edition*. The Benjamin Cummings Publishing Company, Inc., Menlo Park, CA, 1987.
- [32] Larry Yaeger. Computational genetics, physiology, metabolism, neural systems, learning, vision and behavior or polyworld: Life in a new context. In C. G. Langton, editor, *Art. Life III, Proc. Vol. XVII*, pages 263–298. Addison-Wesley, 1994.