

Revising the Evolutionary Computation Abstraction: Minimal Criteria Novelty Search

In: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2010)*. New York, NY:ACM

Joel Lehman
University of Central Florida
4000 University Blvd.
Orlando, FL 32816-2362 USA
jlehman@eecs.ucf.edu

Kenneth O. Stanley
University of Central Florida
4000 University Blvd.
Orlando, FL 32816-2362 USA
kstanley@eecs.ucf.edu

ABSTRACT

Though based on abstractions of nature, current evolutionary algorithms and artificial life models lack the drive to complexity characteristic of natural evolution. Thus this paper argues that the prevalent fitness-pressure-based abstraction does not capture how natural evolution discovers complexity. Alternatively, this paper proposes that natural evolution can be abstracted as a process that discovers *many ways* to express *the same* functionality. That is, *all* successful organisms must meet the same minimal criteria of survival and reproduction. This abstraction leads to the key idea in this paper: Searching for novel ways of meeting the same minimal criteria, which is an accelerated model of this new abstraction, may be an effective search algorithm. Thus the existing novelty search method, which rewards any new behavior, is extended to enforce minimal criteria. Such *minimal criteria novelty search* prunes the space of viable behaviors and may often be more efficient than the search for novelty alone. In fact, when compared to the raw search for novelty and traditional fitness-based search in the two maze navigation experiments in this paper, minimal criteria novelty search evolves solutions more consistently. It is possible that refining the evolutionary computation abstraction in this way may lead to solving more ambitious problems and evolving more complex artificial organisms.

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

General Terms: Algorithms

Keywords: Evolution of Complexity, Artificial Life, Novelty Search, NEAT

1. INTRODUCTION

Paradoxically, although natural evolution lacks volitional guidance, its products exhibit functionality unmatched by the best efforts of human engineering. The field of biology aims primarily to decipher the vast complexity discovered

by the process of natural evolution. A deep philosophical question that has implications for researchers in evolutionary computation (EC) and artificial life, as well as in biology, is what about evolutionary search in nature has allowed such high levels of complexity to be discovered [5, 21]?

This question is important to researchers in EC because current evolutionary algorithms (EAs), although based on abstractions of natural evolution, do not exhibit the same drive towards complexity as natural evolution; EAs will fail to solve a problem if the complexity required to solve is prohibitively high (e.g. an intelligent autonomous robot) [11, 38]. Similarly, artificial life models based on evolution tend eventually to stagnate [6, 15, 30]. An important feature that may begin to explain this disparity between nature and its attempted emulation is the difference between how both methods guide search: EC models typically reward progress towards a fixed goal while natural evolution allows exploration of any design that meets the minimal criteria of survival and reproduction.

The fitness function in EC is often a rough measure of progress to a predetermined objective, as with most *objective functions* common in machine learning and optimization [22]. Optimizing such fitness functions is an abstraction of selection pressure in natural evolution; the implicit assumption is that such optimization is an *accurate* abstraction of the high-level process by which evolution discovers complex features. However, a consensus in biology is building that the pressure to maximize fitness may not be responsible for complexity growth in natural evolution [12, 18, 21].

In fact, by optimizing the fitness function in EC search often *converges* onto deceptive local optima that appear promising but from which no local step in the search space may improve the value of the fitness function. That is, because the fitness function is a heuristic, there is no guarantee that it will reward the *stepping stones* in the search space ultimately necessary to reach the objective of the search. Thus the measure in EC designed to encourage the evolution of complex features may paradoxically inhibit progress.

In part to address this problem, unlike typical EC models, an approach in neuroevolution (i.e. evolving artificial neural networks) called *novelty search* does not reward progress towards an ultimate objective; instead evolution in novelty search is open-ended, like open-ended evolutionary models in artificial life [2, 6, 15, 30]. The main idea of novelty search is to reward behavioral novelty as a proxy for the stepping stones that a fixed fitness function may fail to reward [16]. Novelty search thus circumvents the problem of deceptive local optima because it does not search for a predetermined

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'10 July 7–11, 2010, Portland, Oregon USA.

Copyright 2010 ACM 978-1-4503-0072-8/10/07 ...\$10.00.

objective, but instead only for what is behaviorally different from prior individuals. Importantly, in prior research novelty search has been shown to effectively solve deceptive EC problems [16, 23, 25].

Like novelty search, natural evolution also has no ultimate goal and tends to accumulate novel forms. However, natural evolution is not a direct or raw search for novelty either; in natural evolution there may be no direct encouragement for novelty, nor will a novel organism necessarily be reproductively viable. Instead, search in natural evolution is guided by a binary criterion applied to each organism: By logical necessity, the lineage of an organism that reproduces will continue, while the lineage of an organism that fails to reproduce will end. In other words, an interesting idea for computational purposes is to abstract natural evolution as a passive drift over all organisms meeting the *minimal criteria* of survival and reproduction.

Thus a unifying aspect of all natural life is that every successful organism from the very beginning exhibits this same critical functionality of reproduction. Furthermore, because novel organisms that meet the minimal criteria are preserved by evolution, in nature different ways of life tend to accumulate even though such accumulating diversity may not be explicitly encouraged as in novelty search. These considerations lead to a potentially useful new computational abstraction of evolution: Natural evolution can be abstracted as a process that discovers *many different ways to do the same thing*, i.e. to survive and reproduce.

Although in natural evolution the process of discovering many ways to live is generally passive, i.e. novel organisms may be preserved although they are not directly encouraged, in a computational model the process can be *accelerated* by actively rewarding novelty (as in novelty search). Additionally, while the minimal criterion in natural evolution must always be reproduction, there is no such restriction in a computational model; *any* criteria can be chosen, allowing a search subject to minimal criteria to relate to real-world problems. Thus this paper introduces *minimal criteria novelty search* (MCNS), an extension to the established novelty search method and a new artificial abstraction of natural evolution that searches for novel ways to satisfy the same minimal criteria.

In addition to exemplifying a different abstraction of natural evolution, MCNS also addresses a practical limitation of novelty search: In vast behavior spaces it may take a long time for novelty search to reach a particular goal behavior. While the temptation in such cases is to partially re-introduce the guidance of fitness, if the problem is deceptive then the false compass of fitness may just as likely impair search even more instead of help it. As a practical alternative, it may be more effective to reduce the effective size of the behavior space by enforcing minimal criteria such that an individual is not allowed to reproduce unless it meets all such minimal criteria. In contrast to searching for higher fitness, no progress is directly encouraged by enforcing minimal criteria in MCNS because such criteria are met from the first generation onwards.

To demonstrate the potential of MCNS, in this paper it is compared to traditional objective-based search and to the raw search for novelty in a robot maze navigation task and a more difficult two-point navigation task. In preliminary validation of the argument for MCNS, it solves both problems more consistently than either of the other methods.

The conclusion is that enforcing minimal criteria in the search for novelty can mitigate the problem of massive behavior spaces by exploiting a practical new abstraction of natural evolution. Because this new abstraction may coarsely capture the process by which natural evolution discovers complex organisms, MCNS may thus begin to address the complexity disparity between computational abstractions of evolution and natural evolution itself.

2. BACKGROUND

This section reviews the NEAT algorithm used in the experiments, complexity in natural evolution, and the novelty search method that will be extended with minimal criteria.

2.1 NEAT

In experiments in this paper, behaviors are evolved that are controlled by artificial neural networks (ANNs). Thus a neuroevolution (NE) method is needed to underpin these experiments. The NEAT method is appropriate because it is widely applied [1, 3, 31, 32, 34, 35, 37] and well understood.

The NEAT method was originally developed to evolve ANNs to solve difficult control and sequential decision tasks [31, 32, 34]. Evolved ANNs control agents that select actions based on their sensory inputs. Like the SAGA method [13] introduced before it, NEAT begins evolution with a population of small, simple networks and *complexifies* the network topology into diverse species over generations, leading to increasingly sophisticated behavior. A similar process of gradually adding new genes has been confirmed in natural evolution [19, 36]. This section briefly reviews the NEAT method; for comprehensive introductions see Stanley and Miikkulainen [32, 34].

To keep track of which gene is which while new genes are added, a historical marking is uniquely assigned to each new structural component. During crossover, genes with the same historical markings are aligned, producing meaningful offspring efficiently. Speciation in NEAT protects new structural innovations by reducing competition among differing structures and network complexities, thereby giving newer, more complex structures room to adjust. Networks are assigned to species based on the extent to which they share historical markings. Complexification, which resembles how genes are added over the course of natural evolution [19], is thus supported by both historical markings and speciation, allowing NEAT to establish high-level features early in evolution and then later elaborate on them. Because the approach in this paper is in part motivated by increasing complexity in nature, NEAT's principled ability to add new genes makes a good match for such an investigation.

The next section reviews the phenomenon of increasing complexity in natural evolution to begin to uncover its causes.

2.2 Complexity in Natural Evolution

The observation of complexity in natural evolution is important in both EC and artificial life because it hints at potential missing ingredients in the algorithms that we employ today that aim to abstract from nature. Our algorithms still do not display the same tendency towards long-term increasing complexity that we observe in nature.

This tendency is sometimes called the *arrow of complexity* [5, 21]. For example, it is hard to dispute that humans are more complex than their bacterial ancestors. What about evolutionary search in nature causes complexity to increase?

This question is relevant 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 debated across biology, artificial life, and evolutionary computation [4, 12, 18, 20, 21, 24, 33]. One important question is whether there is a selective pressure towards complexity in evolution (i.e. does selection drive natural evolution to create “higher” forms). While even Darwin may have believed that evolution was a progressive process (“as natural selection works solely by and for the good of each being, all corporeal and mental endowments will tend to progress towards perfection” [7, p. 489]), such beliefs, especially when human-centric (i.e. human beings are the pinnacle or inevitable result of evolution), have since been disputed [8, 12, 20, 26]. A potentially heretical view that is more recently 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 [12, 18, 21].

Researchers like Miconi [21] in artificial life, Sigmund [28] in evolutionary game theory, and Gould [12], McShea [20], and Lynch [17, 18] in biology are arguing that natural selection does not always explain increases in evolutionary complexity. In fact, some argue that to the extent that fitness (i.e. in nature, the differential ability to survive and reproduce) determines the direction of evolution, it can be deleterious to increasing complexity [18, 21, 28]. 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.

Gould [12] argues that a trend of increasing complexity is not due to selection, but simply a byproduct of a hard lower bound on the possible complexity of organisms (i.e. the minimal complexity needed for a single cell to reproduce) without a corresponding upper bound. Lynch [18], another biologist, argues that selection pressure in general does not explain innovation, and that non-adaptive processes (i.e. those not driven by selection) are often undeservedly ignored.

These arguments lead to the idea that pressure to adapt may not be the key ingredient in natural evolution that facilitates complexity growth. This insight is important because most EC models are motivated by an abstraction of natural evolution based on selection pressure. Thus in this paper an alternative abstraction of natural evolution is suggested that aims to coarsely abstract the elements of natural evolution that lead to increasingly complex purposeful behavior.

This alternate abstraction is implemented as an extension of novelty search, which is reviewed next.

2.3 Novelty Search

The problem with the objective-based search paradigm that is common in EC models is that an objective function (e.g. the fitness function) 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 approach suggested by Lehman and Stanley [16] 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*.

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 [27].

EAs such as NEAT are well-suited to novelty search because the population that is central to such algorithms naturally covers a range of 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*.

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. 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.

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

Once objective-based fitness is replaced with novelty, the underlying NEAT algorithm operates as normal, selecting the most novel individuals to reproduce. Over generations, the population spreads out across the space of possible behaviors, sometimes encountering an individual that solves the problem even though progress towards the solution is not directly rewarded.

In fact, there have been several successful applications of novelty search in neuroevolution [16, 23, 25]. Novelty search was introduced in Lehman and Stanley [16] and applied to a deceptive maze task; these results were replicated in Mouret [23] in combination with a multi-objective evolutionary algorithm. Two other investigations have also demonstrated that encouraging behavioral novelty is useful in evolving

adaptive ANNs (i.e. ANNs that learn during their lifetimes) [25, 29]. These results were surprising because they established that an algorithm with no knowledge of its objective can often outperform one specifically rewarded for achieving that objective! The culprit behind this result is deception, which is often profoundly deleterious to traditional EC.

Nevertheless, novelty search is so open-ended that the question naturally arises whether it can be more effectively constrained without losing its open-ended character. The main idea in this paper, introduced next, is such an extension of novelty search.

3. MINIMAL CRITERIA NOVELTY SEARCH

Although natural evolution is *not* overall explicitly a search for novelty, there are several mechanisms in nature that do encourage novelty, both explicitly and implicitly. Explicit novelty search in natural evolution does sometimes result from *negative frequency dependent selection*, in which phenotypically rare individuals are more successful [10, 14]. On a more fundamental level, two properties of natural evolution, *novelty preservation* and *niching*, lead to an implicit search for novelty over evolutionary time scales.

First, novelty in natural evolution is generally preserved as long as a novel organism’s lineage continues to meet the minimal criteria for selection. Second, a complementary process, niching, encourages novelty. That is, by founding a new niche (i.e. a new way to make a living), an organism may *avoid* competition and exploit untapped resources [15]. A new niche may also lead to other new niches that are only reachable from or depend upon the precursor niche. The result of preservation and niching is an accumulation of novel lifestyles; thus natural evolution can be abstracted at a high level as a passive kind of novelty search subject to minimal criteria, which is an alternative to the more common interpretation of evolution as an adaptive competition that underpins most of EC.

This view illustrates that biological evolution is not always about competition; selection merely specifies that a lineage that does not meet the minimal criterion of reproduction will end. Competition, either from its own or competing species, is only one reason that an individual may fail to meet this minimal criterion; if the environment precludes an individual from reproducing or mutations render an organism sterile, then its lineage will also lose to natural selection.

The critical functionality of all organisms is their ability to survive until they are able to reproduce. Thus, from this point of view, all successful organisms ultimately have the *same* functionality (i.e. to reproduce) but may execute this functionality in vastly *different* ways (e.g. by being a beet or a bear, but nonetheless reproducing). Furthermore, unlike the objective of a typical fitness function in EC, this functionality has existed since the very *beginning* of evolution.

Contrary to the discredited egocentric view of humanity as the pinnacle of evolution [8, 12, 20, 26], the complexity of macro-scale organisms such as humans is orthogonal to the fundamental evolutionary goal of efficient reproduction. That is, it is unnecessary for an organism to grow from a single cell (the fertilized egg) to the scale of trillions of cells (a developed human), when ultimately the organism’s contribution to its offspring is yet again only the size of one cell (i.e. a sperm or egg cell) [9, p. 259]. Thus, the complex development cycle of large animals is a kind of Rube-Goldbergian *digression*. Furthermore, when compared to

the offspring of macro-scale organisms, offspring of bacteria are more robust [12] (i.e. they are able to survive in many environments unsuitable for larger animals), and are produced much more quickly (i.e. on the order of minutes instead of years) and efficiently (i.e. without the need to build trillions of additional cells).

Hence, any *forward progress* in evolution is orthogonal to meeting the minimal criterion of reproduction. However, the diversity of living forms has increased. The reason such a diversity including inefficient replicators such as mammals continues to exist is precisely because natural evolution is *not* a free-for-all competition: There are many niches that support life, but only limited competition among them (e.g. grass does not compete with grasshoppers). In fact, filling one niche often leads implicitly to further niches that depend on the first. A proliferation of niches, as seen in natural evolution, implies a corresponding proliferation of varied lifestyles.

The conclusion is that perhaps natural evolution can be abstracted computationally as a process that finds many *novel* ways to express the *same* functionality. While in nature the search is guided only through the passive filter provided by the minimal criteria of survival and reproduction, in a computational abstraction the passive accumulation in nature can be *accelerated* by searching directly for novelty through the novelty search method. While in nature the minimal criterion is always reproduction, in a computational abstraction, the minimal criteria can be *anything*. In fact, it may be possible to choose minimal criteria in a practical domain such that the “epiphenomenal complexity” that results from novelty search subject to such criteria actually solves real-world problems.

While the choice of minimal criteria in a particular domain may require careful consideration because they add restrictions to the search space, such minimal criteria may provide a principled way to reduce the behavior space relative to novelty search alone, which may be sometimes infeasible. This perspective motivates studying the search for novelty subject to minimal criteria.

Extending novelty search to incorporate minimal criteria is straightforward. The evaluation of an individual in the domain need only additionally return whether the minimal criteria are met, thereby complementing the individual’s behavioral characterization. If the minimal criteria are not met, then novelty search simply assigns a novelty score of zero to the individual and changes the failed individual’s behavioral characterization to a dummy value that reflects its failure. For those individuals that meet the minimal criteria, novelty search operates as normal, assigning novelty scores to individuals based solely on relative behavioral novelty.

A practical concern when enforcing minimal criteria is that until an individual is found that meets the criteria, search is effectively random. That is, all individuals not meeting the minimal criteria are viewed equally as failures and only considered for reproduction if there are no individuals that do pass the criteria. Thus, if an individual meeting the minimal criteria is unlikely to appear in the initial population, it may be necessary to *seed* MCNS with a genome specifically evolved that does meet the criteria.

While enforcing minimal criteria involves only minor changes to the novelty search algorithm, the conceptual advance is significant. It is possible that natural evolution’s ability to create an explosion of *interesting* complexity and diversity

is an inevitable result of a search for many ways to do the same thing. That “something” is the minimal criteria.

The next section describes experiments in a maze navigation domain designed to test the potential of minimal criteria novelty search (MCNS).

4. EXPERIMENTS

One practical motivation for enforcing minimal criteria during a search for novelty is to reduce the size of a vast behavior space so that finding the goal behavior can become more tractable. To verify the hypothesis that MCNS can be more efficient than either the raw search for novelty or the direct search for the objective in such large spaces, the maze navigation domain introduced in Lehman and Stanley [16] is revisited with a maze containing vast uninteresting areas orthogonal to successful navigation.

Another motivation for MCNS is that, as a new abstraction of natural evolution, the search for many ways to do the same thing (upon which MCNS is based) can sometimes efficiently solve difficult problems by discovering complex policies. To explore this hypothesis, a more difficult extension of the maze navigation domain is considered in which a robot must navigate to two points instead of just the goal. This task is hard because after reaching one of the points, the robot must effectively adopt a different policy to reach the remaining point, which is non-trivial. If the minimal criterion is that an individual must always reach one of the points to be reproductively viable, MCNS will in effect search for many ways to reach that point, some of which may become more complex and eventually reach the other point as well.

This paper’s experiments utilize NEAT, which has been proven in many control tasks [3, 16, 31, 32, 34, 35], including maze navigation [16, 31], which is the basis of the experiments in this paper as well as the original novelty search paper [16]. In both experiments, NEAT with MCNS will be compared to NEAT with novelty search and traditional fitness-based NEAT.

The next sections describes these experiments in detail.

4.1 Maze Navigation Experiment

The maze navigational domain in Lehman and Stanley [16] is a good model for deception and behavior spaces in general because it is easy to understand and visualize [16, 23]. In this domain, a robot controlled by an ANN must navigate in a maze from a starting point to an end point in a fixed time. The robot (figure 1) has six rangefinders that indicate the distance to the nearest wall within the maze, 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.

If fitness is rewarded proportionally to how close the robot ends from the goal, cul-de-sacs in the maze that lead close to the goal but do not reach it are deceptive local optima to which an objective-based search may converge. Additionally, if the behavior of a robot is characterized as its final Cartesian coordinate within the maze, large empty areas of the maze represent uninteresting areas of the behavior space that a raw search for novelty may waste evaluations exploring. Thus the basic idea of this experiment is to create a map in which there exists a large area orthogonal to reaching the goal that can be pruned from the behavior space through enforcing minimal criteria.

Therefore, as a preliminary test for MCNS, this experi-




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 the goal falls within the infinite projection of that pie-slice. The solid arrow indicates the robot’s heading.




Figure 2: Navigation Maps. In both maps, unfilled circles represent the starting position of the robot and filled circles represent goal points. The maze navigation map (a) is unenclosed, which enlarges the behavior space because the robot can travel outside the maze without restriction. The two-point navigation map (b) has two goal points that a successful navigational robot must reach; the first goal point is the larger filled circle, the second is the smaller filled circle.

ment introduces a variation of the “hard maze” from Lehman and Stanley [16] that has a larger behavior space (figure 2a). The only difference from the original hard maze is that the new maze is no longer fully enclosed; a navigational robot may travel into the vast empty region outside the maze.

Problematically, a raw search for novelty may expend unnecessary effort exploring this vast empty region of the map. Additionally, because the deceptive cul-de-sacs in the maze may confound an objective-based search, fitness-based NEAT may often fail to discover a solution as well. On the other hand, if a minimal criterion for reproduction is enforced such that an individual must end within the original bounds of the maze to be viable, the reduced behavior space may be tractable for MCNS, which may thus more consistently solve the maze than the other approaches.

Fitness-based NEAT, which will be compared to novelty search and MCNS, requires a fitness function to reward maze-navigating robots. The same fitness function from the original formulation is used, which rewards a robot’s nearness to the goal at the end of an evaluation [16, 23]; that is, the fitness f of a robot is defined as: $f = b_f - d_g$, where b_f is a constant bias and d_g is the distance of the robot to the goal at the end of the evaluation. This measure reflects that the objective of the robot is to reach the goal. Furthermore, given a maze with no deceptive obstacles, this fitness function defines a monotonic gradient for search to follow.

NEAT with novelty search, on the other hand, requires a novelty metric to distinguish between maze-navigating robots.

Following the original formulation, the behavior of a robot is defined as its location in the maze at the end of the evaluation [16, 23]; this measure reflects that what is important in navigating a maze is where in the maze a robot ends up at the end of an evaluation. The novelty metric is then the squared Euclidean distance between the ending positions of two individuals.

Finally, NEAT with MCNS distinguishes between behaviors by using the same novelty metric as NEAT with novelty search, but additionally requires a set of minimal criteria to restrict the behavior space. Because behaviors that end outside the maze generally will have no relation to behaviors that navigate the maze, a simple criterion is that to be reproductively viable, an individual must end an evaluation within the maze. In this way, the behavior space is reduced to a subset that has a clear relation to successful maze navigating, although no behavior within this subset is a priori favored over another. It is important to note that because individuals that meet this particular minimal criterion are common in the initial population, it is unnecessary to seed evolution with a specifically-evolved genome.

The next section describes in detail the experiment in the two-point navigation domain.

4.2 Two-point Navigation Experiment

This second experiment elaborates on the maze navigation domain by introducing a more ambitious objective; instead of navigating to a single goal, the navigational robot must reach two separate goal points in the maze. This subtle change to the task relegates the objective behavior to a smaller portion of the behavior space, especially if both goal points are distant from each other and the starting location. It is important to note that which goal point the robot reaches first does not matter as long as both are reached.

To facilitate the robot’s ability to navigate to both goal points, the network in figure 1 is augmented with an additional set of four pie-slice radar sensors that indicate the direction of the second goal point. Also, an input for both goal points is added that is continually activated after the corresponding goal is reached, enabling the robot to potentially change its policy and proceed to the unreached goal.

A new map is introduced (figure 2b) that has two goal points; the first point is in the lower-right corner of the map while the second is in the upper-left. The maze is deceptive because there are cul-de-sacs that lead close to both goal points; these cul-de-sacs create barriers in the fitness landscape that may stymie a traditional objective-based search. In contrast, the raw search for novelty must explore a vast behavior space, and thus may often fail to find a solution in a reasonable amount of time. For example, the raw search for novelty will likely discover separate behaviors, i.e. one that leads to one goal point and one that leads to the other goal point, but may take unreasonably long to discover the highly-specific combined behavior in the vast sea of potential combinations. The main hypothesis is that only MCNS can search a tractable area of the behavior space to more consistently solve the problem.

Fitness-based NEAT, which will be compared to novelty search and MCNS, requires a fitness function to reward the navigational robots. Because the objective is to navigate to both goal points, the fitness f is awarded based on how close during an evaluation the robot approaches both goal points: $f = b_f - d_{g_1} + b_f - d_{g_2}$, where b_f is a constant bias, d_{g_1} is the

closest the robot navigated to the first goal point, and d_{g_2} is the closest the robot navigated to the second goal point.

NEAT with novelty search, on the other hand, requires a novelty metric to distinguish between navigational robots. Because the trajectory of the robot is important in the two-point navigation domain, the behavior of the robot is defined as its location in the maze sampled uniformly a fixed number of times over the course of an evaluation. The novelty metric is then the squared Euclidean distance between the position vectors of two individuals. Thus, an individual that reaches both points, although there is no direct bias to find such an individual, will appear novel.

Finally, NEAT with MCNS uses the same novelty metric as NEAT with novelty search, but additionally requires a set of minimal criteria to restrict the behavior space. If the minimal criterion is that the navigator must at least reach the first goal point, then the knowledge of how to traverse the maze and complete one of the two objectives is always maintained. MCNS will in effect search for novel ways to navigate that preserve the crucial functionality of reaching the first goal point.

By enforcing this minimal criterion, the effective behavior space is reduced to a smaller subset with clear relation to the ultimate objective. The hypothesis is that MCNS can efficiently explore such a subset to discover a behavior that reaches both goal points, while a raw search for novelty may take longer to succeed because the full behavior space is significantly larger and contains mainly behaviors *completely orthogonal* to two-point navigation.

Note that this minimal criterion, unlike the criterion in the first experiment, requires seeding MCNS with an individual that can already reach the first goal point because such behaviors are unlikely to be found in the starting population. Thus, prior to running this experiment, for each of the 100 runs a seed genome was evolved that can reach the first goal point (seen in the lower right of figure 2b) by using the raw search for novelty and the same procedure from the first experiment. To make the comparison fair, all three of the compared methods are started with the seed genome able to reach the first goal point.

4.3 Experimental Parameters

Because NEAT with MCNS, NEAT with novelty search, and fitness-based NEAT differ only in whether they enforce minimal criteria or in substituting a novelty metric for a fitness function, they use the same parameters. All experiments were run with a modified version of the real-time NEAT (rtNEAT) package (available from <http://nn.cs.utexas.edu/keyword?rtneat>) with a population size of 500. The steady-state rtNEAT evolutionary algorithm performs equivalently to generational NEAT [31].

Offspring had a 2% chance of adding a node, a 5% chance of adding a link, and the weight mutation power is 0.4. Parameter settings are based on standard NEAT defaults and were found to be robust to moderate variation. Runs consisted of 500,000 evaluations, which is equivalent to 1,000 generations of 500 individuals in a generational EA.

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 the first experiment, as in the original formulation in Lehman and Stanley [16], the robot has 400 timesteps to navigate through the maze and the fitness bias f_b was 300. In the second experiment, the robot has 1,200 timesteps to reach both points, and the fitness bias f_b was 500. The position of the robot is sampled three times during an evaluation to construct the behavioral characterization.

5. RESULTS

The maze navigation task and the two-point navigation task were run 100 times with fitness-based NEAT, NEAT with novelty search, and NEAT with MCNS. Figure 3 highlights that in both tasks NEAT with MCNS evolves solutions with significantly higher probability than the other two approaches after 500,000 evaluations ($p < 0.01$; Fisher's exact test).

6. DISCUSSION

The results support the hypothesis that reducing the behavior space through enforcing minimal criteria can sometimes increase the efficiency of novelty search. They also further highlight the vulnerability of traditional fitness-based evolution. Thus, MCNS is a way to focus the search for novelty on an interesting subset of behaviors while still remaining agnostic to which of the behaviors in the viable subset will ultimately lead to the objective.

Furthermore, the two-point navigation experiment demonstrates that the search for many ways to do the same thing, an abstraction of how natural evolution accumulates novelty and complexity, can discover policies that exhibit complexity orthogonal to meeting the minimal criteria but essential to solving the problem. Thus, MCNS connects a coarse abstraction of natural evolution to solving real world problems.

While the empirical results show that MCNS can sometimes solve problems more consistently than other methods, the claim is not that MCNS will always work well; if the minimal criteria is too stringent or if evolution is seeded with an individual that is fragile and cannot be mutated without violating the minimal criteria, then most evaluations may be wasted and search will be inefficient. The minimal criteria and the seed genome must be carefully considered just like the fitness function in standard EC models.

Although MCNS is a practical search technique for solving difficult EC problems, it also demonstrates that EC models need not always be derived from the prevalent fitness-pressure-based abstraction of natural evolution. Because the drive to complexity in natural evolution may in fact be opposed by selection pressure, further pursuing the pervasive objective-based paradigm may yield little progress. Instead, alternate abstractions such as MCNS that more accurately abstract the high-level process by which natural evolution discovers complexity may be a more appropriate basis for solving highly ambitious problems.

Both the raw search for novelty and MCNS exhibit an open-ended search process that in principle could continue indefinitely to discover novel policies of increasing complexity, which makes them relevant to artificial life research as well. Yet the maze navigation and two-point navigation domains offer limited potential for behaviors that a human observer would find *interesting*. This limitation is in stark

contrast to natural evolution, which exhibits a vast diversity of interesting forms and behaviors. Thus, in future research, MCNS will be applied in a richer domain with more potential for captivating our intuitive notions of interestingness in artificial life. The ambitious goal of such an experiment is to create an explosion of complexity and diversity reminiscent, though limited in scale, of that found in natural evolution.

In summary, paradoxically, the best way to learn to do something new may often be to try to do the same thing again.

7. CONCLUSIONS

This paper introduced *minimal criteria novelty search*, an extension of the existing novelty search method that requires individuals in novelty search to meet minimal criteria before being eligible to reproduce. Inspired by a new abstraction of natural evolution, MCNS addresses both practical issues when searching through vast behavior spaces and the complexity disparity between EC models and natural evolution. In both a maze navigation task and a more difficult two-point navigation task, MCNS evolved solutions more consistently than either the raw search for novelty or objective-based search. The conclusion is that by refining the abstraction upon which the evolutionary algorithm is based, the resulting search method can solve practical problems while exhibiting dynamics more akin to natural evolution.

8. ACKNOWLEDGEMENTS

This research was supported by DARPA under grant HR0011-09-1-0045 (Computer Science Study Group Phase II).

References

- [1] T. Aaltonen et al. Measurement of the top quark mass with dilepton events selected using neuroevolution at CDF. *Physical Review Letters*, 2009.
- [2] Christoph Adami, Charles Ofria, and Travis C. Collier. Evolution of biological complexity. *Proc. of the Natl. Academy of Sciences, USA*, 97:4463–4468, 2000.
- [3] Brian Allen and Petros Faloutsos. Complex networks of simple neurons for bipedal locomotion. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009.
- [4] WB Arthur. On the evolution of complexity. In G Cowan, D Pines, and D Meltzer, editors, *Complexity: Metaphors, Models and Reality*, pages 65–81. Addison-Wesley, Reading, MA, 1994.
- [5] Mark Bedau. Four puzzles about life. *Artificial Life*, 4: 125–140, 1998.
- [6] A. D. Channon and R. I. Damper. Towards the evolutionary emergence of increasingly complex advantageous behaviours. *International Journal of Systems Science*, 31(7):843–860, 2000.
- [7] Charles Darwin. *On the Origin of Species by Means of Natural Selection or the Preservation of Favored Races in the Struggle for Life*. Murray, London, 1859.
- [8] R. Dawkins. Human chauvinism. *Evolution*, 51(3): 1015–1020, 1997.
- [9] Richard Dawkins. *The Selfish Gene*. Oxford University Press, Oxford, UK, 3rd edition, 1989.
- [10] J. A. Endler and J. J. D. Greenwood. Frequency-Dependent Predation, Crypsis and Aposematic Coloration [and Discussion]. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 319 (1196):505–523, 1988.




Figure 3: Comparing MCNS, Novelty Search and Fitness-based Search. The change in probability of success over time (i.e. number of evaluations) is shown for NEAT with MCNS, NEAT with novelty search, and fitness-based NEAT on the maze navigation (a) and two-point navigation (b) tasks, both averaged over 100 runs of each approach. The main result is in both tasks MCNS evolves solutions with significantly higher probability after 500,000 evaluations ($p < 0.01$; Fisher’s exact test).

- [11] Sevan Ficci and Jordan B. Pollack. Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In *Proceedings of the Sixth International Conference on Artificial Life*, pages 238–247. MIT Press, 1998.
- [12] Steven Jay Gould. *Full House: The Spread of Excellence from Plato to Darwin*. Harmony Books, 1996.
- [13] Inman Harvey. *The Artificial Evolution of Adaptive Behavior*. PhD thesis, School of Cognitive and Computing Sciences, University of Sussex, Sussex, 1993.
- [14] Kimberly A. Hughes, Linh Du, F. Helen Rodd, and David N. Reznick. Familiarity leads to female mate preference for novel males in the guppy, *pomcilia reticulata*. *Animal Behavior*, 58(4):907–916, 1999.
- [15] George Kampis and László Gulyás. Full body: The importance of the phenotype in evolution. *Artificial Life*, 14(3):375–386, 2008.
- [16] Joel Lehman and Kenneth O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *Proc. of the Eleventh Intl. Conf. on Artificial Life (ALIFE XI)*, Cambridge, MA, 2008. MIT Press.
- [17] Michael Lynch. The evolution of genetic networks by non-adaptive processes. *Nature Reviews Genetics*, 8: 803–813, 2007.
- [18] 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.
- [19] Andrew P. Martin. Increasing genomic complexity by gene duplication and the origin of vertebrates. *The American Naturalist*, 154(2):111–128, 1999.
- [20] Daniel W. McShea. Complexity and evolution: What everybody knows. *Biol. and Phil.*, 6(3):303–324, 1991.
- [21] Thomas Miconi. Evolution and complexity: The double-edged sword. *Artificial Life: Special Issue on the Evolution of Complexity*, 2007.
- [22] Tom M. Mitchell. *Machine Learning*. McGraw-Hill, New York, 1997.
- [23] Jean-Baptiste Mouret. Novelty-based multiobjectivization. In *Proc. of the Workshop on Exploring New Horizons in Evolutionary Design of Robots, 2009 IEEE/RSJ Intl Conf on Intelligent Robots and Systems*, 2009.
- [24] C. L. Nehaniv and J. L. Rhodes. On the manner in which biological complexity may grow. In *Mathematical and Computational Biology*, volume 26 of *Lectures on Mathematics in the Life Sciences*, pages 93–102. American Mathematical Society, 1999.
- [25] Sebastian Risi, Sandy D. Vanderbleek, Charles E. Hughes, and Kenneth O. Stanley. How novelty search escapes the deceptive trap of learning to learn. In *Proc. of the Genetic and Evolutionary Computation Conf. (GECCO-2009)*, New York, NY, 2009. ACM.
- [26] Michael Ruse. Evolution and progress. *Tree*, 8(2):55–60, 1993.
- [27] J. Schmidhuber. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18(2):173–187, 2006.
- [28] Karl Sigmund. *Games of Life: Explorations in Ecology, Evolution and Behaviour*. Oxford University Press, Inc., New York, NY, USA, 1993.
- [29] Andrea Soltoggio and Ben Jones. Novelty of behaviour as a basis for the neuro-evolution of operant reward learning. In *Proc. of the Genetic and Evol. Comp. Conf. (GECCO-2009)*, New York, NY, 2009. ACM.
- [30] Russell Standish. Open-ended artificial evolution. *Intl. Journal of Comp. Intell. and Applications*, 3(167), 2003.
- [31] 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 Evol. Comp. and Games*, 9(6):653–668, 2005.
- [32] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.
- [33] Kenneth O. Stanley and Risto Miikkulainen. A taxonomy for artificial embryogeny. *Artificial Life*, 9(2): 93–130, 2003.
- [34] Kenneth O. Stanley and Risto Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Art. Intell. Research*, 21:63–100, 2004.
- [35] Kenneth O. Stanley and Risto Miikkulainen. Evolving a roving eye for Go. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2004)*, Berlin, 2004. Springer Verlag.
- [36] James D. Watson, Nancy H. Hopkins, Jeffrey W. Roberts, Joan A. Steitz, and Alan M. Weiner. *Molecular Bio. of the Gene Fourth Edition*. The Benjamin Cummings Publishing Company, Inc., 1987.
- [37] Shimon Whiteson and Peter Stone. Evolutionary function approximation for reinforcement learning. *Journal of Machine Learning Research*, 7:877–917, May 2006.
- [38] N. Zaera, D. Cliff, and J. Bruten. (Not) evolving collective behaviours in synthetic fish. In *From Animals to Animats 4: Proc. of the Fourth Intl. Conf. on Sim. of Adaptive Behavior*. MIT Press Bradford Books., 1996.