

# Improving Evolvability through Novelty Search and Self-Adaptation

In: *Proceedings of the 2011 IEEE Congress on Evolutionary Computation (CEC 2011)*. Piscataway, NJ: IEEE

Joel Lehman

Department of Electrical Engineering  
and Computer Science  
University of Central Florida  
Orlando, Florida 32826–2363  
Email: jlehman@eecs.ucf.edu

Kenneth O. Stanley

Department of Electrical Engineering  
and Computer Science  
University of Central Florida  
Orlando, Florida 32826–2363  
Email: kstanley@eecs.ucf.edu

**Abstract**—A challenge for current evolutionary algorithms is to yield highly evolvable representations like those in nature. Such evolvability in natural evolution is encouraged through selection: Lineages better at molding to new niches are less susceptible to extinction. Similar selection pressure is not generally present in evolutionary algorithms; however, the first hypothesis in this paper is that novelty search, a recent evolutionary technique, also selects for evolvability because it rewards lineages able to continually radiate new behaviors. Results in experiments in a maze-navigation domain in this paper support that novelty search finds more evolvable representations than regular fitness-based search. However, though novelty search *outperforms* fitness-based search in a second biped locomotion experiment, it proves no more *evolvable* than fitness-based search because delicately balanced behaviors are more fragile in that domain. The second hypothesis is that such fragility can be mitigated through self-adaptation, whereby genomes influence their own reproduction. Further experiments in fragile domains with novelty search and self-adaptation indeed demonstrate increased evolvability, while, interestingly, adding self-adaptation to fitness-based search *decreases* evolvability. Thus, selecting for novelty may often facilitate evolvability when representations are not overly fragile; furthermore, achieving the potential of self-adaptation may often critically depend upon the reward scheme driving evolution.

## I. INTRODUCTION

The philosopher Diderot [5] once wrote, “It seems that Nature has taken pleasure in varying the same mechanism in an infinity of different ways ... She abandons one class of production only after having multiplied the individuals of it in all possible forms.” Naturalists of all generations have been similarly inspired by the vast productivity and creativity of evolution. Such diversity in nature highlights the significant evolvability of natural organisms; the tree of life branches ever outwards with high degree.

Evolvability, or the capacity of an individual or a population to further evolve, has often been studied in the context of natural evolution [3, 12, 20, 31]. A potential explanation for natural evolution’s propensity to generate a wide diversity of organisms is that greater evolvability is selected for. That is, evolvable lineages that are able to modify into more niches have less chance of going entirely extinct [3, 12]. Because of the apparent chasm between the evolvability of organisms in natural evolution and in artificial systems, measuring and

increasing evolvability has also been the subject of many studies within evolutionary computation (EC) [6, 9, 22, 31].

Traditional evolutionary algorithms (EAs), which typically are driven to *converge* to a fitness optimum, contrast with natural evolution and its consistent exploratory diffusion through niches. Such EAs thus often lack selection pressure towards greater evolvability. That is, lineage-level selection is confounded when evolution converges to a single lineage; in such cases there is no *variety* of lineages upon which a higher level of selection pressure can act. Thus when optimizing static fitness functions like those prevalent in EC, fitness may be an impoverished indicator of evolvability.

In contrast, a recent technique in EC, *novelty search* [16], directly rewards individuals exhibiting novel behaviors and is inspired by natural evolution’s continual production of novel forms. Such selection for novelty naturally maintains a wide diversity of behaviors and lineages more similar in spirit than traditional EAs to the accumulation of niches seen in nature.

The first hypothesis in this paper is thus that novelty search may also indirectly select for *evolvable* lineages just as natural evolution does. That is, in the long run, lineages in novelty search more predisposed to producing novelty will generate more offspring than those lineages that are less evolvable. The converse does not hold for traditional objective-based EAs: Once such EAs converge to a local optimum, selection rewards individuals with *decreased* variability that are able to achieve ever slighter fitness increases [31]; the prototypical fitness curve in EC illustrates that in general as fitness increases it becomes increasingly difficult to further increase it.

In this paper, following Kirschner and Gerhart [12], evolvability is defined as the capacity of an organism to “generate heritable phenotypic variation.” While evolvability is often discussed in relation to adaptation, the chosen definition reflects a growing consensus in biology that phenotypic variability in its own right deserves study in the context of evolvability [3, 12, 20, 31]. Thus the evolvability of an individual is accordingly quantified in this paper by measuring its propensity for phenotypic variation [3]. Supporting the first hypothesis, experiments in a maze navigation domain demonstrate that novelty search does discover more evolvable

individuals, populations, and solutions than traditional fitness-based search. However, in a challenging biped domain, though it *outperforms* fitness-based search, novelty search is no more *evolvable* than fitness-based search.

This discrepancy leads to the insight that fragile domains, i.e. those in which nearly every mutation is fatal, can preclude search methods from finding evolvable representations. That is, the weak indirect selection pressure towards evolvability in novelty search can be overpowered by the combination of uniform mutation operators with an unforgiving domain.

Yet the second hypothesis is that *self-adaptation* can restore the evolvability of novelty search. Thus to remedy domain fragility, both novelty search and fitness-based search are extended with the capability to self-adapt mutational parameters, effectively giving the genome greater control over the creation of its offspring. An important result is that the greater control provided by such self-adaptation does not help fitness-based search, which exploits self-adaptation to more effectively *converge* as it nears local fitness optima [4, 7, 24]. However, it *does* allow novelty search to discover more evolvable representations and sometimes to perform better.

There are two conclusions. The first is that because novelty search indirectly selects for evolvable lineages, it may often discover more evolvable representations than traditional objective-based search. The second conclusion is that while self-adaptation has the potential to increase both performance and evolvability, achieving the latent potential of self-adaptation may be more effective in evolutionary algorithms that encourage novelty.

## II. BACKGROUND

Evolvability in both natural and artificial systems is reviewed in this section, as is the neuroevolution method applied in the experiments in this paper. Finally, novelty search, the evolutionary technique inspired by natural evolution’s continual production of diverse novel forms, is explained.

### A. Evolvability in Natural Evolution and EC

Natural evolution has discovered flexible, highly evolvable representations that have facilitated the discovery of widely diverse organisms. An important question that could inform EC is what properties of natural evolution led to such evolvability? At the same time, EC enables studies of evolvability that can potentially inform biology through computational experiments impossible in nature.

From a biological perspective, Kirschner and Gerhart [12] describe evolvability as “an organism’s capacity to generate heritable phenotypic variation,” and suggest that evolvability in natural evolution results from an accumulation of flexible building blocks that are heavily conserved. However, they can be combined and regulated in many ways to yield substantial phenotypic variety with few mutations. Further, they argue that such evolvable lineages may be selected for; a lineage able to discover and exploit new niches or to quickly radiate through existing niches following extinctions will itself be less likely to go completely extinct.

Examining evolvability from both biology and EC, Wagner and Altenberg [31] similarly describe evolvability as relating to the phenotypic variability of a genome, and argue that EC can potentially provide insight into evolvability that biology cannot; the effect of alternate genetic representations on evolvability can be tested within EC but are not easily explored in nature. Furthermore, the authors argue that the structure of the genotype to phenotype mapping is fundamental to evolvability. This mapping, which includes the mechanism for the reproduction and mutation of an organism, is itself subject to selection and evolution in nature.

Studies in EC have described a lack of evolvability in practice [22, 31] and have noted possible directions for increasing evolvability. Some argue that static fitness functions, which are prevalent in most of EC, do not encourage evolvability [6, 22] and instead suggest that fitness functions should change over the course of evolution. Other suggestions for improving evolvability are to allow adaptation of the genetic operator set [9], to increase the extent of neutral networks [6], or to employ indirect encodings that allow more plastic genotype to phenotype mappings [22, 31]. However, prior studies have not examined the effects of alternate reward schemes like novelty search on evolvability, as is undertaken in this paper.

### B. NeuroEvolution of Augmenting Topologies (NEAT)

In experiments in this paper, behaviors are evolved that are controlled by artificial neural networks (ANNs). Thus a neuroevolution (NE; [32]) method is needed to implement these experiments. The NEAT method is appropriate because it is widely applied [1, 2, 28, 29], well understood, and often run in conjunction with novelty search.

The NEAT method was originally developed to evolve ANNs to solve difficult control and sequential decision tasks [28, 29]. Evolved ANNs control agents that select actions based on their sensory inputs. Like the SAGA method [10] introduced before it, NEAT begins evolution with a population of small, simple networks and *complexifies* them. The population expands into diverse species over generations, leading to increasingly sophisticated behavior. A similar process of gradually adding new genes is seen in natural evolution [17]. This section briefly reviews the NEAT method; for comprehensive introductions see Stanley and Miikkulainen [28, 29].

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 [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.

The next section reviews novelty search, a technique in EC that rewards only behavioral novelty.

### C. Novelty Search

The problem with the objective-based search paradigm that is common in EC models is that an objective function (i.e. the fitness function) does not necessarily reward the intermediate stepping stones that lead to the objective, nor does it necessarily distinguish more evolvable individuals. The more ambitious the objective, the harder it is to identify *a priori* these evolvable stepping stones.

The approach suggested by Lehman and Stanley [16] is to identify novelty as a *proxy* for stepping stones. 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 differs from objective-based search by rewarding stepping stones. That is, anything that is genuinely different is rewarded and promoted as a jumping-off point for further evolution. Additionally, those lineages discovering *more* stepping stones will be indirectly rewarded, encouraging evolvability. 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 [25].

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, 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  $\rho$  at point  $x$  is given by

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

where  $\mu_i$  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. Note that novelty search is not an exhaustive or random search because many genotypes map to the same behavior, and the number of novel behaviors is reasonable and limited in many practical domains.

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 finding 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, 19, 23] and genetic programming [8, 15]. Novelty search was introduced in Lehman and Stanley [14] and applied to a deceptive maze task; these results were replicated in Mouret [19]. Experiments have also shown that encouraging novelty is useful in evolving adaptive ANNs (i.e ANNs that learn during their lifetimes) [23, 26].

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. However, an interesting question orthogonal to performance that is addressed in this paper is whether abandoning pressure to optimize the objective can also affect higher-level properties such as evolvability.

### III. MEASURING EVOLVABILITY

The primary hypothesis in this paper is that rewarding novelty increases evolvability when compared with rewarding progress towards a fixed objective. Therefore, an important aspect of evaluating this hypothesis is quantifying evolvability.

One definition of evolvability is “an organism’s capacity to generate heritable phenotypic variation” [12]. Brookfield [3] similarly suggests evolvability is “the proportion of radically different designs created by mutation that are viable and fertile.” These definitions reflect a growing consensus in biology that the ability to generate phenotypic variation is fundamental to evolvability [3, 12, 20, 31].

From these definitions, a way of estimating an individual’s evolvability is to generate many children from it and then measure the degree of phenotypic variation among those offspring. In effect, this measure quantifies how well the organization of the individual’s representation has internalized domain information to enable more behaviorally diverse mutations. This measure is similar in intention to that in Reisinger et al. [22] but is more granular because it quantifies how well an

*individual's* evolved representation exploits domain structure instead of measuring an *encoding's* ability to do the same.

Accordingly, the measure of evolvability in this paper is to count how many distinct behaviors there are among a series of generated offspring of a particular individual. An individual whose offspring tend to display many distinct behaviors is capable of generating much heritable phenotypic variation, and is thus evolvable. In practice, such a measure requires a domain-fitted test between two individuals to determine whether their behaviors are distinct.

The domain-specific novelty metric in novelty search that measures behavioral distance between two individuals can naturally be adapted for this purpose: Two behaviors are distinct if the behavioral distance between them is greater than a fixed threshold.

Using this test of behavioral distinction, the number of distinct behaviors among a list of behaviors can be calculated by the following greedy algorithm that accumulates a list of distinct behaviors: The first behavior is added to this list by default, and each subsequent behavior is added if it is distinct from each behavior already in the distinct behavior list. The size of this filtered list is the number of distinct behaviors, an estimate of the individual's evolvability.

At regular intervals during a run, the evolvability of each individual in the population is measured. That is, for each individual in the population many offspring are sequentially generated by first cloning the individual and then mutating the clone; the idea is to sample the genotypic space around an individual and examine the distribution of those samples in behavior space. The behaviors of these perturbed clones are then concatenated to form a list, with the number of distinct behaviors in the concatenated list acting as an indicator of that individual's phenotypic variability, i.e. its evolvability.

#### IV. EXPERIMENTS

To compare the *evolvability* of individuals evolved by both novelty search and traditional objective fitness-based search, experiments are conducted in two domains previously employed to compare the *performance* of the two varieties of search, i.e. maze-navigation and biped locomotion [16].

This paper's experiments utilize NEAT, which has been proven in many control tasks [2, 14, 16, 28, 29], including maze navigation [14] and biped locomotion [2], the test domains in this paper. In both experiments, the evolvability of individuals evolved with NEAT with novelty search is compared to that of traditional fitness-based NEAT.

The next sections describe these maze navigation and biped locomotion experiments in detail.

##### A. Maze Experiment

The maze navigation domain in Lehman and Stanley [16] is a good model for behavior spaces in general because it is easy to understand and visualize [16, 19]. In this domain, a robot controlled by an ANN must navigate in a maze (figure 1) from a starting point to an end point in a fixed time. The robot (figure 2) has six rangefinders that indicate the distance




Fig. 1: **Maze Navigation Map.** In this map, the larger circle represents the starting position of the robot and the smaller circle represents the goal.

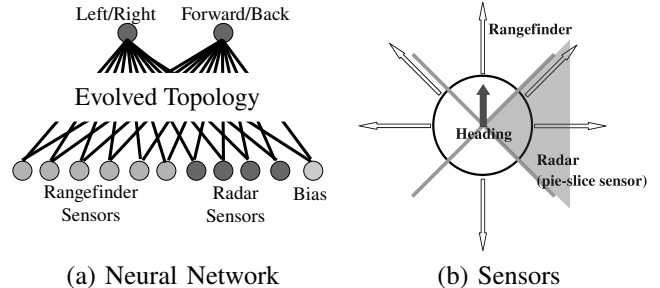


Fig. 2: **Maze-Navigating Robot.** The artificial neural network that controls the maze navigating robot is shown in (a). The layout of the sensors is shown in (b). Each arrow outside of the robot's body in (b) is a rangefinder sensor that indicates the distance to the closest obstacle in that direction. 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.

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.

Fitness-based NEAT, which will be compared to novelty search, 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, 19]; 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.

NEAT with novelty search, on the other hand, requires a novelty metric to distinguish between maze-navigating robots. Following Lehman and Stanley [16], the behavior of a robot is defined as its location in the maze at the end of the evaluation [14, 19], reflecting what is important in navigating a maze. The novelty metric is then the squared Euclidean distance between the ending positions of two individuals. This same novelty metric also distinguishes behaviors when calculating evolvability; two behaviors are distinct if the metric between them is greater than a fixed threshold.

Population size for the hard maze was 250 and each run lasted 1,000 generations. When measuring evolvability, 200 offspring were created for each individual in the population and evolvability was measured every 50 generations. Solution criteria in the maze domain are the same as in Lehman and Stanley [16] (i.e. ending within five units of the goal). All other parameters are the same as in Lehman and Stanley [16].

## B. Biped Experiment

A more challenging domain is biped locomotion, a difficult control task that is popular within machine learning [2, 18, 21, 30]. The key issue is whether the brittleness in the biped domain from the need for balance and oscillation [11] might prevent an evolutionary algorithm from discovering evolvable individuals able to walk far.

In this domain, a three-dimensional biped robot in a realistic physics simulation is controlled by a type of ANN called a *continuous time recurrent neural network* (CTRNN) that is able to express the non-linear dynamics found in natural gaits and is common in other biped experiments [18, 21]. The objective is to walk as far as possible within a given time limit. The task is difficult because it requires coordination, balance, and the discovery of oscillatory patterns. Discovering such oscillatory patterns critical to cyclic walking is particularly deceptive because they may initially provide no immediate fitness advantage; fitness-based search quickly converges to robots that fall the farthest, and the nascent discovery of oscillatory patterns may actually *impede* such strategies.

The biped domain works as follows. A biped robot is controlled by an ANN for a fixed duration (15 simulated seconds). The evaluation is terminated if the robot falls or after the allocated time expires. The objective is that the robot travel the greatest possible distance from the starting location.

The ANN that controls the biped has only two inputs, which are contact sensors that signal for each foot whether it is touching the ground. The sparsity of input makes the problem more difficult because the ANN has *no information* on the orientation of the robot or on the current angles of its joints.

The biped robot (figure 3) has a total of six degrees of freedom (DOF): two degrees in each hip joint (pitch and roll) and one degree in each knee joint (pitch). Simple sphere-shaped feet make ankle joints unnecessary in this model, although the lack of typical feet or a torso (which could provide a counterbalance) require knees to bend backwards (as in birds) to balance, adding to the challenge of the model.

The ANN outputs movement requests for each degree of freedom (DOF) in the model, i.e. for each independent axis of rotation for all joints in the model. The outputs are scaled to match the angular range of the corresponding DOF, which is interpreted as the angle that the neural controller is requesting. The difference between the requested angle and the current orientation of the DOF denotes the disparity between the state the ANN is requesting and the current state of the model. A proportional controller applies torque to reduce this disparity. In other words, the ANN directs the low-level controllers towards a particular state. This model and method of control are similar to those in Reil and Husbands [21].

This experiment again compares the evolvability of fitness-based NEAT to NEAT with novelty search. It is important to note that NEAT was extended for these experiments to evolve CTRNNs (i.e. it also evolves the time constant assigned to each node). A natural fitness function is the squared distance traveled from the starting location. This distance is measured by recording the location of the center of mass of the biped




Fig. 3: **Biped Robot.** A visualization of the biped robot controlled by an evolved ANN in the biped experiment in this paper.

robot before and after evaluation, and then calculating the distance between the two points. Distance traveled is a standard measure among evolutionary biped locomotion experiments [18, 21, 30], and matches the intuitive notion of learning to walk with increasing stability and efficiency.

In contrast, NEAT with novelty search requires a behavioral characterization to distinguish biped gaits. The behavioral characterization in this domain is the offset of the biped’s center of mass sampled at one second intervals during the evaluation:

$$x'_k = \text{sign}(x_k - x_0) * (x_k - x_0)^2, \quad (2)$$

$$y'_k = \text{sign}(y_k - y_0) * (y_k - y_0)^2, \quad (3)$$

where  $x_0$  and  $y_0$  correspond to the initial planar center of mass (i.e. ignoring the vertical  $z$  component) of the biped robot, and  $x_k$  and  $y_k$  correspond to the center of gravity sample taken after the  $k$ th second of simulation. The magnitude of the offsets is squared just as the fitness measure is to make the comparison uniform. If the robot falls, the center of gravity for all remaining samples is set to the robot’s center of gravity when it fell. The behavioral characterization of a particular controller is then the vector that concatenates all pairs  $(x'_i, y'_i)$ , where  $1 \leq i \leq m$  and  $m$  is the final sample taken, to form  $(x'_1, y'_1, x'_2, y'_2, \dots, x'_m, y'_m)$ .

The novelty metric for two different controllers is the same as in the maze domain, i.e. the sum of the squared distances between the controllers’ behavioral characterization vectors. Unlike in the maze domain, temporal sampling is necessary because the temporal pattern is fundamental to walking. This additional information allows the novelty metric to differentiate two gaits that end up at the same location by different means.

Again, this novelty metric also acts as a behavioral distinction test when measuring evolvability; two behaviors are distinct if the behavioral distance between them is greater than a fixed threshold.

The population had 500 individuals and each run lasted 2,000 generations. When measuring evolvability, 200 offspring were cloned from each individual in the population and evolvability was measured every 100 generations. The walker traveling the farthest in each run is considered the solution. Other parameters are the same as in Lehman and Stanley [16].

## V. RESULTS

In the hard maze experiment, novelty search discovered solutions in every one of the 20 runs while fitness-based




Fig. 4: Comparing the evolvability of novelty search and fitness-based search. The average evolvability of evolved solutions is shown in (a). The average and maximum evolvability of all individuals in the population are shown for novelty search and fitness-based search in (b) the hard maze and (c) the biped walking experiments. Results are averaged over the initial 20 runs. Error bars indicate standard deviation.

search solved it only once. Thus an additional 580 runs of fitness-based search were conducted to accumulate a more representative sample of solutions. Out of the combined 600 runs, 41 were successful. The average evolvability of solutions (figure 4a) discovered by novelty search in the hard maze ( $\mu = 26.03$ ,  $sd = 2.026$ ) was significantly higher ( $p < 0.05$ ; Mann-Whitney U-test) than that of solutions found by fitness-based search ( $\mu = 24.76$ ,  $sd = 2.97$ ). In addition, reflecting evolvability across evolution, after the 25th generation, the evolvability of the most evolvable individual and the average evolvability of individuals (figure 4b) in the population were significantly higher ( $p < 0.01$ ; Mann-Whitney U-test) in novelty search than in fitness-based search.

In the biped experiment, novelty search outperformed fitness-based search by evolving walkers that traveled farther on average ( $p < 0.01$ ; Mann-Whitney U-test). Yet, unlike in the maze, the evolvability of biped champions (figure 4a) was not significantly higher in novelty search than in fitness-based search. Neither were the evolvability of the most evolvable individual or the average evolvability of individuals in the population (figure 4c) significantly different between novelty search and fitness-based search (Mann-Whitney U-test).

## VI. SELF-ADAPTATION EXPERIMENTS

The results in the hard maze conflict with those in the biped domain; though novelty search evolved more evolvable individuals in the maze domain, those evolved in the biped domain were not significantly more evolvable. Further examination shows that individuals evolved by novelty search and fitness-based search in the biped domain are very fragile; nearly all mutations of such individuals immediately fall, suggesting that the mutation and recombination operators are ill-fitted to the evolved representations.

A potential remedy to such fragility is self-adaptation: Augmenting a genome with additional parameters can allow evolution to adapt parameters governing reproduction such as mutation rate [4, 13]. A genome with such self-adaptive capabilities has greater influence over how its offspring are created. The idea is that evolution can select for self-adaptive parameters that better complement a genome’s evolved representation and mitigate domain fragility.

An additional fragile domain is crafted to more generally test the new hypothesis that self-adaptation may encourage evolvability in fragile domains. In the original conception of the maze domain there is no penalty for colliding with walls, and all evolved solutions collide with walls as they navigate to the goal. Thus, a more difficult variant of the maze domain is created that terminates an evaluation when an individual collides with a wall; a crashed individual receives a very low fitness (for fitness search) and a dummy behavior characterization to a particular point outside the map (for novelty search). The motivation is to make gradients of both improvement (for fitness-based search) and behavioral change (for novelty search) more brittle, mirroring the inherent brittleness of the biped domain.

To implement self-adaptation in NEAT, genomes are augmented with floating point parameters that influence the mutation of connection weights. Two such parameters are considered here: Connection mutation power, which modulates the magnitude of a potential change to a weight, and connection mutation rate, which specifies the probability that a particular connection weight will be changed at all.

Initial experiments revealed that evolving a single mutation power and rate for each entire genome is too coarse; certain connections in the neural network may be more or less sensitive to perturbation, a possibility impossible to exploit with only a single set of mutation parameters per genome [7]. At the opposite extreme, evolving separate sets of mutation parameters for every connection in an ANN greatly expands the parameter space evolution searches through. A balance is achieved by augmenting NEAT’s genome with a fixed-length set of mutation parameters for each genome among which each connection within the ANN can choose. That is, with this kind of self-adaptation each connection weight in a genome is mutated according to a particular rule within the set of rules the genome contains. This setup is similar to that employed by Stanley and Miikkulainen [27] in the context of rules for adaptation in plastic networks.

Thus, each genome is given a list of three pairs of mutation settings and each connection is augmented with an integer parameter that indexes within the list. When an individual is mutated, every one of its links are usually mutated by

adding to its weight a random number chosen from the same uniform distribution. With self-adaptation, each connection of the neural network may be *exempted* from mutation with probability determined by the connection mutation rate. If it is chosen to be mutated, the power of the mutation is modulated by multiplying it with the connection mutation power. Thus, a self-adaptive genome can potentially craft an architecture of mutation that complements the genome’s evolved representation.

It is important to note that reproduction in these self-adaptation experiments is only asexual, as in some other studies [4, 7]; the reason here is that crossover facilitates mixing of self-adaptive parameters between different lineages, which can overpower weaker lineage-level selection pressure on such parameters that have no immediate impact on fitness [20, 31]. The previous experiments, with the addition of the new fragile hard maze domain, are repeated with self-adaptation to investigate the effect of such self-adaptation on the evolvability of both fitness-based NEAT and NEAT with novelty search.

The added self-adaptive rules that influenced mutation of ANN connection weights themselves changed probabilistically at *fixed* rates [4, 13]. During mutation of a genome, a randomly-selected rule would be changed with probability 0.75. Each of the two floating point parameters making up such a rule are then perturbed with probability 0.2 by a random number uniformly selected between  $-0.2$  and  $0.2$ . Each genome also had a 10% chance of changing to which rule a randomly-chosen connection was linked.

The population size for the fragile hard maze was 250 and evolution lasted 2,000 generations. Evolvability was measured every 100 generations. Other parameters were the same as for the normal hard maze experiment.

## VII. SELF-ADAPTATION RESULTS

Figure 5 shows how self-adaptation effects novelty search and fitness-based search in the hard maze, biped, and fragile hard maze domains. The main result is that self-adaptation allows novelty search to discover significantly more evolvable representations in the biped and fragile hard maze ( $p < 0.05$ ; Mann-Whitney U-test), while combining fitness-based search and self-adaptation is *universally* detrimental across every measure and domain ( $p < 0.05$ ; Mann-Whitney U-test).

In addition, novelty search combined with self-adaptation increased *performance* in the fragile hard maze as well; in 300 additional runs it was solved 31 times by that setup, while it was solved at best only 14 times out of 300 by other setups. The difference is significant ( $p < 0.05$ ; Fisher’s exact test). There was no significant difference in the average distance traveled by biped champions between novelty search with self-adaptation and without it, but the two novelty search setups significantly outperformed the two fitness-based setups ( $p < 0.05$ ; Mann-Whitney U-test). Additionally, combining fitness-based search with self-adaptation hurt performance when compared to fitness-based search without self-adaptation ( $p < 0.05$ ; Mann-Whitney U-test).

## VIII. DISCUSSION

The results support the idea that novelty search may often discover more evolvable representations than fitness-based search, although the ability of any method to discover evolvable representations is likely stymied by a poor fit between reproductive operators, the representation, and the domain. In addition, while novelty search can effectively mold mutational parameters to match the representations it discovers, fitness-based search instead suffers from such an additional capability.

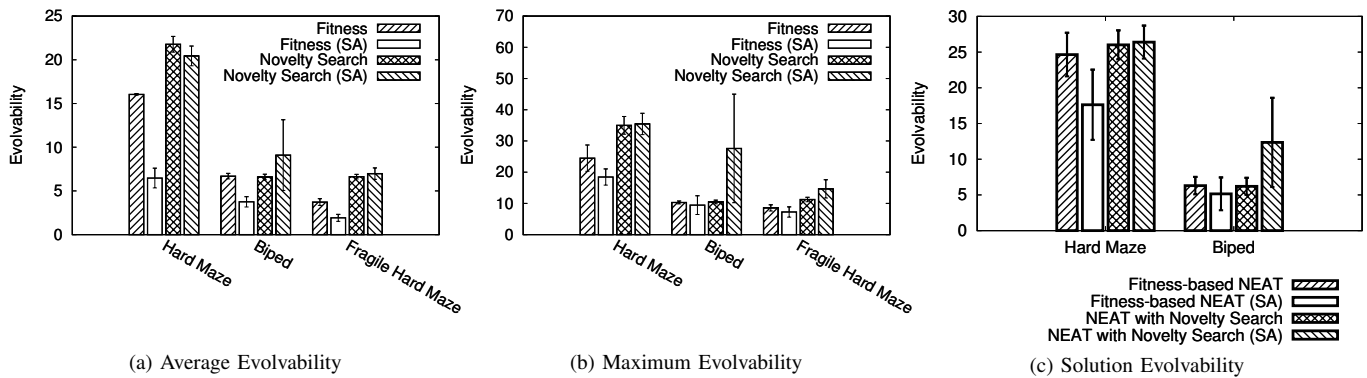
More significantly, the inability of fitness-based search to exploit self-adaptation and its lesser evolvability may illustrate the greater *conflict* between greedily pursuing the objective and desirable representation. That is, a lack of representational variety prevents discovering *better* representations. This handicap is a significant challenge that evolution may often face when employed as an optimizer; by considering only current fitness, traditional EAs often converge to a particular lineage and its corresponding representation without regard to higher-level properties such as evolvability. Evolution must then elaborate on what may be a poor representation.

Interestingly, the results in this study may suggest revisiting prior disappointing results from testing EAs with self-adaptive mutation rates on non-trivial fitness landscapes [4, 7, 13, 24]. Often in such cases self-adaptive mutation rates are prematurely minimized (which accelerates premature convergence) because mutations leading to further fitness increases become exceedingly rare [4, 7, 13, 24]. In contrast, coupling self-adaptation with novelty search facilitates an *improved* balance between mutation rate and innovation, even in the deceptive and challenging biped locomotion domain. Thus prior experimental results that might be interpreted as illustrating the *limitations* of self-adaptation could be partially misleading. That is, the tendency of self-adaptation to exacerbate premature convergence in prior experiments may be an incidental byproduct of the reward scheme currently dominant in EC.

An interesting direction for future research is to conduct a similar study with indirect encodings that have more powerful genotype to phenotype mappings than NEAT’s direct encoding of neurons and connections in an ANN. The representations learned by such indirect encodings may more heavily constrain and directionally influence the behaviors of their offspring, allowing for greater degrees of evolvability. For example, in natural evolution organizational principles themselves can evolve, enabling structural regularities like symmetry or modularity and exploratory mechanisms such as adaptive immune systems or behavioral plasticity [12]. Thus, the gain in evolvability from rewarding novelty may be greater and more critical in indirect encodings also able to exploit similar regularities.

## IX. CONCLUSIONS

The conclusion is that oftentimes pursuing an objective may not only hinder evolvability but also prevent the achievement of self-adaptation’s full potential. The lesson is that the type of evolutionary search applied can influence important properties beyond performance, including the quality of representation discovered and the ability of search to exploit selection on



**Fig. 5: Effects of self-adaptation on evolvability.** In all figures, methods with self-adaptation are indicated by the abbreviation *SA*. The average evolvability of all individuals in the final population for each variant are shown in (a), the maximum evolvability of all individuals in the final population are shown in (b), and the average evolvability of discovered solutions are in (c). Results are averaged over 20 runs and error bars indicate standard deviation. The main implication is that self-adaptive parameters help evolvability in novelty search but not in fitness-based search.

higher levels. At heart, natural evolution is not a method for black-box optimization; by shackling EC to that end we may inadvertently obfuscate the emergence of natural evolution's characteristic properties.

#### ACKNOWLEDGMENTS

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

#### REFERENCES

- [1] T. Aaltonen et al. Measurement of the top quark mass with dilepton events selected using neuroevolution at CDF. *Phys. Rev. Letters*, 2009.
- [2] 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.
- [3] J.F.Y. Brookfield. Evolution: The evolvability enigma. *Current Biology*, 11(3):R106 – R108, 2001.
- [4] J. Clune, D. Misevic, C. Ofria, R.E. Lenski, S.F. Elena, and R. Sanjuán. Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS Computational Biology*, 4(9): 1145–1146, 2008.
- [5] D. Diderot. *Pensées sur l'interprétation de la nature*. 1754.
- [6] M. Ebner, M. Shackleton, and R. Shipman. How neutral networks influence evolvability. *Complexity*, 7(2):19–33, 2001.
- [7] M.R. Glickman and K. Sycara. Reasons for premature convergence of self-adapting mutation rates. In *Evol. Comp., 2000. Proc. of the 2000 Congress on*, volume 1, pages 62–69. IEEE, 2002.
- [8] H.J. Goldsby and B.H.C. Cheng. Automatically Discovering Properties that Specify the Latent Behavior of UML Models. In *Proceedings of MODELS 2010*, 2010.
- [9] J.J. Grefenstette. Evolvability in dynamic fitness landscapes: A genetic algorithm approach. In *Evolutionary Computation, 1999. CEC 99. Proceedings of the 1999 Congress on*, volume 3. IEEE, 2002.
- [10] Inman Harvey. *The Artificial Evolution of Adaptive Behavior*. PhD thesis, School of Cognitive and Computing Sciences, University of Sussex, Sussex, 1993.
- [11] Duško Katić and Miomir Vukobratović. Survey of intelligent control techniques for humanoid robots. *Journal of Intelligent Robotics Systems*, 37(2):117–141, 2003.
- [12] M. Kirschner and J. Gerhart. Evolvability. *Proceedings of the National Academy of Sciences of the United States of America*, 95(15):8420, 1998.
- [13] Oliver Kramer. Evolutionary self-adaptation: a survey of operators and strategy parameters. *Evolutionary Intelligence*, 3:51–65, 2010.
- [14] 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 Art. Life (ALIFE XI)*, Cambridge, MA, 2008. MIT Press.
- [15] Joel Lehman and Kenneth O. Stanley. Efficiently evolving programs through the search for novelty. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2010)*. ACM, 2010.
- [16] Joel Lehman and Kenneth O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evol. Comp.*, 2011. To appear.
- [17] Andrew P. Martin. Increasing genomic complexity by gene duplication and the origin of vertebrates. *The American Naturalist*, 154(2):111–128, 1999.
- [18] Gary McHale and Phil Husbands. Gasnets and other evolvable neural networks applied to bipedal locomotion. In *From Animals to Animats 8*, 2004.
- [19] Jean-Baptiste Mouret. Novelty-based multiobjectivization. In *Proc. of the Workshop on Exploring New Horizons in Evol. Design of Robots, 2009 IEEE/RSJ Intl. Conf. on Intell. Robots and Systems*, 2009.
- [20] M. Pigliucci. Is evolvability evolvable? *Nature Reviews Genetics*, 9(1): 75–82, 2008.
- [21] Torsten Reil and Phil Husbands. Evolution of central pattern generators for bipedal walking in a real-time physics environment. *IEEE Transactions on Evolutionary Computation*, 6(2):159–168, April 2002.
- [22] Joseph Reisinger, Kenneth O. Stanley, and Risto Miikkulainen. Towards an empirical measure of evolvability. In *Genetic and Evolutionary Computation Conference (GECCO2005) Workshop Program*, pages 257–264, Washington, D.C., 2005. ACM Press.
- [23] Sebastian Risi, Sandy D. Vanderbleek, Charles E. Hughes, and Kenneth O. Stanley. How novelty search escapes the deceptive trap of learning to learn. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2009)*, New York, NY, 2009. ACM.
- [24] G. Rudolph. Self-adaptive mutations may lead to premature convergence. *Evol. Comp., IEEE Trans. on*, 5(4):410–414, 2002.
- [25] J. Schmidhuber. Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Conn. Sci.*, 18(2):173–187, 2006.
- [26] 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.
- [27] Kenneth O. Stanley and Risto Miikkulainen. Efficient evolution of neural network topologies. In *Proceedings of the 2002 Congress on Evolutionary Computation (CEC'02)*, Piscataway, NJ, 2002. IEEE.
- [28] Kenneth O. Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evol. Comp.*, 10:99–127, 2002.
- [29] Kenneth O. Stanley and Risto Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21:63–100, 2004.
- [30] Michiel Van de Panne and Alexis Lamouret. Guided optimization for balanced locomotion. In *Sixth Eurographics Workshop on Anim. and Sim.*, pages 165–177. Springer Verlag, 1995.
- [31] G.P. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, 1996.
- [32] Xin Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, 1999.