

Evolvability Search: Directly Selecting for Evolvability in order to Study and Produce It

Henok Mengistu
Evolving AI Lab
Dept. of Computer Science
University of Wyoming
hmgengist@uwyo.edu

Joel Lehman
Robotics, Evolution, and Art
Laboratory
IT University of Copenhagen
jleh@itu.dk

Jeff Clune
Evolving AI Lab
Dept. of Computer Science
University of Wyoming
jeffclune@uwyo.edu

ABSTRACT

One hallmark of natural organisms is their significant evolvability, i.e., their increased potential for further evolution. However, reproducing such evolvability in artificial evolution remains a challenge, which both reduces the performance of evolutionary algorithms and inhibits the study of evolvable digital phenotypes. Although some types of selection in evolutionary computation indirectly encourage evolvability, one unexplored possibility is to *directly* select for evolvability. To do so, we estimate an individual's future potential for diversity by calculating the behavioral diversity of its immediate offspring, and select organisms with increased offspring variation. While the technique is computationally expensive, we hypothesized that direct selection would better encourage evolvability than indirect methods. Experiments in two evolutionary robotics domains confirm this hypothesis: in both domains, such *Evolvability Search* produces solutions with higher evolvability than those produced with Novelty Search or traditional objective-based search algorithms. Further experiments demonstrate that the higher evolvability produced by Evolvability Search in a training environment also generalizes, producing higher evolvability in a new test environment without further selection. Overall, Evolvability Search enables generating evolvability more easily and directly, facilitating its study and understanding, and may inspire future practical algorithms that increase evolvability without significant computational overhead.

Keywords

Evolvability; Evolutionary Algorithms; Divergent Search

1. INTRODUCTION

An important factor in the prolific creativity of natural evolution is its drive towards increased evolvability, which involves increasing an organism's or population's capacity for further evolution [3, 8]. In addition to accelerating the evolutionary process as a whole, such evolvability confers advantages to individual species by allowing them to avoid

extinction or to establish new niches [3, 13]. For all these reasons, encouraging evolvability is an important goal in evolutionary computation. Yet producing the evolvability of natural organisms remains a challenge [5, 16, 18, 35].

Optimization in traditional objective-based evolutionary search, i.e. where the fitness function is a static estimate of solution quality based on only one or a few objectives, need not result in increased evolvability [18]. The reason is that there is no necessary connection between increasing performance and increasing evolutionary potential; often, pressure to adapt in particular ways implicitly penalizes deviations which may nonetheless increase evolvability.

Because evolvability does not tend to emerge unaided in Evolutionary Computation (EC), a variety of mechanisms for indirectly increasing evolvability have been explored [12, 16, 18, 20, 27, 30]. Some approaches change the fitness function over time in various ways to encourage the ability for genomes to more quickly adapt to such changes [12, 20, 30]. Another approach is to include a vast number of objectives, which approximates the infinite number of niches created by the ever-changing natural world [27]. Further approaches include Novelty Search [18], and periodically extinguishing most organisms in niches [16]. While such indirect methods can be effective, an unexplored possibility is instead to *directly* select for evolvability. This paper introduces *Evolvability Search*, a new class of evolutionary algorithm wherein the fitness function is a direct measure of an individual's evolvability.

There is no overall agreement on evolvability's definition or measurement [28], which complicates operationalizing it as a criterion for optimization. The definition adopted here follows one mainstream conception of evolvability as phenotypic variability (irrespective of fitness) [3, 13, 28], i.e. the tendency of an individual's offspring to exhibit phenotypic diversity. However, it is important to acknowledge that there are other definitions, and that evolvability itself may represent a cluster of related concepts [28]. For example, another major class of definitions stress evolvability's relation to adaptation, e.g. the rate of adaptation or the production of *adaptive* variation [35]. While distinct, these conceptions are also connected: Increasing the ability to produce phenotypic variation can also increase how quickly a lineage can adapt [3, 13, 16]. One advantage of the definition chosen here is that it affords a concrete measure of evolvability that is independent of the particular goal of search, i.e. how much phenotypic variety exists within an organism's mutational neighborhood [13, 16, 18, 19]. Evolvability Search optimizes this measure directly by generating many offspring from

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. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '16, July 20-24, 2016, Denver, CO, USA

© 2016 ACM. ISBN 978-1-4503-4206-3/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2908812.2908838>

an individual, and setting its fitness to a measure of the phenotypic variety among the offspring.

The primary purpose of Evolvability Search is to provide a tool for researchers in EC to more easily produce evolvable organisms, thereby aiding the study of evolvability. An additional motivation is to investigate whether searching *only* for evolvability is itself an effective search algorithm. That is, optimizing evolvability alone may indeed result in increased evolvability, but such search may not be effective for solving tasks. However, the pressure to increase diversity among an individual’s offspring will also increase population diversity as a byproduct, i.e. individuals selected for reproduction by definition will have more diverse offspring. As a result, Evolvability Search may in performance resemble diversity-seeking algorithms such as Novelty Search, which are often effective in practice [7, 25, 26, 31].

We investigate the potential of the proposed method in two Evolutionary Robotics domains from previous evolvability studies [16, 18], i.e. maze navigation and biped locomotion. In these domains, we compare the performance of Evolvability Search with two other Evolutionary Algorithms (EAs), both a traditional objective-driven algorithm and Novelty Search, which has previously been shown to increase evolvability [16, 18]. We test the hypothesis that Evolvability Search produces more evolvable solutions (as measured by the ability to generate phenotypic variability) as well as the hypothesis that Evolvability Search can effectively solve tasks. Finally, we test whether the evolvability produced is *general*, by transferring organisms from the environment they were evolved within to a new environment, and validating that they remain evolvable after such transfer. Indeed, the experimental results show that Evolvability Search significantly increases evolvability and produces *generalizable* evolvability. Similar to other divergent EAs [7, 25, 26, 31], it also increases population diversity and evolves solutions to deceptive problems.

Because evolvability has proven difficult to produce indirectly, scientists have lacked the ability to study highly evolvable, digital phenotypes. Our results show that through direct selection, Evolvability Search can produce highly evolvable organisms, thereby providing scientists the opportunity to study them. Such opportunity may aid both the design of EAs and the study of evolvability itself. In this way, Evolvability Search may serve as a stepping stone for considering the relationship between EAs and evolvability in new ways. What distinguishes this algorithm is that it is explicitly ‘long-sighted,’ unlike natural evolution and most EAs, which normally short-sightedly select by criteria related only to current performance (and not future potential). The main conclusion is that Evolvability Search is a helpful tool for producing and *studying* evolvable organisms; thus it may also help stimulate further algorithms able to more efficiently generate evolvability, which could prove helpful across EC as a whole.

2. BACKGROUND

This section reviews previous approaches to encourage evolvability in EC and also the neuroevolution algorithm employed in the experiments.

2.1 Evolvability in Evolutionary Computation

Reproducing the evolvability of natural organisms is an important yet unmet research goal in EC. Accordingly, re-

searchers have proposed many approaches for encouraging evolvability [5, 6, 12, 16, 18, 20, 27, 30]. However, most such approaches encourage evolvability only through *indirect* mechanisms. For example, Novelty Search selects for individuals demonstrating novel behavior relative to those previously produced by the search [17]; rewarding novelty indirectly encourages evolvability because mechanisms that consistently enable novelty (and will thereby be consistently selected) also enable phenotypic variability [18]. A related approach is to set a vast number of separate, diverse evolutionary objectives, which approximates the infinite number of niches created by the ever-changing natural world [27]. Such optimization leads to frequent “goal-switching,” as lineages fit on one objective invade other objectives, which rewards lineages that produce behaviorally diverse offspring and increases evolvability [27].

Another family of approaches encourage evolvability by systematically altering the fitness function during evolution [12, 30]. For example, the Modularly Varying Goals idea creates indirect selective pressure for evolvability by repeatedly changing the evolutionary goal in a modular way [12]. The motivation is that to persist consistently through changing conditions requires the meta-ability to adapt more quickly between such conditions. Imposing periodic catastrophic events similarly create pressure to adapt more quickly [16]; such events indiscriminately wipe out most of the organisms in niches, thereby creating evolutionary bottlenecks. Evolvable organisms are indirectly favored, because lineages able to radiate through multiple niches have better chances of persisting across such bottlenecks.

A final approach is to add a cost for connections within evolved networks, which encourages sparsity, modularity, and hierarchy. Such properties in turn improve evolvability [5, 9, 11, 22]. This connection cost technique is inspired by the energy costs paid by biological organisms to build, use, and maintain components such as neural connections. Such connection costs may also serve function similar to regularization in machine learning, creating a pressure towards parsimony that prevents over-specific solutions and promotes generalization to new environments (a form of evolvability) [14].

Interestingly, however, to the best of our knowledge, past studies have not investigated the effects of increasing evolvability through *direct* selection, i.e. by quantifying evolvability to serve as the fitness function. Exploring such an approach is the primary contribution of this paper.

2.2 NeuroEvolution of Augmenting Topologies

In the experiments in this paper, robots are controlled by Artificial Neural Networks (ANNs) evolved through the widely used [1, 2, 32, 33] NeuroEvolution of Augmenting Topologies (NEAT) technique [32, 33]. Motivating its choice, NEAT has been used in several previous EC studies of evolvability [16, 18, 19]. Note that this section only briefly reviews NEAT; for further details see [32, 33].

The NEAT algorithm starts from a population of simple ANNs and complexifies them during evolution through mutations that add nodes and connections. The algorithm divides the population into distinct niches in a process called speciation [32, 33]. Each species consists of ANNs that share common ancestry and network topology. Such speciation allows individuals with novel ANN structures an opportunity to adapt by limiting competition to be only *within* species. In this way, speciation avoids premature extinction of poten-

tially promising ANN structures. In NEAT, new nodes and connections are assigned a unique marking number, which allows NEAT to perform meaningful crossover. That is, aligning genes with the same historical markings enables connections to be swapped that are likely to serve similar functional roles. The NEAT implementation from [18] is adapted to implement Evolvability Search.

3. EVOLVABILITY SEARCH

Evolvability Search is motivated by the important biological property of evolvability, which is a key enabler of nature’s ability to create complex, diverse creatures [3, 8]. The main idea is to select more evolvable individuals directly, i.e. the fitness function quantifies an individual’s evolvability.

The evolvability measure in this paper is adopted from previous evolvability studies [16, 18, 19] and stems from one common conception of evolvability, i.e. an organism’s ability to generate heritable phenotypic variation (irrespective of fitness) [3]. We approximate an individual’s capacity to generate future phenotypic variation by measuring phenotypic variability among a sample of the individual’s simulated offspring (which are discarded). Such variability is quantified as the number of unique behaviors; in particular, each offspring is considered sequentially and added to a list of unique behaviors only if its behavior is significantly different from the behaviors of organisms already in the list. Two behaviors are considered different if the distance between them according to a domain-specific behavioral distance metric is above a pre-specified threshold (here 0.01).

Evolvability Search, like divergent search algorithms [18], does not directly pursue the underlying objective. Instead it attempts to broadly reward intermediate stepping stones that could potentially lead to it. However, the heuristic of Evolvability Search differs from that of methods like Novelty Search because it selects individuals with a greater potential for diversity, rather than diversity itself.

Interestingly, the pressure to increase diversity among an individual’s offspring indirectly elevates population diversity. That is, selected individuals will by definition tend to have more diverse offspring, which will form the next generation’s population, and thus further diversify the population. This tendency to diversify suggests that Evolvability Search may share performance characteristics with diversity-driven algorithms, which often perform well when applied to deceptive problems [7, 24–26, 31].

One clear liability of Evolvability Search is that fitness calculation is computationally expensive; one must evaluate a sufficient sample of offspring to estimate a single individual’s evolvability. However, the algorithm can be easily parallelized and can therefore leverage benefits from increasing computational power and multi-core machines. While computational efficiency is often an important consideration, ultimately in some domains—or for some types of research (e.g. into the nature of evolvability)—achieving highly-evolvable solutions may be more important. The next section introduces experiments that highlight the potential of Evolvability Search to achieve such ends.

4. EXPERIMENTS

The experiments described in this section compare Evolvability Search with two representative search methods also used in previous evolvability studies [16, 18]: (1) traditional

objective-based search, which is a baseline control; and (2) Novelty Search, which is representative of divergent search algorithms. All three algorithms are compared in the two evolutionary robotics domains described next, i.e. maze navigation and biped locomotion.

4.1 Maze Navigation Experiment

The maze navigation domain is an appropriate domain for testing Evolvability Search because it is complex enough to be interesting, but is also easy to visualize and understand [17, 23]. Furthermore, the domain is deceptive (i.e. it is difficult to solve by only moving closer to the goal), a property that commonly complicates search in complex domains [17].




Figure 1: The neural network controller (a) and map (b) for the maze navigation domain. Robots must navigate the map from the start location (larger circle) to the goal location (smaller circle). Figures from [18].

In the maze navigation domain, a simulated robot controlled by an ANN (Fig. 1a) must successfully navigate a maze (the so called “hard maze” [18]; Fig. 1b) within 400 time steps. The robot (Fig. 2) has six rangefinder sensors that enable it to estimate its distance to nearby walls, and has four pie-slice radar sensors that activate when the goal is detected within a pie-slice. The robot also has two effectors, one enabling it to turn left or right, and the other to move forward or backward.




Figure 2: The maze-navigating robot. The outgoing arrows indicate the six rangefinder sensors, which allow the robot to estimate its distance to nearby obstacles. At the front, back, and both sides of the robot, there are four pie-slice sensors that activate when the goal falls within that sensor’s angular range. The robot’s heading is indicated by the solid, blue arrow. Figure from [18].

The objective-based search algorithm employs a fitness function that favors robots that end up closer to the goal. The fitness f of a robot is given by: $f = b_f - d_g$, where b_f is a constant value that ensures positive fitness, and d_g is the robot’s final distance to the goal at the end of its evaluation.




Figure 3: A simulated biped robot. In the biped locomotion domain, artificial neural networks are evolved to control the robot shown in this visualization. Figure from [18].

Novelty Search instead applies a novelty metric to quantify how novel a robot’s behavior is relative to the current and past populations. As in [17], the behavior of a robot is defined as its final Cartesian location within the maze. For a new individual, the novelty metric computes the average distance between it and its k nearest neighbors. Such calculation requires a distance metric, which in this domain is defined as the squared Euclidean distance between the final locations of any pair of robots. Evolvability Search applies this same behavioral distance to determine the number of distinct behaviors among an organism’s offspring, i.e. the organism’s evolvability.

For all algorithms, 200 offspring are generated and evaluated to calculate the evolvability of an organism. Similar to other evolvability studies in the maze navigation domain [16, 18], the evolvability of the population in the Novelty and objective-based treatments is measured every 100 generations. The population size is 250 and each run lasts 1000 generations. A maze-navigating solution is defined as one that finishes within five units of the goal. All other parameters are the same as in [18].

4.2 Biped Locomotion Experiment

The second domain, biped locomotion, is adapted from [17]. Neural networks are evolved with the goal of controlling a simulated biped robot to walk as far as possible (Fig. 3). The domain is challenging because the controller must combine balancing with oscillatory motion to produce a stable walking gait [17].

The simulated biped robot has two degrees of movement in each hip joint (pitch and roll) and one degree in each knee joint (pitch), resulting in a total of six degrees of freedom. The robot is controlled by a Continuous-Time Recurrent Neural Network (CTRNN). CTRNNs, which are commonly applied in such experiments [21, 29], are suitable for expressing cyclic dynamics like those found in natural gaits. Reflecting the goal of an effective robot gait, the fitness function for objective-driven search is how far the robot walks, measured as the Euclidean distance from the robot’s start location. Evaluation of a controller is terminated if the robot falls or when the allotted time period (15 seconds) expires.

For Novelty and Evolvability Search, the behavioral characterization is formed by sampling the biped’s center of mass once per second. Each biped’s behavior is represented by 15 points within an $x - y$ plane (ignoring the vertical z component). Temporal information is recorded in this domain to enable Novelty and Evolvability Search because temporal dynamics are a key aspect of walking gaits. For this domain, behavioral distance is measured as the square root of the sum of squared differences between two individuals’ behavioral characterizations, as is in [17].

Following [18], in this domain NEAT is augmented with self-adapting mutation rates, which enable evolution to con-

trol how strong mutations affect particular connections. The reason is that self-adaptation has been shown to enable greater capacity for evolvability in the biped domain [16, 18]. To implement self-adaptation, the NEAT genome is extended with three pairs of mutation settings, which themselves mutate at fixed rates (to escape infinite regress) [4, 15]. Each pair contains two parameters: a connection mutation power, which regulates the *magnitude* by which a mutation changes the weight of a connection; and a connection mutation rate, which encodes the *probability* of a mutation occurring to a particular connection. Thus, with self-adaptation, a connection could be exempted from mutational changes, or conversely, prone to many more (and intense) mutations. All self-adaptation parameters are the same as in [18]. See [18] for a detailed description of this feature.

The population size for the biped experiment is 500 individuals and evolution proceeds for 2,000 generations. When measuring evolvability, 200 offspring are created from each individual. As in [18], the evolvability of individuals in the Novelty Search and objective-driven treatments is measured every 200 generations. The robot that travels the farthest in each run is considered the solution for that run. Other algorithmic parameters are the same as in [18], and all biped simulation parameters are from [17].

5. RESULTS

5.1 Maze Navigation Results

Deception in the hard maze prevents objective-driven search from often solving the problem [17], a result replicated here: Objective-based search solved the task only three times in 50 runs. In contrast, Evolvability Search produced solutions in 49 out of 50 runs. Moreover, individuals in Evolvability Search exhibited diverse navigational behaviors (Fig. 4a), affirming that population diversity increases as a byproduct of encouraging evolvability. As in previous work [17], Novelty Search evolved solutions in all 50 runs. The number of successful runs of Evolvability Search and Novelty Search was significantly higher than that of objective-driven search ($p < 0.0001$, Fisher’s exact test), but did not significantly differ from each other.

Fig. 4b compares the average evolvability of populations across different methods, showing that Evolvability Search produces significantly more evolvable populations ($p < 0.05$). Similarly, if the maximum evolvability of individuals produced during runs is compared across methods, Evolvability Search again outperforms the controls (data not plotted, $p < 0.001$). To compare the evolvability of solutions, i.e. individuals that solve the maze, we performed as many runs as were necessary to obtain 45 distinct successful runs per method. No additional runs were required for Evolvability Search or Novelty Search (solutions are taken from the first 45 runs). However, 550 additional runs were required for objective-based search (thus 600 objective-based runs were conducted in total). As shown in Fig. 4c, Evolvability Search again produces significantly higher evolvability by this metric. Note that unless otherwise stated, all statistical tests are Mann-Whitney-Wilcoxon rank-sum tests.

Another aspect of performance is the number of evaluations taken to produce solutions (irrespective of such solutions’ evolvability). The average number of macro-evaluations (i.e. counting only the number of organism evaluations, which discounts the 200 offspring evaluations required for each or-




Figure 4: Maze navigation results. (a) The final locations visited by individuals are shown for one final population from Evolvability Search. The distribution of individuals shows that the algorithm avoids deception and can produce controllers that solve the problem. (b) Evolvability Search produces significantly higher evolvability than Novelty Search and objective-based search. As previously shown [18], Novelty Search produces significantly more evolvable populations than objective-based search. The curves represent means over all runs. Significant differences ($p < 0.05$) are shown comparing Evolvability Search to other methods, and are indicated below the plot every 100 generations for Novelty Search (circles) and objective-based search (squares). (c) Averaging evolvability over *solutions* (the first organism from each run that reaches the goal) instead of over the population, produces the same qualitative results.

ganism in Evolvability Search) needed for objective-based search to solve the task is significantly higher than that required by Novelty Search or Evolvability Search (Fig. 5a, $p < 0.001$). However, the macro-evaluations needed to solve the task do not differ significantly between Novelty and Evolvability Search (Fig. 5a, $p > 0.05$). A more accurate reflection of computational expense is given by Fig. 5b, which counts how many total evaluations of genomes in the maze simulation are conducted, i.e. the offspring evaluations in Evolvability Search are also accounted for. By such accounting, Evolvability Search requires significantly more evaluations to produce solutions than Novelty Search or objective-based search ($p < 0.001$).

Because Evolvability Search is afforded more computation per generation, we tested whether Novelty Search or objective-based search would produce comparable levels of evolvability if they were run longer. We thus performed extended experiments that lasted 5,000 generations (the default runs lasted 2,000 generations), which is as many as were possible given limited computational resources. A plot of evolvability over time shows that average evolvability in both Novelty Search and objective-based search appears to plateau at levels much lower than that produced by Evolvability Search. Thus the expectation is that such search methods could not match Evolvability Search even if granted the same amount of computational resources (Fig. 6a).

5.2 Biped Locomotion Results

In the biped locomotion domain, Evolvability Search produced controllers able to travel significantly farther than the best from objective-based search (Fig. 7a, $p < 0.001$). However, Novelty Search produced controllers significantly outperforming both Evolvability Search and objective-based search (Fig. 7a, $p < 0.001$). As in the maze domain, however, Evolvability Search produced significantly higher evolvability than the control methods, whether measured as the average evolvability of individuals in the population (Fig. 7b, $p < 0.05$), the maximum evolvability individual evolved in a given run ($p < 0.001$), or in champion individuals only




Figure 5: Evaluations required to solve the hard maze. (a) Objective-based search requires significantly more macro-evaluations than Novelty Search and Evolvability Search. The difference between the number of macro-evaluations required by Evolvability Search and Novelty Search is not significant. (b) However, the number of effective evaluations (i.e. actual maze simulations) required by Evolvability Search is significantly higher than both Novelty Search and objective-based search. Significance is measured by the Mann-Whitney-Wilcoxon rank-sum test.

(Fig. 7c, $p < 0.001$). These results suggest that in some domains there may be a trade-off between performance and evolvability. As in the maze navigation domain, when extending the duration of evolution for Novelty Search and objective-based search, evolvability seems to plateau below that produced by Evolvability Search (Fig. 6b). Overall, the biped domain results are qualitatively similar to those in the maze navigation domain, suggesting they may represent a general tendency.

6. EVolvABILITY GENERALITY TEST

While the measure of evolvability adopted here (i.e. phenotypic variability within a fixed environment) captures one important aspect of evolvability, another intuition about evolvability is that it enables individuals to adjust more quickly to new environments, i.e. adaptability. One possi-

bility is that the offspring from an evolvable individual will yield a greater diversity of behaviors than a less evolvable one even when both are transferred to a new environment. Importantly, among such greater variety may be found the seeds of useful adaptations. Because adaptability is an important property in EC, this section describes a test designed to investigate whether evolvability particular to one environment also generalizes to one unseen by evolution. The original and new environments are analogous to the concepts of *training set* and *testing set* in machine learning.

As an initial exploration, this evolvability generality test is performed only in the maze navigation domain. We extract the most evolvable organism from each run for each treatment evolved on the hard maze. For each organism, we generate 200 offspring and evaluate them without further evolution in a new maze (the mega maze [34]; Fig. 8a). The name of the mega maze reflects that is more than twice as large (450 units \times 450 units) as the hard maze (200 units \times 200 units). The mega maze also contains many independent pathways. Thus, it is unlikely that a controller that has memorized a particular trajectory on the hard maze will result in substantial exploration of the mega maze.

If an organism’s evolvability generalizes to the new environment, then the behaviors exhibited by its offspring will be more diverse. One way to measure behavioral diversity in the new environment is to count how many areas of the mega maze are explored by the 200 transferred offspring of each evolvable organism [34]. Following [34], we overlay a 450×450 grid on the mega maze and track the percentage of grid cells that are visited at least once by any of an organism’s transferred offspring. This metric captures some aspect of generalizable evolvability, because it measures the phenotypic diversity of an organism’s offspring in a novel environment. The results of this test show that Evolvability Search scores higher on this metric than both Novelty Search and objective-based search ($p < 0.001$, Fig. 8b–d & 9)




Figure 6: Extended evolutionary runs suggest that the evolvability produced by Novelty Search and objective-based search would not rise to the level of that produced by Evolvability Search irrespective of how long the algorithms are run. Average evolvability is shown per generation for both (a) the maze-navigation domain and (b) the biped locomotion domain. The dashed line at the top of each plot indicates the level of evolvability that Evolvability Search produced after 2000 generations.

7. DISCUSSION

The results of these experiments highlight that directly selecting for evolvability, although computationally expensive, is a viable approach to produce evolvable and adaptable individuals. Because there remains uncertainty about the exact indirect selective forces by which evolvability emerges

in natural evolution, replicating nature’s evolvability in EC remains an unmet challenge [5, 16, 18, 35]. However, unlike in nature, the results here show that EAs can gain advantage from leveraging foresight to select individuals with greater potential to generate future diversity.

Furthermore, there may be ways to apply Evolvability Search that reduce computational cost. For example, it may benefit indirect encodings to encourage evolvability of *neural structures*, i.e. in what variety of ANN topologies are reachable from a current individual. While the mapping between ANN structure and behavior is inconsistent, the cost of computing ANN evolvability is likely trivial compared to full behavioral evaluations of offspring; and such neural evolvability may be an important ingredient to improve search with indirect encodings.

Evolvability Search may also be applied intermittently, especially when domain evaluations are computationally expensive. For example, solutions produced through other algorithms (e.g. through Novelty Search) could be refined afterwards, applying Evolvability Search to optimize evolvability under the constraint of retaining the ability to solve the task. The organisms could then be further evolved in the same or different environments.

Another interesting avenue for future work is to encourage different distributions of offspring behaviors. That is, instead of maximizing variance in offspring behaviors, novelty can be maximized in the behavioral *distributions* of offspring. Doing so may cause organisms to canalize in different ways (e.g. some produce limbs whose lengths are all correlated, others produce limbs with uncorrelated limb lengths, still others could correlate hind legs, but not forelegs, etc.). The overall product may be a significantly more evolvable *population*, consisting of evolvable individuals predisposed to different types of variation [36]. In this way, explicitly encouraging search in the space of canalizations may benefit search, reflecting the perceived importance of canalization for evolvability in biology [10].

Apart from direct practical applications, the algorithm may help shed light on the structure of evolvability. For example, optimizing evolvability in more ambitious domains may reveal if evolvability is often a deceptive property. In other words, do the stepping stones to increased evolvability tend to be evolvable? With Evolvability Search it may be possible to empirically examine how often the gradients of evolvability align with the gradients of novelty or of objective-driven fitness. In this way, Evolvability Search may be a useful tool to help understand the evolution of evolvability and how our current algorithms encourage it (or why they fail to). Such understanding could lead to further algorithmic improvements in EC.

The cost of Evolvability Search may be somewhat mitigated by the trend of increasing computational power and access to many CPU cores, because the algorithm is trivially parallelizable. Regardless, our results show that the computational cost may be necessary if researchers wish to reach levels of evolvability higher than those produced by other known algorithms; in some cases such computational cost may be a small price to pay relative to the benefits of easily and consistently generating highly evolvable organisms, which also enables their study.

Finally, some readers may object to the definition of evolvability adopted here because it measures phenotypic variation without including any measure of adaptation or fitness. We




Figure 7: Biped Locomotion Results. (a) Novelty Search produces robots that travel significantly farther than those from Evolvability Search and objective-based search. However, impressively, Evolvability Search does significantly outperform objective-based search. (b) Most importantly, Evolvability Search produces significantly more evolvability than Novelty Search or objective-based search. That result holds when averaged across the entire population (b) and when averaged across the best organism from the final population of each run (c). As in [18], the corresponding statistics for Novelty Search are significantly higher than objective-based search. Plotted are means over all runs. p values < 0.05 vs. Evolvability Search are indicated below the plot for Novelty Search (circles) and objective-based search (squares) every 200 generations.

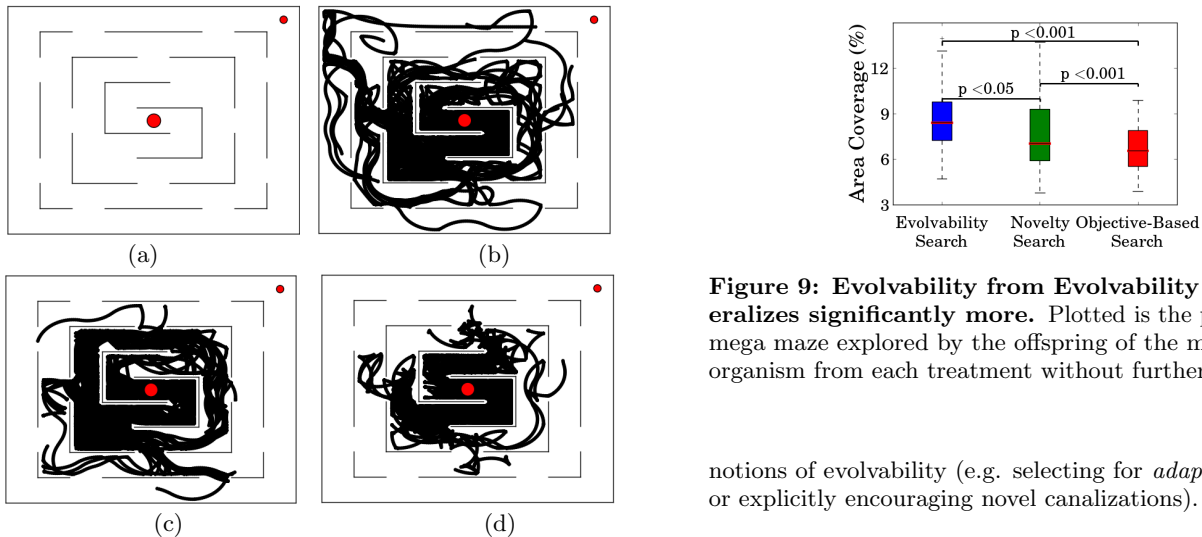


Figure 8: Evolvability Generality Test. (a) The mega maze. The larger and smaller circles represent the start and goal, respectively. The area covered by offspring of the most evolvable Evolvability Search individual (b) is larger than for Novelty Search (c) and objective-based search (d).

note that Evolvability Search could also be with alternative definitions of evolvability. While we chose the definition in this paper because it is principled, easy to measure, and enjoys support in the biological literature, an exciting research question for future work is to how to encourage other types of evolvability; such alternate evolvability measures could include aspects of improving adaptability both within an environment or problem, or to new ones. For example, the current method selects for maximizing production of *any* type of variation. However, such impartiality may undermine evolving useful constraints on such variation, like canalizing useful dimensions of variation and thereby reducing the generation of unhelpful behaviors. Thus an interesting idea for future work is to investigate how to promote more complex

Figure 9: Evolvability from Evolvability Search generalizes significantly more. Plotted is the percent of the mega maze explored by the offspring of the most evolvable organism from each treatment without further evolution.

notions of evolvability (e.g. selecting for *adaptive* variation or explicitly encouraging novel canalizations).

8. CONCLUSIONS

This paper introduces Evolvability Search, a method for producing evolvability by selecting for it directly. The results demonstrate that Evolvability Search is an effective method to produce organisms and solutions with increased evolvability, thereby providing scientists with a new tool to create and study more evolvable organisms. By introducing an EA with explicit foresight, the hope is also to foster research into new kinds of effective long-sighted algorithms. While computationally expensive, Evolvability Search may be the most viable approach when highly evolvable solutions and populations are essential.

Acknowledgements: Support came from an NSF CAREER award (CAREER: 1453549) to Jeff Clune and via computers provided by the US NSF and DOE’s Open Science Grid. We thank J.Huizinga and R.Velez for helpful feedback and discussions.

9. REFERENCES

- [1] T. Aaltonen et al. Measurement of the top-quark mass with dilepton events selected using neuroevolution at CDF. *Physical Review Letters*, 102(15):152001, 2009.

- [2] B. Allen and P. Faloutsos. Complex networks of simple neurons for bipedal locomotion. In *Intelligent Robots and Systems (IROS)*, pages 4457–4462. IEEE, 2009.
- [3] J. Brookfield. Evolution: the evolvability enigma. *Current Biology*, 11(3):R106–R108, 2001.
- [4] J. Clune, D. Misevic, C. Ofria, R. Lenski, S. Elena, and R. Sanjuán. Natural selection fails to optimize mutation rates for long-term adaptation on rugged fitness landscapes. *PLoS Computational Biology*, 4(9):e1000187, 2008.
- [5] J. Clune, J.-B. Mouret, and H. Lipson. The evolutionary origins of modularity. *Proceedings of the Royal Society B*, 280(20122863), 2013.
- [6] J. Clune, K. Stanley, R. Pennock, and C. Ofria. On the performance of indirect encoding across the continuum of regularity. *IEEE Transactions on Evolutionary Computation*, 15(4):346–367, 2011.
- [7] S. Doncieux and J.-B. Mouret. Behavioral diversity measures for evolutionary robotics. In *Evolutionary Computation (CEC), 2010 IEEE Congress on*, pages 1–8. IEEE, 2010.
- [8] D. J. Earl and M. W. Deem. Evolvability is a selectable trait. *Proceedings of the National Academy of Sciences of the USA*, 101(32):11531–11536, 2004.
- [9] K. O. Ellefsen, J.-B. Mouret, J. Clune, and J. C. Bongard. Neural modularity helps organisms evolve to learn new skills without forgetting old skills. *PLoS Comput Biol*, 11(4):e1004128, 2015.
- [10] S. Gould. *Ontogeny and phylogeny*. Harvard University Press, 1977.
- [11] J. Huizinga, J.-B. Mouret, and J. Clune. Evolving neural networks that are both modular and regular: Hyperneat plus the connection cost technique. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 697–704, 2014.
- [12] N. Kashtan and U. Alon. Spontaneous evolution of modularity and network motifs. *Proc. of the National Academy of Sciences*, 102(39):13773–13778, 2005.
- [13] M. Kirschner and J. Gerhart. Evolvability. *Proc. of the National Academy of Sciences*, 95(15):8420–8427, 1998.
- [14] K. Kouvaris, J. Clune, L. Kounios, M. Brede, and R. A. Watson. How evolution learns to generalise: Principles of under-fitting, over-fitting and induction in the evolution of developmental organisation. *arXiv preprint arXiv:1508.06854*, 2015.
- [15] O. Kramer. Evolutionary self-adaptation: a survey of operators and strategy parameters. *Evolutionary Intelligence*, 3(2):51–65, 2010.
- [16] J. Lehman and R. Miikkulainen. Enhancing divergent search through extinction events. In *Proceedings of the 2015 on Genetic and Evolutionary Computation Conference*, pages 951–958. ACM, 2015.
- [17] J. Lehman and K. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, 2011.
- [18] J. Lehman and K. O. Stanley. Improving evolvability through novelty search and self-adaptation. In *Evolutionary Computation (CEC), 2011 IEEE Congress on*, pages 2693–2700. IEEE, 2011.
- [19] J. Lehman and K. O. Stanley. Evolvability is inevitable: increasing evolvability without the pressure to adapt. *PLoS one*, 8(4):e62186, 2013.
- [20] P. J. B. Lipson, H. and N. P. Suh. On the origin of modular variation. *Evolution, International Journal of Organic Evolution*, 56(8):1549–1556, 2002.
- [21] G. McHale and P. Husbands. Gasnets and other evolvable neural networks applied to bipedal locomotion. *From Animals to Animats*, pages 163–172, 2004.
- [22] H. Mengistu, J. Huizinga, J.-B. Mouret, and J. Clune. The evolutionary origins of hierarchy. *PLoS Computational Biology*, To appear, 2016.
- [23] J.-B. Mouret. Novelty-based multiobjectivization. In *New horizons in evolutionary robotics*, pages 139–154. Springer, 2011.
- [24] J.-B. Mouret and J. Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.
- [25] J.-B. Mouret and S. Doncieux. Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity. In *IEEE Congress on Evolutionary Computation*, pages 1161–1168. IEEE, 2009.
- [26] J.-B. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary computation*, 20(1):91–133, 2012.
- [27] A. Nguyen, J. Yosinski, and J. Clune. Innovation engines: Automated creativity and improved stochastic optimization via deep learning. In *Proc. of the Genetic and Evolutionary Computation Conference*, 2015.
- [28] M. Pigliucci. Is evolvability evolvable? *Nature Reviews Genetics*, 9(1):75–82, 2008.
- [29] T. Reil and P. Husbands. Evolution of central pattern generators for bipedal walking in a real-time physics environment. *Evolutionary Computation, IEEE Transactions on*, 6(2):159–168, 2002.
- [30] J. Reisinger and R. Miikkulainen. Selecting for evolvable representations. In *Proceedings of the 8th annual conference on Genetic and evolutionary computation*, pages 1297–1304. ACM, 2006.
- [31] S. Risi, S. D. Vanderbleek, C. E. Hughes, and K. O. Stanley. How novelty search escapes the deceptive trap of learning to learn. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 153–160, 2009.
- [32] K. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.
- [33] K. Stanley and R. Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 21(1):63–100, 2004.
- [34] R. Velez and C. J. Novelty search creates robots with general skills for exploration. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pages 737–44, 2014.
- [35] G. Wagner and L. Altenberg. Complex adaptations and the evolution of evolvability. *Evolution*, 50(3):967–976, 1996.
- [36] B. Wilder and K. Stanley. Reconciling explanations for the evolution of evolvability. *Adaptive Behavior*, 23(3):171–179, 2015.