

Enhancing Divergent Search through Extinction Events

Joel Lehman
University of Texas at Austin
joel@cs.utexas.edu

Risto Miikkulainen
University of Texas at Austin
risto@cs.utexas.edu

ABSTRACT

A challenge in evolutionary computation is to create representations as evolvable as those in natural evolution. This paper hypothesizes that extinction events, i.e. mass extinctions, can significantly increase evolvability, but only when combined with a *divergent* search algorithm, i.e. a search driven towards diversity (instead of optimality). Extinctions amplify diversity-generation by creating unpredictable evolutionary bottlenecks. Persisting through multiple such bottlenecks is more likely for lineages that diversify across many niches, resulting in indirect selection pressure for the capacity to evolve. This hypothesis is tested through experiments in two evolutionary robotics domains. The results show that combining extinction events with divergent search increases evolvability, while combining them with convergent search offers no similar benefit. The conclusion is that extinction events may provide a simple and effective mechanism to enhance performance of divergent search algorithms.

Categories and Subject Descriptors

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

Keywords

Extinction events; evolvability; divergent search

1. INTRODUCTION

Biological organisms are evolvable, i.e. they have significant capacity to evolve further [6, 30]. In contrast, evolvability remains an ambitious challenge in evolutionary computation (EC) [4, 18, 30], hinting that some important mechanisms of natural evolution may be missing. Confounding the issue, some such mechanisms may be synergistic and provide little benefit in isolation. In such cases, the usual methodology of introducing and testing mechanisms individually may fail to identify promising modifications to evolutionary algorithms (EAs).

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 '15, July 11 - 15, 2015, Madrid, Spain

© 2015 ACM. ISBN 978-1-4503-3472-3/15/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2739480.2754668>

This paper hypothesizes that *extinction events*, i.e. mass extinctions, are such a mechanism. Extinctions can increase evolvability, but only when coupled with a divergent reward scheme, which is another non-canonical EA mechanism. That is, most EAs aim to optimize towards an underlying fixed objective, which leads to *convergence* to a single phenotype. In contrast, natural evolution inherently *diverges*, accumulating a growing diversity of novel solutions to the problems of surviving and reproducing [27]. Extinction events may enhance the effectiveness only of EAs that similarly diverge.

Although well-studied in evolutionary biology, extinction events are rarely included in EAs [8, 14]. One reason is that in the short term they are often idiosyncratic and destructive, selecting individuals arbitrarily or for reasons uncharacteristic of its evolutionary history as a whole. In this view, extinction events may act merely as upheavals or impediments to the evolutionary process. Thus beyond increased biological accuracy, there is little reason to model them.

However, although extinction events exterminate stochastically a large proportion of ways of life, they may still produce beneficial evolutionary regularities. In particular, extinction events can create unpredictable evolutionary bottlenecks by filtering out most pre-existing phenotypic niches. Lineages that are able to diversify more quickly than others, i.e. those that are more evolvable, may have a greater chance of persisting across multiple such bottlenecks. By radiating through multiple niches, such evolvable lineages in effect have multiple tickets in the lottery of extinction. Thus extinction events indirectly select for evolvability.

Interestingly, diversity-driven EAs also select for evolvability, but through a different mechanism. For instance, if selection rewards novelty, those lineages that *consistently* produce diversity over time will be selected. In this way, adaptations enabling evolvability hitchhike through the immediate advantage they provide in selection.

Because extinction events and divergent EAs encourage evolvability in a different manner, there may be an additive benefit when they are combined. In contrast, convergent search, such as a traditional objective-seeking EA, does not necessarily encourage evolvability [18]. Thus the hypothesis of this paper is that combining extinction events with divergent EAs can increase evolvability, whereas no such benefit may arise when extinctions are combined with a traditional convergent EA.

To test this hypothesis, representative divergent and convergent search methods are augmented with extinction events and compared to unaugmented ones in two evolution-

ary robotics (ER) domains. The results demonstrate that when divergent EAs are combined with extinction events, the result is more evolvable populations and more effective solutions. In contrast, the tested convergent EA shows no similar benefit. In this way, extinction events may provide a simple and powerful mechanism to enhance divergent EAs.

2. BACKGROUND

This section first reviews the neuroevolution technique and the measure of evolvability applied in the experiments. Extinction in EC and divergent search are then described.

2.1 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] and well understood.

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 [12] introduced before it, NEAT begins evolution with a population of small, simple networks and *complexifies* them. The population expands into diverse with diverse ANN topologies over generations, leading to increasingly sophisticated behavior. A similar process of gradually adding new genes is seen in natural evolution [21]. 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 assigned uniquely 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 [21], is thus supported by both historical markings and speciation, allowing NEAT to establish high-level features early in evolution and then later to elaborate on them.

Note that speciation and crossover in NEAT are disabled in this paper’s experiments to simplify the EA. This simplification facilitates integration of NEAT with the behavioral grid search algorithm described later, which allows uniform comparison between all implemented search algorithms. Additionally, NEAT is augmented in all experiments with self-adaptive mutation rates, as supported by Lehman and Stanley [18]. Such self-adaptation was found to enhance performance in initial experiments.

To implement self-adaptation, each genome is augmented with a list of three pairs of mutation settings, and each connection is augmented with an integer parameter that indexes within the list. Typically, when an individual is mutated, every one of its connections are altered by adding to its weight a random number chosen from the same uniform distribution. In contrast, with self-adaptation, each connection can be *exempted* from mutation with probability determined by a floating point parameter in its associated mutation set-

tings. If it is not exempted from mutation, the uniform distribution from which the weight is perturbed is scaled by a separate floating point parameter in its associated mutation settings. In this way, an individual can encode an architecture of mutation that potentially complements its evolved representation, thus making it more evolvable.

2.2 Measuring Evolvability

Natural evolution has produced flexible, highly evolvable representations that facilitate its prolific discovery of diverse organisms. An important question that could inform EC is: What properties of natural evolution led to such evolvability? Investigating this question is aided by quantitative measures of evolvability. Such measures enable empirical studies of how different evolutionary features affect evolvability.

While there is no overall consensus on evolvability’s definition or its measurement [26], one common conception is to consider evolvability as an organism’s phenotypic variability [3, 7, 13, 30]; that is, the capacity of an organism’s lineage to generate novel phenotypic traits captures some significant part of what enables some lineages to adapt more quickly than others, although there exist alternative definitions that focus on different or overlapping aspects of evolvability [26]. This conception (of evolvability as phenotypic variability) is adopted in this paper, as in previous related studies [18, 19].

The evolvability measure used in this paper was originally proposed by Lehman and Stanley [18]. At regular intervals during evolution, a sample of individuals is taken from the population and their evolvability is estimated in the following way: First, the mutational neighborhood of that individual is approximated through creating 1,000 clones and mutating each clone independently once. Second, each clone is evaluated in the experimental domain and its behavior is quantified. Third, a regular grid is superimposed over the entire space of such quantified behaviors, and each of the mutated clones is mapped into the grid square that contains its behavior. The original individual’s evolvability is then the number of populated grid squares, i.e. a quantification of the propensity of the genotype’s offspring to realize diverse behaviors.

This measure is applied in this paper’s experiments to test whether evolvability is enhanced by algorithmic extinction, which is described in the next section.

2.3 Extinction in Evolutionary Computation

Because on their surface extinction events appear simply destructive, they are not often integrated into EAs. When extinction-like mechanisms are included in EAs, they often serve to *focus* search [9, 11, 14, 31] and remove seemingly unpromising classes of solutions [8, 28]. As a result, extinctions are often initiated when the search for higher performance stagnates [11, 28]. Overall, typical incarnations of algorithmic extinctions implicitly or explicitly protect the overall best-performing individual [8, 11, 14, 20].

In contrast, the implementation of extinction events in this paper does not inherently favor any particular behavior in the population over any other and is not triggered by stagnating performance. Instead, extinctions wipe out particular classes of behaviors irrespective of performance. The idea is to shift the focus from particular high-performing individuals to the overall evolvability of the population as a whole. Giving credence to this idea, Palmer and Feldman [25] demonstrated in an abstract model that indiscrim-

inate extinctions acting upon spatially-divided populations can lead to increasing evolvability, which is similar to the approach in this paper.

A more typical example of extinction-like strategies in EC is *burst mutation* [9, 11]. Burst mutation (which is similar to delta-coding [31]), is initiated when performance stagnates for a fixed number of generations, and extinguishes all but the highest performing solution. Evolution then focuses on the neighborhood in the search space centered on this solution. Weight mutations are applied from the Cauchy distribution, which result in mainly small tweaks, with occasional large changes. The overall idea is that sometimes greater *precision* is needed to solve some tasks, which is facilitated by corresponding greater *focus* on promising individuals.

Another relevant approach is random restarts [10], which can be seen as total extinction. The entire population is re-initialized if the task is not solved within a fixed number of evaluations. The hope is that a random restart may allow evolution to uncover a different (and more promising) attractor. The approach in this paper is somewhat similar to random restarts in that nearly all individuals are extinguished; however, a critical difference is that a small diverse set of individuals persists extinction.

Importantly, the effect of adding extinctions to an EA may depend on characteristics of the underlying EA. In particular, while most EAs converge as they seek to optimize towards a fixed objective, other EAs are instead driven to explore many divergent possibilities simultaneously. The next section reviews such divergent EAs.

2.4 Divergent Evolutionary Algorithms

A signature characteristic of natural evolution is its tendency to accumulate diversity over time [27], expanding through niches and discovering a wide range of solutions to the problems of life. In this way, natural evolution is *divergent*, i.e. it is inherently driven to explore many divergent possibilities in parallel. In contrast, optimization algorithms are most often *convergent*, i.e. they are driven to seek the singular highest value of a provided objective function. Similarly, most EAs apply an abstraction of natural evolution as an optimizer, i.e. the concept of biological fitness is often abstracted as a static fitness function that the EA optimizes. Thus, in a striking contrast to their natural inspiration, EAs themselves are often convergent [15, 16].

However, while dominant in EC, abstracting evolution as an optimizer is only one out of many possible foundations for creating an EA. By focusing on different aspects of natural evolution, other abstractions may produce EAs with different characteristics, e.g. those that agree with natural evolution’s tendency towards divergence. For example, EAs that abstract nature as an accumulator of novelty [16, 19, 24] diverge as they continually uncover novel forms. Because these kinds of algorithms are not inherently driven towards optimization, they are less susceptible to deception (often outperforming objective-based search on such problems [16]), and are more likely to increase evolvability [18]. Interestingly, even when objective-based search is as efficient, representations provided by divergent search are often more evolvable [15].

3. APPROACH

Just as there are many ways to perform convergent search, there are also many ways to create an EA with a drive to-

wards diversity [5, 16, 19, 24]. Two representative methods, novelty search and behavioral grid search, will be used in this paper. They are each reviewed in this section, followed by a description of how extinction events are implemented.

3.1 Novelty Search

Novelty search is inspired by natural evolution’s drive towards novelty, and rewards novel behavior directly *instead* of progress towards a fixed objective [16]. Tracking novelty requires little change to any evolutionary algorithm aside from replacing the objective-based fitness function with a *novelty metric*. Such a metric measures how different an individual is from other individuals, thereby creating a constant pressure to produce something new. The key idea is that instead of rewarding performance on an objective, novelty search rewards diverging from prior behaviors. Therefore, novelty in behavior needs to be *measured*.

The novelty metric characterizes 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. Intuitively, if the average distance to a given point’s nearest neighbors is large then it is in a sparse area; if the average distance is small, it is in a dense region. The sparseness ρ at point x is given by

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

where μ_i is the i th-nearest neighbor of x with respect to the distance metric dist , which is a domain-dependent measure of behavioral difference between two individuals in the search space. Candidates from more sparse regions of the behavior space thus receive higher novelty scores.

With fixed probability an individual is entered into the permanent archive that characterizes the distribution of prior solutions in behavior space. The current generation plus the archive constitute 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. However, even without an explicit objective, novelty search is still driven by meaningful information; that is, behaving in a novel way often requires learning the structure of the domain.

Once objective-based fitness is replaced with novelty, the underlying EA operates as usual, selecting the most novel individuals to reproduce. Over generations, the population spreads out across the space of possible behaviors.

Another divergent search algorithm used in this paper is behavioral grid search, described next.

3.2 Behavioral Grid Search

Behavioral grid search represents grid-based divergent search algorithms such as those described in Lehman and Stanley [19] and Cully et al. [5], and is tested in the experiments to demonstrate that the results are not specific to one particular divergent search algorithm.

Similar to novelty search, behavioral grid search is inspired by natural evolution’s tendency to explore many

niches in parallel. One simple way to facilitate diversity is to divide search resources explicitly across a variety of behavioral niches. Instead of encouraging behavioral diversity directly as in novelty search, a regular grid can be superimposed over the space of behaviors, wherein each grid square acts as a discrete niche. That is, each new individual is mapped to the grid square that contains the behavior it demonstrates. If each grid square only supports a limited capacity of organisms then the search as a whole will not converge, but will be driven to expand through the space of niches [19]. Note that this algorithm is in effect a steady-state version of the limited capacity niche model explored by Lehman and Stanley [19]; it is also similar to the MAP-ELITES algorithm of Cully et al. [5, 23].

While each such niche could impose competitive selection pressure among organisms with similar behavior (as in MAP-ELITES [5, 23] or novelty search with local competition [17]), for simplicity both divergent search algorithms in this paper are driven *only* by diversity. In particular, behavioral grid search applies steady-state replacement that is biased towards less-populated niches. Selection chooses an individual to reproduce asexually by picking one at random from a populated niche that is also chosen at random. To replace this new offspring, a randomly-chosen individual is removed from the most populous niche. Thus individuals in less-populated niches are more likely to be chosen to reproduce and will not be replaced. The overall effect is that the EA is driven to colonize unoccupied niches through evolving diverse behaviors, resulting in divergent search.

3.3 Extinction Events

The approach in this paper is to augment EAs with extinction events. The idea is that mass extinctions may indirectly select for the ability to diversify quickly. That is, if extinctions vacate a large proportion of occupied niches stochastically, then lineages that radiate through many niches are more likely to survive repeated such extinctions.

Thus, while there are many ways to implement extinction events [8, 14, 20], the one in this paper is motivated by extinction events as probabilistic filters biased towards evolvable lineages. In particular, the approach is at regular intervals to extinct a large proportion of evolved behaviors, without any consideration for their performance. That is, even the current champion is not immune from extinction.

From this general idea, the concrete implementation of extinctions is separately fitted to the three search methods applied in the experiments: behavioral grid search, novelty search, and traditional objective-based search. The idea is to implement extinctions to respect the distinct features of each search method.

Behavioral grid search defines explicit discrete niches. Thus in that model, an extinction event acts to vacate all but a fixed number of randomly-chosen niches; one individual from each selected niche is randomly-chosen to survive. Because each niche supports a different set of similar phenotypic behaviors, evolvable lineages (with higher phenotypic variability) that colonize many such niches are more likely to persist. Individuals from these few surviving niches are amplified as they repopulate vacated niches.

In contrast, there are no discrete niches in the novelty search and objective-based fitness EAs. To establish a similar effect, a greedy algorithm chooses a few surviving individuals spread across the space of behaviors. First, a set of

accumulated survivors is initialized with a random member of the population. Then, the set is augmented incrementally until it is of the desired size. To decide which individual to add to the set, the population is sorted by each individual’s minimum behavioral distance to an existing member of the set; an individual is then chosen at random from the top 20% of the sorted list. In this way, a small but diverse set of behaviors will survive each extinction. The stochasticity of the greedy algorithm ensures that extreme behaviors that are ranked highest will not always persist multiple extinctions; otherwise, because the underlying EA is a steady-state one, the same extreme behavior might always survive extinctions regardless of its evolvability or ability to diversify.

In all search methods, after the population is decimated, replacement is disabled temporarily until the population is replenished to its previous size. The experiments utilizing these techniques will be described next.

4. EXPERIMENTS

To explore whether EAs augmented with extinction events produce more evolvable populations, experiments were conducted in two representative ER domains: maze navigation [18] and biped locomotion [16]. These domains will be described first, followed by experimental setup and results.

4.1 Maze Navigation Domain

In the maze navigation domain, a simulated wheeled robot (figure 1) is embedded in a two-dimensional maze (figure 2). The objective for the robot is to traverse the maze and arrive at a fixed goal point. Thus, the fitness f of an individual for objective-based search is $f = b_f - d_g$, 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. For novelty search and behavioral grid search, evolution instead requires a characterization of behavior. Because ending location is a critical factor in navigating mazes, the behavior of a robot is defined as its location in the maze at the end of the evaluation [16, 22]. For behavioral grid search, each grid square within a regular grid superimposed over all ending locations acts a discrete niche. Individuals are mapped into the niche that contains the behavior they exhibit when evaluated.

The particular instantiation of the domain is the fragile hard maze of Lehman and Stanley [18]. It was chosen because it is one of the well-studied hard maze domains [16, 18, 19, 24] that offer significant capacity for increased evolvability relative to random initial populations [18, 19].

In the hard maze’s original conception [16], a robot is not penalized when it collides with a wall. However, in the fragile version, a robot’s evaluation is immediately terminated in such a case. For objective-based search, a colliding robot receives a fitness value f_{low} which is a minuscule value that is lower than otherwise possible. For novelty search, such a robot receives a minimal novelty score, n_{low} . For behavioral grid search, a colliding robot is mapped to a special niche that fills up quickly and renders colliding behavior inviable after the first few evaluations. The motivation for such penalties is to make the domain more *fragile*, i.e. more sensitive to mutation. That way the domain accentuates the challenge of discovering evolvable representations [18].

4.2 Biped Locomotion Domain

As the second domain, the biped domain of Lehman and Stanley [16] is adapted for these experiments (figure 3). In

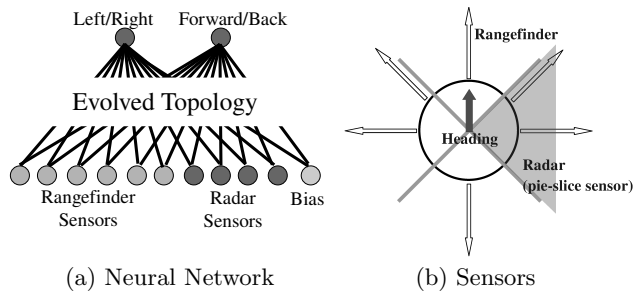


Figure 1: **A Maze-Navigating Robot.** The artificial neural network that controls the maze navigating robot is shown in (a). The layout of the sensors is shown in (b). Each arrow outside of the robot’s body in (b) is a rangefinder sensor that indicates the distance to the closest obstacle in that direction. The robot has four pie-slice sensors that act as a compass towards the goal, activating when the goal falls within the infinite projection of that pie-slice. The solid arrow indicates the robot’s heading.




Figure 2: **Maze Navigation Map.** In this map, the larger circle represents the starting position of the robot and the smaller circle represents the goal. To solve the task, the robot must navigate a circuitous path, which requires the evolution of non-trivial behavior. Further complicating the task, colliding with any wall immediately terminates the evaluation and renders the individual inviable. In this way, the domain presents a significant challenge for discovering evolvable representations.

this domain, the goal is to evolve a controller for a simulated biped robot that results in a stable walking gait. The problem is challenging because both balance and oscillation are needed. Like the maze domain it is also fragile: Most mutations to controllers are fatal and cause the robot to fall over quickly [18]. Thus biped locomotion provides a natural challenge domain for exploring how to increase evolvability and performance of EAs.

The domain works as follows (for more details see [16]). A biped robot in a physically realistic three-dimensional simulation is controlled by an ANN for a fixed duration (15 seconds). The evaluation is terminated if the robot falls or after the allocated time expires. The objective is for the robot to travel the greatest possible distance from the starting location.

The fitness of a biped controller for objective-based search is evaluated as the squared distance the robot walks before it falls. Its behavior for calculating novelty is derived from sampling its center of gravity each second it walks. The additional information provided by temporal sampling allows novelty search to differentiate two gaits that end up at the same location by different means.

For behavioral grid search, behavior is characterized simply by the robot’s center of gravity at the simulation’s completion. Such simplification is applied because the size of the grid of niches grows exponentially with increasing dimensionality. As in the maze domain, a regular two-dimensional




Figure 3: **Biped Robot.** In the biped locomotion domain, the ANN in (a) controls the biped robot that is visualized in (b). The robot has motors that apply forces to achieve the joint angles that are output by the ANN. In particular, there are motors for each of the six degrees of freedom: One in its left and right knees (LK and RK), and two in each hip (LH1, LH2, RH1, and RH2). Additionally, the robot’s ANN receives input from sensors in its feet that activate when they touch the ground. The challenge for the robot is to locomote as far as possible.

grid is superimposed over this behavior space, where each grid square defines a discrete niche shared by robots demonstrating behaviors encompassed by that square.

4.3 Experimental Setups

Across these two domains, three representative search algorithms are applied. Divergent search is represented by behavioral grid search and novelty search, while convergent search is represented by a traditional objective-based EA. While behavioral grid search is based on a geometric space of niches, novelty search and objective-based search share the same underlying steady-state EA and differ only in their underlying incentive scheme. All three search algorithms evolve ANNs represented by NEAT.

For each search method, four experimental setups are considered. In the *Control* setup there are no extinction events, i.e. the underlying EA operates as usual. In the *Extinction 100k*, *Extinction 200k* and *Extinction 400k* setups, extinction events occur every 100,000, 200,000, and 400,000 evaluations, respectively. The idea is to explore the effect of varying the inclusion and frequency of extinctions on the resulting evolvability of the different EAs.

In the maze domain, the population size was 250, while in the biped domain it was 500. Novelty search and objective-based search both used the same underlying steady-state EA with tournament selection and tournament size of five. Over both search algorithms and both domains, evolution ran for 3,000,000 evaluations and extinction events spare only 10 individuals; 40 independent runs were conducted for each combination of domain, setup, and algorithm.

For behavioral grid search, the resolution of the grid that was superimposed over the space of behaviors in the maze domain was 20×20 . The resolution was 40×40 in the biped domain over possible ending positions between $[-8.0, 8.0]$ meters for both planar coordinates. In both domains these same grids served to calculate evolvability, i.e. for all methods the evolvability of an individual was measured by the amount of unique grid squares exhibited by mutants of that individual. In particular, the evolvability of the population was estimated every 250,000 evaluations by applying the evolvability measure to 200 individuals chosen at random from the population.

4.4 Results

Figure 6 shows how the evolvability of final populations is distributed in all setups. Supporting this paper’s hypothesis, the main result is that extinction events result in increased evolvability only when combined with divergent search algorithms. A representative comparison of the temporal dynamics of evolvability increase between convergent and divergent search is shown in figure 4. Additionally, figure 5 visualizes how increasing evolvability allows evolution to rebound more quickly. (Corresponding figures for the maze domain and behavioral grid search are shown in the supplemental website <http://nn.cs.utexas.edu/downloads/papers/lehman.gecco2015-supplement.html>). Reflecting its label, in all experimenter observations across domains and setups the convergent objective-based search algorithm consistently converged to a single champion behavior.

Beyond evolvability, extinction events also affect the *performance* of search methods. Figure 7 highlights differences in evolutionary success, showing that extinction events lead to more effective adaptation when combined with divergent search. Independently of setup, objective-based search never solved the maze. This result is likely due to the deceptive nature of this domain [16]. Additionally, with objective-based search in the biped domain no significant differences in performance result from introducing extinction events. In contrast, extinctions increase performance of divergent search: In nearly all combinations of domains and divergent search algorithms, performance is highest with the Extinction 100k setup, suggesting that more frequent extinction events are the most effective. In the maze domain, differences between the divergent algorithms are not significant, while in the biped domain, each novelty search setup significantly outperforms its corresponding behavioral grid search setup (Mann-Whitney U-test; $p < 0.05$). The likely reason is that biped behaviors are differentiated better by the high-dimensional behavior characterization used in novelty search than with the simplified behavior characterization necessary for behavioral grid search.

5. DISCUSSION

This paper provides evidence for the counter-intuitive insight that repeated short-term destruction may enhance the long-term potential of an evolutionary process. Such robustness through upheaval is reminiscent of creative destruction in business and wildfires in ecosystems. Thus such events can serve as simple and effective enhancements to EAs.

However, echoing previous results with self-adaptation [18], the results also demonstrate that extinction events accelerate evolution only when combined with divergent search. In this way, powerful mechanisms of natural evolution may often be interdependent, and some such mechanisms may provide little benefit in isolation. In particular, a search reward scheme facilitating divergence (e.g. as in behavioral grid search or novelty search) may prove a critical ingredient enabling many mechanisms seen in natural evolution. Furthermore, differential benefit of this kind suggests that best practices learned for convergent search may not apply to divergent search. Thus, the conclusions forged in the context of the more dominant convergent search paradigm may need to be re-evaluated for divergent search.

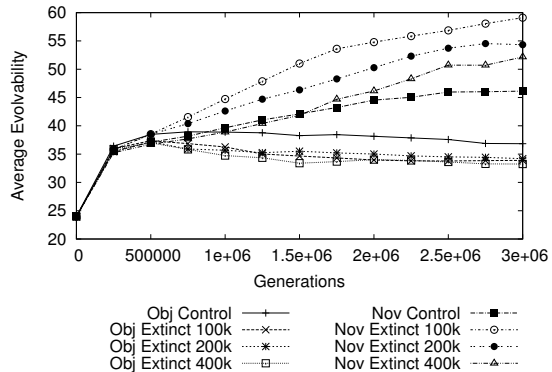


Figure 4: **The Change in Evolvability over Evolution.** The average (mean) evolvability of individuals in the population is shown for objective-based search (labeled Obj) and novelty search (labeled Nov) in the biped domain. Extinction setups combined with novelty search result in increased evolvability and an overall increasing trend. Extinction setups combined with objective-based search result in decreased evolvability and an overall stagnating trend. The conclusion is that there is a qualitative difference in how evolvability changes over time between the two search methods.

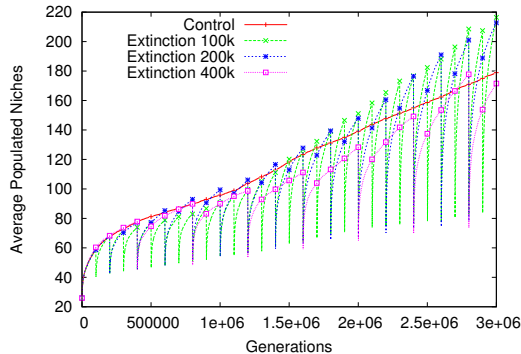


Figure 5: **Dynamics of Niche Occupation.** The average (mean) number of niches occupied over evolution is shown for behavioral grid search in the biped domain. The Control setup accumulates niches monotonically, whereas the Extinction setups are decimated at regular intervals. Because each extinction event spares only 10 niches, increasingly quick repopulation in the Extinction setups suggest that the representations are more evolvable in these setups than in the others. Note that each extinction always brings the model to exactly 10 niches, although this is not visible due to sampling error.

6. CONCLUSION

This paper forwards the hypothesis that extinction events can accelerate evolution in divergent search algorithms. The hypothesis is supported by results from two evolutionary robotics domains, in which extinction events benefit divergent search but not a more traditional convergent objective-based search. Thus extinction events may provide a simple and effective mechanism for improving the long-term performance of divergent EAs.

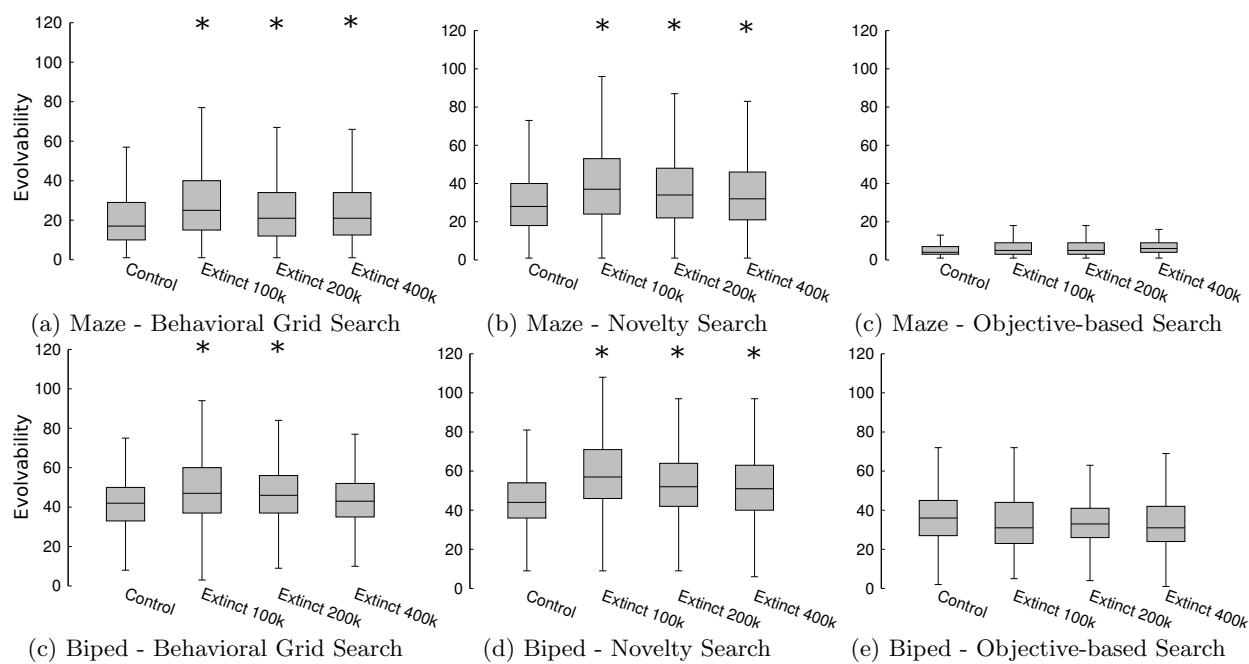


Figure 6: **Distribution of Evolvability in Final Populations.** The box plots show the evolvability of individuals in final populations for the four setups of the three search methods in the maze and biped domains. First, protective one-way ANOVA tests showed that there were significant differences between the extinction setups and the control with divergent search methods (i.e. behavioral grid search and novelty search; $p < 0.05$), but not with objective-based search. Second, pairwise Mann-Whitney U tests showed that with the divergent search algorithms, the average final evolvability in the extinction setups was significantly higher than in the control in 11 of the 12 pair-wise comparisons (those indicated with “*”; $p < 0.05$). The conclusion is that extinction events enhance evolvability only when paired with divergent search.

7. ACKNOWLEDGMENTS

This research was supported in part by NSF grants DBI-0939454, IIS-0915038, and SBE-0914796, and by NIH grant R01-GM105042.

References

- [1] T. Aaltonen et al., Measurement of the top quark mass with dilepton events selected using neuroevolution at CDF, *Physical Review Letters* (2009).
- [2] B. Allen and P. Faloutsos, Complex networks of simple neurons for bipedal locomotion, in: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2009).
- [3] J. Brookfield, Evolution: The evolvability enigma, *Current Biology*, 11(3):R106 – R108 (2001).
- [4] J. Clune, J.-B. Mouret, and H. Lipson, The evolutionary origins of modularity, *Proceedings of the Royal Society of London B: Biological Sciences*, 280(1755):20122863 (2013).
- [5] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret, Robots that can adapt like natural animals, *arXiv preprint arXiv:1407.3501* (2014).
- [6] R. Dawkins, The evolution of evolvability, *On Growth, Form and Computers*, 239–255 (2003).
- [7] M. Dichtel-Danjoy and M. Félix, Phenotypic neighborhood and micro-evolvability, *Trends in Genetics*, 20(5):268–276 (2004).
- [8] G. B. Fogel, G. W. Greenwood, and K. Chellapilla, Evolutionary computation with extinction: experiments and analysis, in: *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*, volume 2, 1415–1420, IEEE (2000).
- [9] F. Gomez and R. Miikkulainen, Incremental evolution of complex general behavior, *Adaptive Behavior*, 5:317–342 (1997).
- [10] F. J. Gomez and R. Miikkulainen, Solving non-markovian control tasks with neuroevolution, in: *IJCAI*, volume 99, 1356–1361 (1999).
- [11] F. J. Gomez and R. Miikkulainen, Active guidance for a finless rocket using neuroevolution, in: *Proceedings of the Genetic and Evolutionary Computation (GECCO 2003)*, 2084–2095, Springer (2003).
- [12] I. Harvey, *The Artificial Evolution of Adaptive Behavior*, Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex, Sussex (1993).
- [13] M. Kirschner and J. Gerhart, Evolvability, *Proceedings of the National Academy of Sciences of the United States of America*, 95(15):8420 (1998).
- [14] T. Krink and R. Thomsen, Self-organized criticality and mass extinction in evolutionary algorithms, in: *Evolutionary Computation, 2001. Proceedings of the 2001 Congress on*, volume 2, 1155–1161, IEEE (2001).
- [15] J. Lehman, S. Risi, and K. O. Stanley, On the benefits of divergent search for evolved representations, in: *Proceedings of the EvoNet 2012 Workshop at ALIFE XIII* (2012).
- [16] J. Lehman and K. O. Stanley, Abandoning objectives: Evolution through the search for novelty alone, *Evol. Comp.*, 19(2):189–223 (2011).




Figure 7: **Performance of Evolved Solutions.** The ability of evolution to generate well-adapted solutions is shown for the four setups of the three search methods in the maze and biped domains. In the maze navigation domain, a successful robot can navigate the full extent of the maze. In the biped domain, better solutions walk a longer distance. First, protective one-way ANOVA tests showed that there were significant differences between the extinction setups and the control with divergent search methods (i.e. behavioral grid search and novelty search; $p < 0.05$), but not with objective-based search. Second, pairwise Mann-Whitney U tests showed that with each divergent search method in each domain, the Extinction 100k setup significantly outperformed the control (those indicated with “*”; $p < 0.05$; the lines indicate standard error), and no extinction setups significantly underperformed the control. The conclusion is that extinction events enhance divergent search’s ability to uncover well-adapted solutions, but provide no comparable enhancement for convergent search.

- [17] J. Lehman and K. O. Stanley, Evolving a diversity of virtual creatures through novelty search and local competition, in: *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2011)*, ACM (2011).
- [18] J. Lehman and K. O. Stanley, Improving evolvability through novelty search and self-adaptation, in: *Evolutionary Computation (CEC), 2011 IEEE Congress on*, 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] J. Marin and R. V. Sole, Macroevolutionary algorithms: a new optimization method on fitness landscapes, *Evolutionary Computation, IEEE Transactions on*, 3(4):272–286 (1999).
- [21] A. P. Martin, Increasing genomic complexity by gene duplication and the origin of vertebrates, *The American Naturalist*, 154(2):111–128 (1999).
- [22] J.-B. Mouret, Novelty-based multiobjectivization, in: *Proceedings of the Workshop on Exploring New Horizons in Evolutionary Design of Robots, 2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (2009).
- [23] J.-B. Mouret and J. Clune, Illuminating search spaces by mapping elites., *arXiv preprint arXiv:1504.04909* (2015).
- [24] J.-B. Mouret and S. Doncieux, Encouraging behavioral diversity in evolutionary robotics: an empirical study, *Evolutionary Computation*, 20(1):91–133 (2012).
- [25] M. E. Palmer and M. W. Feldman, Spatial environmental variation can select for evolvability, *Evolution*, 65(8):2345–2356 (2011).
- [26] M. Pigliucci, Is evolvability evolvable?, *Nature Reviews Genetics*, 9(1):75–82 (2008).
- [27] R. Standish, Open-ended artificial evolution, *International Journal of Computational Intelligence and Applications*, 3(167) (2003).
- [28] K. O. Stanley and R. Miikkulainen, Evolving neural networks through augmenting topologies, *Evolutionary Computation*, 10:99–127 (2002).
- [29] K. O. Stanley and R. Miikkulainen, Competitive coevolution through evolutionary complexification, 21:63–100 (2004).
- [30] G. Wagner and L. Altenberg, Complex adaptations and the evolution of evolvability, *Evolution*, 50(3):967–976 (1996).
- [31] D. Whitley, K. Mathias, and P. Fitzhorn, Delta coding: An iterative search strategy for genetic algorithms, in: *ICGA*, volume 91, 77–84 (1991).
- [32] X. Yao, Evolving artificial neural networks, *Proceedings of the IEEE*, 87(9):1423–1447 (1999).