On the Benefits of Divergent Search for Evolved Representations

To appear in: *Proceedings of the EvoNet2012 Workshop at the Thirteenth International Conference on Artificial Life (ALIFE XIII).*

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

Sebastian Risi
Department of Electrical Engineering
and Computer Science
University of Central Florida
Orlando, Florida 32826–2363
Email: risi@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*—Evolved representations in evolutionary computation are often fragile, which can impede representation-dependent mechanisms such as self-adaptation. In contrast, evolved representations in nature are robust, evolvable, and creatively exploit available representational features. This paper provides evidence that this disparity may partially result from a key difference between natural evolution and most evolutionary algorithms: Natural evolution has no overarching objective. That is, nature tends to continually accumulate novel forms without any final goal, while most evolutionary algorithms eventually converge to a point in the search space that locally maximizes the fitness function. The problem is that individuals that maximize fitness do not *need* good representations because a representation's *future* potential is not reflected by its *current* fitness. In contrast, search methods without explicit objectives that are consequently divergent may implicitly reward lineages that continually diverge, thereby indirectly selecting for evolvable representations that are better able to diverge further. This paper reviews a range of past results that support such a hypothesis from a method called novelty search, which explicitly rewards novelty, i.e. behaviors that diverge from previously encountered behaviors. In many experiments, novelty search demonstrates significant representational advantages over traditional fitness-based search, such as evolving more compact solutions, uncovering more evolvable representations, and more fully exploiting representational features. The conclusion is that divergent evolutionary algorithms like novelty search may exert selection pressure towards higher quality representations than traditional convergent approaches to search.

***Index Terms*—Evolutionary computation, neural nets, heuristic methods**

I. Introduction

The representations of organisms evolved in nature are notable for their significant robustness and evolvability [3, 15]. In contrast, the representations uncovered by evolutionary algorithms are often fragile [11, 15]. The significant potential of representational mechanisms such as self-adaptation are often therefore unrealized in practice [2, 5, 13]. Thus an important question in evolutionary computation is what causes such a noticeable disparity. This question is particularly significant to the field of neuroevolution [10, 14, 17], i.e. evolving artificial neural networks (ANNs), where the domain itself may induce fragility [6] and poor evolvability can hinder efforts to evolve complex behaviors [6, 18].

One hypothesis is that the representational disparity between nature and EC results from a key difference between the dominant reward schemes in EC and natural evolution: The fitness-based search paradigm in EC often rewards progress towards a fixed objective, while natural evolution instead accumulates phenotypically diverse solutions to the problems of life with no overall final objective. In this way, evolution can be viewed as a *divergent* process, while the fitness-based abstraction of evolution common in evolutionary computation is often *convergent*.

This difference is important to representation because it can indirectly impact the kinds of representations that are rewarded. In particular, a divergent search like natural evolution may implicitly reward lineages better at diverging (e.g. genes that enable new species to emerge are less likely to go entirely extinct). The critical factor is that in a divergent search a representation that continually facilitates diverging from the past can distinguish itself over time from a less evolvable representation. Put another way, more evolvable representations correlate with continuing phenotypic divergence. As a result, divergent searches also may often incentivize exploiting representation-dependent features like mutational self-adaptation because such features can enhance the ability to diverge. That is, they can better align reproduction with representation to more consistently produce novelty. For example, self-adaptation in an ANN might limit mutations to connections critical to functionality and exaggerate mutations on connections that can effectively modulate behavior. The idea is that this kind of elegant self-adaptive outcome might evolve when using a divergent search for the very reason that it is more likely to create further divergence.

While it might seem like a similar argument could be made for convergent fitness-based search (it may reward lineages better at generating higher fitness), that argument is undermined by two fundamental problems. First, relatively high fitness in a population does not necessarily correlate with ability to further increase fitness, because fitness in evolutionary algorithms (with a static fitness function) is absolute. That is, unlike evolutionary novelty, a high-fitness individual has high fitness no matter when it is discovered. So the tendency to stay fit (a static property) does not reflect as much about underlying representation as the tendency to produce change. The second counterargument to fitness is that because fitness-based search tends to converge, by definition it will demonstrate a natural tendency to eliminate representational variation. Without such variation there cannot be any indirect selection on representation at all. Thus when converged to a local optimum a fitness-based search will likely become fixated on an arbitrary representation. Escaping from the local optimum is then predicated on modifying this arbitrary representation to increase fitness further, which may require ad-hoc, patchwork-like changes [16]. In other words, such convergence to a single representation and subsequent pressure to shoehorn that representation towards higher fitness may often oppose elegant or evolvable solutions, and may only be able to exploit representation-dependent features like self-adaptation to the extent that they can greedily increase or maintain fitness.

In fact, supporting such an idea and the hypothesis presented in this paper, evidence is accumulating that divergent searches in evolutionary computation may often benefit representation [1, 6, 8, 9, 12, 16]. To examine such evidence in more detail, this paper reviews past results from a representative example of divergent search called novelty search [7, 8] that explicitly rewards behavioral novelty (i.e. diverging from behaviors previously seen during search). Many of these experiments compare novelty search to a more traditional fitness-based search in either neuroevolution or genetic programming. The idea is to isolate the effect of changing the reward scheme from a convergent search for an objective to a divergent search for behavioral novelty. In a wide range of experiments such novelty search has proven beneficial to evolved representations. Thus instead of focusing as usual on the quality of the *solutions*, this paper focuses instead on the quality of the *representations*, even in cases when both approaches can find solutions.

II. Novelty Search

This section briefly reviews novelty search; for a more comprehensive introduction see Lehman and Stanley [8]. In contrast to most evolutionary algorithms, which tend to converge, novelty search is a *divergent* evolutionary technique. It is inspired by natural evolution's drive towards novelty, and directly rewards novel behavior *instead* of progress towards a fixed objective [7, 8].

The main idea in novelty search is to quantify a new individual's divergence from past behaviors. The greater the divergence from the past, the more promising an individual is considered by novelty search. To facilitate measuring such divergence, each individual (e.g. an ANN or evolved program tree) is mapped to a point in *behavior space* through a domain-specific characterization of behavior. A good behavior characterization should succinctly capture the dimensions of behavioral variation that are fundamental to a particular domain in the hope of reflecting an intuitive distance between evolved behaviors. In this way, the motivation is to align the algorithmic idea of novelty with our own intuitions on what the concept means.

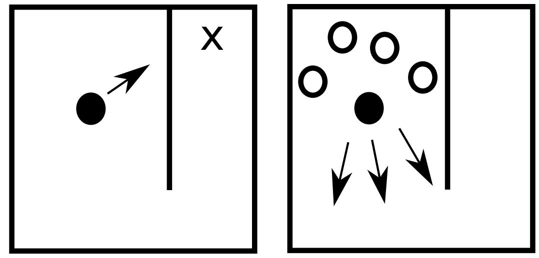

(a) Fitness-based Search (b) Novelty Search

Fig. 1: **Reward Scheme in Novelty Search.** The gradients of improvement (illustrated by arrows) from a particular behavior (the filled circle) instantiated by (a) convergent fitness-based search and (b) novelty search are shown for a simplified representation of a two-dimensional behavior space. The solid lines indicate behavioral constraints that mutation cannot cross and the objective of the search is represented by an "X". For novelty search, the open circles represent previously explored behaviors. Note that novelty search's gradients diverge (and can be explored by diverging evolutionary lineages) while fitness-based search in this scenario is driven to converge towards a local optimum in the direction of the objective.

After an individual is mapped into the behavior space, its novelty is measured as the sparseness of its neighborhood within that space. This sparseness is approximated in practice by measuring a new individual's average distance to its closest neighbors among the current population and an archive of individuals whose behaviors were judged highly novel when they were first encountered. Once objective-based fitness is replaced with novelty, the underlying evolutionary algorithm operates as normal, selecting the most novel individuals to reproduce. In effect, novelty search is driven to explore the behavior space. The gradient of search is simply towards what is *new*, with no other explicit objective. Over generations, the population spreads out across the space of possible behaviors. Figure 1 illustrates how this drive towards novelty differs from more traditional reward schemes in evolutionary computation.

Interestingly, although novelty search applies no direct pressure to accomplish any particular objective, it has often proven to benefit the search for the objective in deceptive domains [8, 9, 12]. However, the *performance* of novelty search is not the focus of this paper. Instead, the next section reviews empirical results that demonstrate novelty search's tendency to benefit evolved *representation*.

III. Empirical Impact of Novelty Search on Representation

To provide evidence that novelty search generally encourages better representations than a more conventional fitness-based search for the objective, this section reviews results in which novelty search discovers more elegant solutions [8, 9],

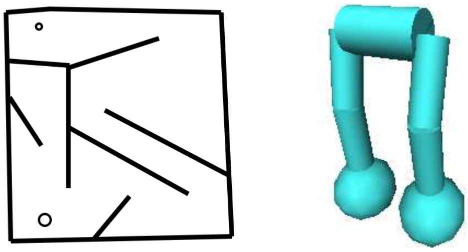

(a) Maze Domain (b) Biped Robot

Fig. 2: **Example Domains.** In the maze domain shown in (a), an ANN controls a maze-navigating robot that starts each evaluation at the location marked by the larger circle. The aim is to discover an ANN that can navigate the robot to the target location within the maze (marked by the smaller circle). Fitness-based search rewards minimizing the distance from the target location to the robot's location at the end of its evaluation. In contrast, novelty search incentivizes robots to end in novel locations within the maze. The goal in the biped domain is to evolve an ANN that can control the biped robot shown in (b) to walk as far as possible during the ANN's evaluation in a three-dimensional physically-realistic simulation. For fitness-based search, the reward function is the distance traveled by the robot before it falls, while novelty search instead incentivizes ANNs that generate walking trajectories *different* from those seen before during search.

more general solutions [1, 12], solutions that better exploit representation-dependent features [6], and more evolvable solutions [6].

First, in several experiments novelty search has evolved significantly more compact solution ANNs or programs than fitness-based search [8, 9], even when both methods solve the same problem. This observation is important because a smaller solution may more elegantly capture the regularities of the domain or be further modified more easily [16]. In particular, Lehman and Stanley [8] demonstrated that neuroevolution with novelty search evolved smaller solution ANNs than fitness-based search in two domains with different ANN models. Note that these first domains are shown in more detail in figure 2 as illustrative examples. In the first domain (figure 2a), a robot controlled by a plain ANN had to navigate a maze in a fixed time limit. Evolved continuous time recurrent neural networks controlled a biped robot (figure 2b) in the second domain with the goal of walking far. Similarly, Lehman and Stanley [9] applied genetic programming to maze navigation and the artificial ant domain [4] (where an artificial ant controlled by an evolved program must navigate a trail of food) and found that in all four experiments (two different mazes and two different ant trails) novelty search evolved smaller solution programs and was less susceptible in general to the significant problem in genetic programming of program bloat.

Additionally, some experiments have hinted that novelty search may lead to solutions that generalize better than those crafted by fitness-based search [1, 12]. Generalization is significant because a more general solution can be applied in contexts not explicitly encountered during evolution, avoiding the need to re-run evolution. In Risi et al. [12], plastic neural networks able to generalize were evolved faster by novelty search than by fitness-based search in a T-maze domain requiring ANNs controlling wheeled robots to learn from experience (similar to T-mazes used for testing the learning of live rats). Furthermore, genetic programs evolved by novelty search in Doucette [1] were more likely to generalize to solve random variations of artificial ant trails than those from fitness-based search.

Another aspect of representation is exploiting representation-dependent features such as self-adaptation (i.e. encoding reproductive parameters within a genome). Ideally, an evolutionary algorithm would fully exploit such parameters because they create greater *potential* for discovering a broad range of possible evolved behaviors. Lehman and Stanley [6] accordingly illustrated three neuroevolution domains where self-adaptation of probability and strength of mutations for ANN connections was exploited by novelty search but hindered fitness-based search (because it accelerated search's ability to converge to local optima [15]).

Finally, Lehman and Stanley [6] explicitly investigated the impact of novelty search on evolved ANNs' evolvability, i.e. the ability of an individual to further evolve [3, 15]. Evolvability is important because it may underlie our appreciation of natural organisms' significant ability to adapt to changing environments [3, 15]. Furthermore, investigating evolvability directly may most fundamentally test the hypothesis that selecting for divergence may encourage representations' ability to diverge. In Lehman and Stanley [6], when applied to maze navigation and biped walking, and compared to fitness-based evolution, novelty search uncovered more evolvable solutions. It also demonstrated higher evolvability on average across evolved populations.

Thus across a wide range of domains and in both neuroevolution and genetic programming, novelty search has been shown to benefit evolved representation, which supports the hypothesis that divergent search may be a key ingredient in the disparity in representational robustness between evolutionary computation and natural evolution.

IV. Discussion

The results reviewed in this paper suggest that applying convergent evolutionary algorithms is often detrimental to evolved representation, supporting the argument in Woolley and Stanley [16]. While this insight is consistent with the ubiquity of divergence and creativity in natural evolution, the lesson remains important because of the near-ubiquity in evolutionary computation of convergent evolution through static fitness functions.

Finally, the success of novelty search at exploiting representation-dependent features like self-adaptation suggests the need to reevaluate previous results in which the promise of a representation-dependent feature was not met by fitness-based search [2, 5, 13]. The true potential of such a feature may not be reflected by the ability of convergent search to exploit it. For experiments comparing or modifying representations in particular it may be important to apply some form of divergent search or to otherwise encourage representational quality. Otherwise, the conclusions reached could be deceptive and poorly reflect a representation's true potential.

V. Conclusions

This paper argues the hypothesis that divergent search, i.e. search that branches towards many continually-diverging goals, may generally encourage better representations than the dominant approach in evolutionary computation of fitness-based search towards a fixed target. This hypothesis was supported by reviewing previous experiments with novelty search, a divergent evolutionary algorithm that often demonstrates representational advantages over fitness-based search. The general conclusion is that exploratory divergent search processes like novelty search may often prove beneficial for representation.

Acknowledgment

This research was supported by DARPA and ARO through DARPA grant N11AP20003 (Computer Science Study Group Phase 3) and US Army Research Office grant Award No. W911NF-11-1-0489. This paper does not necessarily reflect the position or policy of the government, and no official endorsement should be inferred.

References

[1] J. Doucette, "Novelty-based fitness measures in genetic programming," Master of Science in Computer Science, Dalhouise University, 2010.

[2] M. Glickman and K. Sycara, "Reasons for premature convergence of self-adapting mutation rates," in *Evolutionary Computation, 2000. Proceedings of the 2000 Congress on*, vol. 1. IEEE, 2002, pp. 62–69.

[3] M. Kirschner and J. Gerhart, "Evolvability," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 95, no. 15, p. 8420, 1998.

[4] J. R. Koza, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge, MA: MIT Press, 1992.

[5] O. Kramer, "Evolutionary self-adaptation: a survey of operators and strategy parameters," *Evolutionary Intelligence*, vol. 3, pp. 51–65, 2010.

[6] J. Lehman and K. O. Stanley, "Improving evolvability through novelty search and self-adaptation," in *Evolutionary Computation (CEC), 2011 IEEE Congress on*. IEEE, 2011, pp. 2693–2700.

[7] ——, "Exploiting open-endedness to solve problems through the search for novelty," in *Proc. of the Eleventh Intl. Conf. on Artificial Life (ALIFE XI)*. Cambridge, MA: MIT Press, 2008.

[8] ——, "Abandoning objectives: Evolution through the search for novelty alone," *Evol. Comp.*, vol. 19, no. 2, pp. 189–223, 2011.

[9] ——, "Efficiently evolving programs through the search for novelty," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2010)*. ACM, 2010.

[10] S. Nolfi and D. Floreano, *Evolutionary Robotics*. Cambridge: MIT Press, 2000.

[11] J. Reisinger, K. O. Stanley, and R. Miikkulainen, "Towards an empirical measure of evolvability," in *Genetic and Evolutionary Computation Conference (GECCO2005) Workshop Program*. Washington, D.C.: ACM Press, 2005, pp. 257–264.

[12] S. Risi, C. Hughes, and K. Stanley, "Evolving plastic neural networks with novelty search," *Adaptive Behavior*, 2010.

[13] G. Rudolph, "Self-adaptive mutations may lead to premature convergence," *Evolutionary Computation, IEEE Transactions on*, vol. 5, no. 4, pp. 410–414, 2002.

[14] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.

[15] G. Wagner and L. Altenberg, "Complex adaptations and the evolution of evolvability," *Evolution*, vol. 50, no. 3, pp. 967–976, 1996.

[16] B. G. Woolley and K. O. Stanley, "On the deleterious effects of a priori objectives on evolution and representation," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2011)*. ACM, 2011.

[17] X. Yao, "Evolving artificial neural networks," *Proceedings of the IEEE*, vol. 87, no. 9, pp. 1423–1447, 1999.

[18] N. Zaera, D. Cliff *et al.*, "(Not) evolving collective behaviours in synthetic fish," in *In Proceedings of International Conference on the Simulation of Adaptive Behavior*, 1996.