

An Extended Study of Quality Diversity Algorithms

Justin K. Pugh, L. B. Soros, and Kenneth O. Stanley
Department of Computer Science
University of Central Florida
Orlando, FL 32816 USA
{jpugh, lsoros, kstanley}@cs.ucf.edu

ABSTRACT

In a departure from conventional optimization where the goal is to find the best possible solution, a new class of evolutionary algorithms instead search for *quality diversity* (QD) – a maximally diverse collection of individuals in which each member is as high-performing as possible. In QD, diversity of behaviors or phenotypes is defined by a *behavior characterization* (BC) that is typically *unaligned* with (i.e. orthogonal to) the notion of quality. As experiments in a difficult maze task reinforce, QD algorithms driven by such an unaligned BC are unable to discover the best solutions on sufficiently deceptive problems. This study comprehensively surveys known QD algorithms and introduces several novel variants thereof, including a method for successfully confronting deceptive QD landscapes: driving search with multiple BCs simultaneously.

Keywords

novelty search, non-objective search, quality diversity, behavioral diversity, neuroevolution

1. INTRODUCTION

Evolutionary computation (EC) traditionally focuses on abstracting from natural evolution the idea that “survival of the fittest” is a powerful mechanism for optimization. Recognizing that such optimization-oriented applications generally fail to capture the apparent diversifying force observed in nature, a new search paradigm has emerged in EC where the goal is not the single best individual but instead a wide variety of successful individuals. Often applied in evolutionary robotics domains, new *quality diversity* (QD) algorithms such as NSLC [1] and MAP-Elites [2] collect diversity with respect to some *behavior characterization* (BC) while simul-

This work was supported by the National Science Foundation under grant no. IIS-1421925. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO’16 Companion July 20-24, 2016, Denver, CO, USA

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4323-7/16/07.

DOI: <http://dx.doi.org/10.1145/2908961.2909000>




Figure 1: QD-Gauntlet. An egocentric agent must navigate from the start point (bottom) to the goal point (top).

taneously optimizing performance within local behavioral neighborhoods.

Pugh et al. [3] present an initial study comparing QD algorithms in a simulated maze domain called QD-Maze, revealing the importance of the BC’s *alignment* with the notion of quality and additionally suggesting that unaligned BCs may be sub-optimal for driving search.

The goal of this paper is to expand significantly on the scope of Pugh et al. [3] with a more difficult maze domain and a greater variety of QD algorithms, focusing specifically on confronting the challenge of unaligned QD.

2. EXPERIMENT

Agents are controlled by evolved neural networks as in Pugh et al. [3] and tasked with navigating a maze called QD-Gauntlet (Fig. 1) along one of the four possible solution paths (each of which is composed of four successive segments). QD-Gauntlet represents an unprecedented level of difficulty for testing such approaches, containing several dead-ends to deceive quality-seeking algorithms and imposing strict time constraints on agents.

To investigate whether the performance gap between BCs observed in the relatively easy QD-Maze domain from Pugh et al. [3] increases with problem difficulty, the best and worst BCs (EndpointBC and DirectionBC, respectively) from the QD-Maze study are compared on the much more challenging QD-Gauntlet. EndpointBC characterizes agent behavior by its (x, y) coordinate at the end of its trial and is strongly aligned with the notion of quality (proximity to the goal). The unaligned DirectionBC instead characterizes *how* the

agent drives and consists of the directions the agent faced most often during each of five time slices: north (0.125), east (0.375), south (0.625), or west (0.875).

This study compares each of the QD algorithms considered in Pugh et al. [3] along with several new variants introduced for the first time here. In *MAP-Elites + passive genetic diversity* (MEPGD), each bin in the MAP-Elites grid has two slots: the first managed by elitism and the second managed by random replacement (30% chance). In **NS-NS**, two independent novelty scores are combined in a multi-objective formulation: one for EndpointBC and one for DirectionBC. In **NS-NSLC**, a local competition objective is added to NS-NS where neighbors are determined by DirectionBC. In **ME-ME**, two independent MAP-Elites grids exist in parallel (one for each BC); parents are selected equally from each grid and offspring are mapped to both grids. Finally, **MENOV-MENOV** is the same as ME-ME except selection is proportional to novelty (each grid also maintains an NS-style archive). Each treatment (Table 1) is run 20 times, each for 1,000,000 evaluations.

No BC	DirectionBC	EndpointBC	Multi-BC
Fitness	NS _d	NS _e	NS _e NS _d
	NSLC _d	NSLC _e	NS _e NSLC _d
	ME _d	ME _e	ME _e ME _d
	MENOV _d	MENOV _e	MENOV _e MENOV _d
	MEPGD _d	MEPGD _e	

Table 1: Treatments compared on QD-Gauntlet. Columns differ by the BC that drives search.

Treatments are evaluated according to the *QD-score*¹ metric from Pugh et al. [3] and a more optimization-oriented metric called *total maze progress* that measures how close all four maze legs are to being solved (according to the furthest point discovered down each solution path). Importantly, in this study QD-score is always measured with respect to DirectionBC consistent with the observation that in the literature, QD is usually collected with an unaligned BC. Thus the EndpointBC-driven treatments here effectively test whether QD with respect to one BC can be achieved passively by driving search with another BC altogether.

3. RESULTS & DISCUSSION

Figure 2 depicts the performance of each treatment according to total maze progress. Unsurprisingly, Fitness falls victim to deception and struggles to solve the difficult QD-Gauntlet. However, QD algorithms driven by DirectionBC (green) fare even worse, indicating that pursuing unaligned diversity does not help overcome the problem of deception (in fact it may actually make matters worse). While it is well-known that novelty search driven by EndpointBC (NS_e, red) excels at maze-solving tasks, a surprising result is that NSLC_e may perform even better ($p = 0.09$, Student’s t-test). The success of NSLC_e here suggests that strictly local competition might represent a way to rectify the deception-thwarting power of novelty-seeking with the compass of the objective without reintroducing deception.

Figure 3 depicts the ability of each treatment to successfully collect QD in this domain. Consistent with their inability

¹Unlike the “proximity to the goal” heuristic that drives search, the quality measure in QD-score reflects the length of the shortest drivable path to reach the goal.




Figure 2: Total maze progress. Values are averaged over 20 runs and error bars represent standard error. A maximum of 100 points is possible per leg for a total possible score of 400 corresponding to solving all four legs.




Figure 3: Final QD-score. Values are averaged over 20 runs and error bars represent standard error.

to solve the QD-Gauntlet, DirectionBC-driven treatments (green) also fail to achieve QD; in fact, more QD is collected by driving search with EndpointBC (red) even though QD-score is always measured with respect to DirectionBC, suggesting that on hard problems, a BC’s ability to overcome deception is just as important to QD as its expression of the desired notion of diversity. Multi-BC QD algorithms (blue) represent a promising compromise, allowing search to simultaneously be driven by a BC suited to each purpose (e.g. one aligned and one unaligned). While NS_eNSLC_d and NSLC_e are essentially tied ($p = 0.163$) after the full 1,000,000 evaluations, a graph of their QD-score over time (not shown) displays fundamentally different trends: NS_eNSLC_d increases rapidly and reaches an early plateau, while NSLC_e only catches up in the long run (NS_eNSLC_d scores significantly better than NSLC_e until 224,000 evaluations).

This paper is also accompanied by a website where sample interactive behavior space visualizations are available for all treatments: <http://eplex.cs.ucf.edu/QD/GECCO16>

References

- [1] J. Lehman and K. O. Stanley. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, pages 211–218. ACM, 2011.
- [2] J.-B. Mouret and J. Clune. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909*, 2015.
- [3] J. K. Pugh, L. B. Soros, P. A. Szerlip, and K. O. Stanley. Confronting the challenge of quality diversity. In *Proceedings of the 17th Annual Conference on Genetic and Evolutionary Computation*, ACM, 2015.