Blending Notions of Diversity for MAP-Elites

Daniele Gravina
Institute of Digital Games
University of Malta
daniele.gravina@um.edu.mt

Antonios Liapis
Institute of Digital Games
University of Malta
antonios.liapis@um.edu.mt

Georgios N. Yannakakis
Institute of Digital Games
University of Malta
georgios.yannakakis@um.edu.mt

ABSTRACT

Quality-diversity algorithms focus on discovering multiple diverse and high-performing solutions. MAP-elites is such an algorithm, as it partitions the solution space into bins and searches for the best solution possible for each bin. In this paper, multi-behavior variants of MAP-Elites are tested where the MAP-Elites grid partitions the solution space based on a certain dimension, while selection is guided by measures of diversity on another dimension. Four divergent search algorithms are tested for this selection process, targeting novelty or surprise or their combination, and their performance on a soft robot evolution task is discussed.

CCS CONCEPTS

• **Computing methodologies → Search methodologies**; **Genetic algorithms**; *Evolutionary robotics*;

KEYWORDS

MAP-Elites; quality-diversity; surprise search; soft robots

ACM Reference Format:
Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2019. Blending Notions of Diversity for MAP-Elites. In *Genetic and Evolutionary Computation Conference Companion (GECCO '19 Companion), July 13–17, 2019, Prague, Czech Republic.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3319619.3321987

1 INTRODUCTION

Quality-diversity (QD) evolutionary algorithms must balance the search between high-performing solutions (exploitation) and highly diverse solutions (exploration) Of particular interest is the exploration/exploitation strategy of the MAP-Elites algorithm [8], which subdivides the feature space as a discrete *grid* and searches for the best performing individuals at each point of the chosen space. In the original implementation, parents are selected randomly among the best solutions (elites) at each point of the feature space, which guarantees a comprehensive exploration. However, this paper tests two possible ways of biasing the selection drive: (a) *rewarding* different measures of diversity or (b) using different *characterizations* of diversity. To achieve the first, we can reward the novelty of solutions [7], their surprise [2], or a combination of the two [3–5]. The second selection drive suggests that diversity between solutions can be measured in many ways, and rewarding diversity across multiple such dimensions can be beneficial as it further boosts exploration and can lead to more creative solutions to the same problem.

This paper tests the assumption that MAP-Elites can benefit through the use of both different *notions* of divergence (e.g. novelty or surprise) and dissimilar *characterizations* of diversity. Towards that end, we pressure the selection of individuals within MAP-Elites via novelty search, surprise search and combinations of the two. More importantly, we propose a *multi-behavioural characterization* (multi-BC) of the evolving solutions, using a different characterization for selection pressure than the characterization used in replacement (i.e. to discern elites). We investigate the impact of these MAP-Elites diversification mechanisms on the quality and diversity of solutions in a soft robot evolution testbed.

2 EXPERIMENT

Four new multi-behavior variants of the MAP-Elites algorithm are tested on a soft robot evolution task, as discussed below.

2.1 Testbed

Experiments in this paper use VoxCad [6] to simulate soft robots on a lattice of 5×5×5 voxels. Each voxel may have one of four materials, two active and two passive. A CPPN [9] decides which material should be placed in the lattice, using the 3D coordinate as input and the type of voxel as output. The *performance* characterization is the Euclidean distance covered by the evolved soft-robots from a fixed starting point until the end of the allocated simulation ticks [4]. Two different behaviour characterizations (BCs) are used: (a) *structural* BC as the percentage of filled voxels over the total lattice size and the percentage of blue voxels over the total filled voxels, (b) *path* BC as the tick-by-tick distance between two robots' trajectories [4].

2.2 Algorithms

This paper compares the original MAP-Elites (ME) algorithm to four multi-BC approaches. As in [8], in all MAP-Elites variants a batch of solutions (1024 individuals) is generated in every iteration by means of cross-over and mutation according to the principles of the NEAT algorithm [10]. All MAP-Elites variants use the structural BC to partition the space into a 2-dimensional map of 32×32 bins. The baseline ME selects candidate parents uniformly among elites stored in the map. MAP-Elites with Novelty Search (ME-NOV) selects parents based on the novelty score of [7] using the path BC as distance. MAP-Elites with Surprise Search (ME-SUR) selects parents based on the surprise score of [2] using the path BC as distance and by predicting paths of clusters based on a history of the last two generations (see [4]). MAP-Elites with Novelty-Surprise Search (ME-NSS) selects parents based on a linear combination of

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 '19 Companion, July 13–17, 2019, Prague, Czech Republic
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6748-6/19/07.
https://doi.org/10.1145/3319619.3321987

Table 1: Feature maps produced by the five methods, sampling a run of resolution of $5 \times 5 \times 5$. White bins do not have any robots, while colored bins denote respectively the fitness of the best individual, the number of offspring generated and the number of times a bin has been selected (from top to bottom).

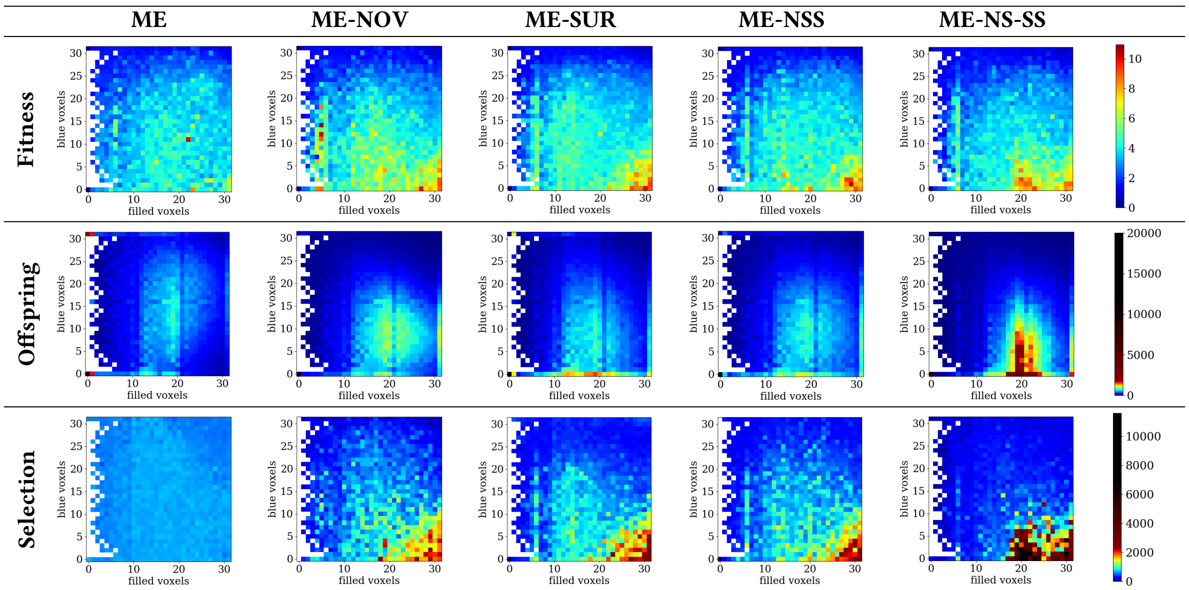

the surprise and novelty score, with 60% contribution of the novelty score and 40% of the surprise score. MAP-Elites with Novelty Search-Surprise Search (ME-NS-SS) finds Pareto-optimal individuals via NSGA-II [1] using novelty score and surprise score as objectives.

2.3 Results

Findings reported in this section are based on 10 independent evolutionary runs, and reported statistical significance is based on a two-tailed Mann-Whitney U-test at $p < 0.05$. Results show that all four multi-BC variants reached a significantly higher global performance [8] compared to ME within $300 \cdot 10^3$ evaluations. Specifically, ME-NS-SS had the highest global performance (12.80 ± 2.70) compared to ME (8.38 ± 0.70). However, the differences between ME and the multi-BC variants' QD score [8] were not significant.

In terms of diversity, all algorithms have the same coverage [8] of the feature space and thus explore along the structural BC in a similar way. However, the nearest-neighbor voxel-by-voxel structural distance shows that robots evolved via ME are significantly more diverse in terms of morphology than multi-BC approaches. On the other hand, the nearest-neighbor distance in path BC is significantly higher for each multi-BC approach compared to ME.

Table 1 shows three different heatmaps collected from a sample run of each approach, showing how the selection process and its outcomes were biased in each ME variant. Notably, offspring of ME-NS-SS are almost exclusively in the mid-bottom of the feature map. The bias towards such structures is reflected in the fitness, as more highly performing solutions are found in less filled structures.

3 CONCLUSION

This paper tested how MAP-Elites benefits from a selection pressure based on behaviour characterizations orthogonal to the features used to partition the space. Testing multi-BC variants on soft robot evolution, global performance and the diversity of robots' locomotion is increased. These findings provide more evidence to the power of a combined drive for novelty and surprise in QD algorithms.

4 ACKNOWLEDGEMENTS

This project has received funding from the European Union's Horizon 2020 programme under grant agreement No 787476.

REFERENCES

[1] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. A. M. T. Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. on Evolutionary Computation* 6, 2 (2002), 182–197.
[2] Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2016. Surprise Search: Beyond Objectives and Novelty. In *Proc. of the Genetic and Evolutionary Computation Conference.* ACM.
[3] Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2017. Coupling Novelty and Surprise for Evolutionary Divergence. In *Proc. of the Genetic and Evolutionary Computation Conference.* ACM.
[4] Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2018. Fusing Novelty and Surprise for Evolving Robot Morphologies. In *Proc. of the Genetic and Evolutionary Computation Conference.* ACM.
[5] Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2019. Quality Diversity Through Surprise. *IEEE Trans. on Evolutionary Computation* (2019).
[6] Jonathan Hiller and Hod Lipson. 2012. Dynamic simulation of soft heterogeneous objects. *arXiv preprint arXiv:1212.2845* (2012).
[7] Joel Lehman and Kenneth O Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19, 2 (2011).
[8] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
[9] Kenneth O Stanley. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines* 8, 2 (2007), 131–162.
[10] Kenneth O Stanley and Risto Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* 10, 2 (2002), 99–127.