Younger Is Better: A Simple and Efficient Selection Strategy for MAP-Elites

Alexandre Coninx
Sorbonne Université, CNRS, Institut des Systèmes Intelligents et de Robotique, ISIR
Paris, France
alexandre.coninx@sorbonne-universite.fr

Stephane Doncieux
Sorbonne Université, CNRS, Institut des Systèmes Intelligents et de Robotique, ISIR
Paris, France
stephane.doncieux@sorbonne-universite.fr

ABSTRACT

What makes a quality-diversity (QD) algorithm effective? This question is increasingly studied as such algorithms gain in popularity. Recent work on Novelty Search linked its efficiency to the evolution of a population under a dynamic selection pressure defined by the novelty metric, which allows the population to converge to highly evolvable individuals, on which spending the evolutionary budget allows to navigate the behavior space. MAP-Elites, another popular QD algorithm, does not use an explicit population, instead sampling the parent population directly from the behavioral archive at each generation. This sampling can be uniform, or biased according to criteria such as curiosity, which favor individuals that previously generated successful offsprings. In this article, we show that such improved selection schemes are efficient because they create a dynamic pseudo-population that mimics the desirable qualities of a Novelty Search population. We do this by proposing another simpler selection scheme, youth, that simply focuses the evolutionary budget on such a dynamic pseudo-population, and by showing it exhibits similar behavior and performance as curiosity without explicitly selecting for previously successful parents. This has important implications for the design of future QD algorithms taking the best of both archive-based and population-based methods.

CCS CONCEPTS

• **Computing methodologies → Search methodologies**; **Evolutionary robotics**; *Neural networks.*

KEYWORDS

Novelty search, evolutionary robotics, evolvability, behavior space

ACM Reference Format:

Alexandre Coninx and Stephane Doncieux. 2021. Younger Is Better: A Simple and Efficient Selection Strategy for MAP-Elites. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion), July 10–14, 2021, Lille, France.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3449726.3459493

How do diversity algorithms achieve evolvability? The answer varies, as those algorithms can be categorised into two main approaches depending on if their parent selection process rely on a population or on a much larger archive of solutions. Population-based approaches based on Novelty Search (NS) [4] follow the principles of most other evolutionary algorithms, with a fixed-size set of individuals engaged in a variation and selection loop. The selection process is driven by a novelty objective measuring how far a solution is from previously evaluated ones in a behaviour space. By contrast, in archive-based approaches like MAP-Elites (ME) [7], the selection process directly relies on the archive, with the parents individuals being directly drawn from it, and the children inserted into it if different enough from or better than existing ones.

For population-based algorithms, evolvability has been shown to be a by-product of the evolutionary process. In NS, an explored solution is not novel anymore once it has been explored once – and will be less and less novel when similar solutions will have been explored – the ability to move rapidly in the behaviour space is indirectly rewarded. Even in an elitist setup, older individuals do not tend to survive for long. NS thus creates an ever-moving population that favours individuals with a good evolvability [3].

In archive-based algorithms, it evolvability may also be a by-product of the algorithm [5]. Since the competition is local to a niche, individuals that invade new niches have no competitors and thus more chances to survive and generate offsprings, resulting in another indirect pressure towards evolvability. However, this model assumes that all individuals of the increasing archive generate new offsprings at each generation. In practice, at each generation, archive-based algorithms draw parents from the archive using a selection scheme (in the original ME, a random uniform selection [7]). The consequence is that the evaluation budget is spread over the whole archive whereas in population based algorithms it is concentrated on a limited set of points.

Evolvability is necessary both to explore new niches and to navigate a dynamic fitness landscape; we can therefore expect it to be promoted by both those pressures, which could be used in a synergistic way. Both NS and ME have a dynamic fitness landscape defined by their archive; what ME misses to take advantage of it for evolvability is a persistent population. Note that we do not need an explicit population distinct from the archive for this; any parent selection scheme that selects the parent population entirely or mostly from the parent and offsprings of the previous generation will exhibit a similar "population effect". A simple parent selection scheme that meets this requirement is to select the individuals that were added to the archive the most recently. We study the behavior of the ME algorithm with this *youth* selection scheme as well as three other schemes from the state of the art, and investigate 1) which selection schemes exhibit a "population effect" and 2) if this "population effect" is correlated with efficient exploration.

GECCO '21 Companion, July 10–14, 2021, Lille, France

ACM ISBN 978-1-4503-8351-6/21/07.
https://doi.org/10.1145/3449726.3459493

We studied four variants of the popular ME algorithm [7] differing only by the parent selection scheme used. We used the *random*, *pseudonovelty* and *curiosity* schemes, which correspond respectively to the *grid_random*, *grid_novelty* and *grid_curiosity* variants from [2], and our *youth* variant where the parents are the N individuals which were the most recently added to the archive. We also ran the NS algorithm from [4] as an exploration-only baseline.

The algorithms were evaluated in two simulated robotics tasks: *Fastsim-16x16*, a mobile robotics exploration task in a maze using the Fastsim 2D simulator [6], and *PyBulletAntDeterministic*, a locomotion task using a four-legged walker in an empty 3D world [1]. In both cases, the behavior space was the final position of the robot, and quality was based on minimizing motor commands.

Each variant of ME and the NS baseline were run 20 times on each task. In order to fairly compare NS to ME methods, we considered the coverage of the set of all the generated offsprings..

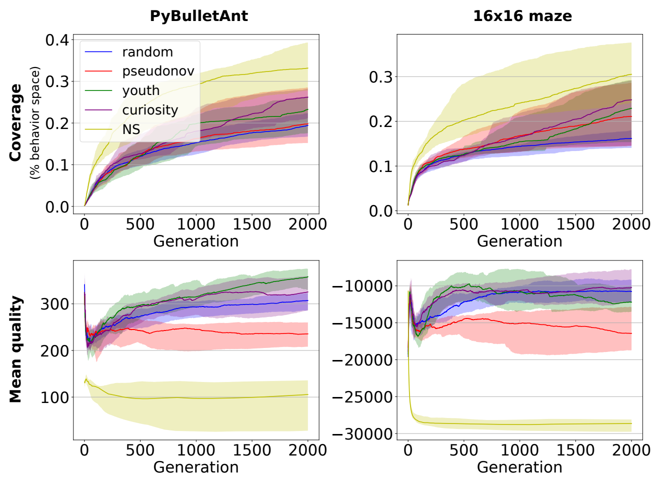

Figure 1: Coverage and mean quality in the archive for the four considered MAP-Elites variants and the NS baseline.

The coverage and quality results (fig. 1) show good performance for the *youth* variant and are consistent with the existing literature [2] for other algorithms. As expected, the pure exploration NS algorithm shows quicker coverage growth. Average quality measures are similar between ME variants, with the exception of the *pseudonovelty* variant which performs significantly worse.

Figure 2 shows the mean curiosity and age metrics on the selected parent set for each of the ME variants. The curiosity metric shows that the *youth* and *curiosity* variants have similar behaviors. This means the *youth* variant selects and generates high-curiosity individuals as well as the *curiosity* variant which explicitly selects individuals with the highest curiosity. The age metric again highlights the similarities between the *youth* and *curiosity* variants. Although the *youth* variant unsurprisingly selects the youngest individuals, the *curiosity* variant is close. This hints at the *curiosity* variant also exhibiting a "population effect", with the parent set being dynamically renewed with the recently added offsprings.

The *curiosity* variant is shown to behave similarly to the *youth* variant. This is because the curiosity score is set to zero in new individuals, and at each generation, the parents' curiosity increases or decreases depending on if their offspring is successful or not. Since the proportion of successful offsprings is often less than 10%, most parents quickly reach a negative curiosity value (see Fig. 2). Therefore, the successful offsprings overperform most older individuals on the curiosity metric with their score of zero, and the parent set is partially renewed with the successful offsprings, as in the *youth* variant. We can therefore at least partially explain the performance of the *curiosity* variant to this "population effect".

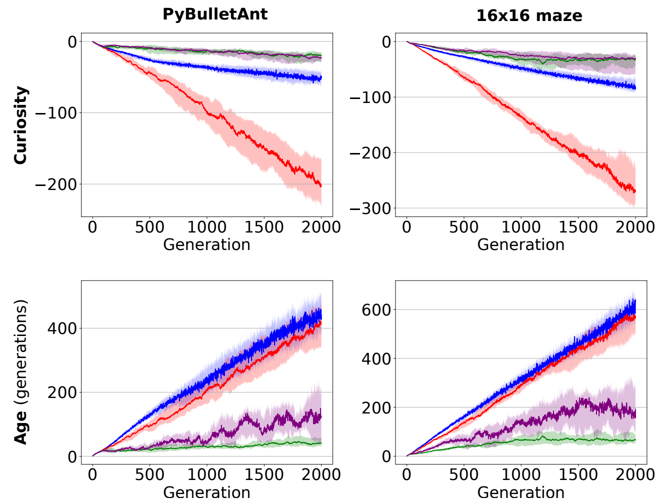

Figure 2: Mean value of the curiosity and age metrics on the parent set for the four considered ME variants.

We therefore propose the *youth* selection scheme as a simpler alternative to curiosity which performs similarly. Some more elaborate QD algorithms exploiting this "population effect" could be designed, combining it with some other criteria such as novelty or pseudonovelty, which are not sufficient to promote evolvability but have other interesting exploration properties.

ACKNOWLEDGMENTS

The authors would like to thank the ISIR Open Ended Learning discussion group and the HPCaVe Sorbonne Université HPC platform.

REFERENCES

[1] Erwin Coumans and Yunfei Bai. 2016–2019. PyBullet, a Python module for physics simulation for games, robotics and machine learning. http://pybullet.org.

[2] Antoine Cully and Yiannis Demiris. 2018. Quality and Diversity Optimization: A Unifying Modular Framework. *IEEE Transactions on Evolutionary Computation* 22, 2 (April 2018), 245–259.

[3] Stephane Doncieux, Giuseppe Paolo, Alban Laflaquière, and Alexandre Coninx. 2020. Novelty Search Makes Evolvability Inevitable. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference* (Cancún, Mexico) *(GECCO '20)*. Association for Computing Machinery, New York, NY, USA, 85–93.

[4] Joel Lehman and Kenneth O. Stanley. 2011. Abandoning Objectives: Evolution Through the Search for Novelty Alone. *Evolutionary Computation* 19, 2 (June 2011), 189–223.

[5] Joel Lehman and Kenneth O Stanley. 2013. Evolvability is inevitable: Increasing evolvability without the pressure to adapt. *PloS one* 8, 4 (2013), e62186.

[6] J. Mouret and S. Doncieux. 2012. Encouraging Behavioral Diversity in Evolutionary Robotics: An Empirical Study. *Evolutionary Computation* 20, 1 (March 2012), 91–133.

[7] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv:1504.04909 [cs, q-bio]* (April 2015). arXiv: 1504.04909.