

Comparing reliability of grid-based Quality-Diversity algorithms using artificial landscapes

Leo Cazenille
Ochanomizu University
Tokyo, Japan
leo.cazenille@gmail.com

ABSTRACT

Quality-Diversity (QD) algorithms are a recent type of optimisation methods that search for a collection of both diverse and high performing solutions. They can be used to effectively explore a target problem according to features defined by the user. However, the field of QD still does not possess extensive methodologies and reference benchmarks to compare these algorithms. We propose a simple benchmark to compare the reliability of QD algorithms by optimising the Rastrigin function, an artificial landscape function often used to test global optimisation methods.

CCS CONCEPTS

• Theory of computation → Evolutionary algorithms;

KEYWORDS

Quality-Diversity, Unconstrained Optimisation, Exploration, MAP-Elites

ACM Reference Format:

Leo Cazenille. 2019. Comparing reliability of grid-based Quality-Diversity algorithms using artificial landscapes. 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/10.1145/3319619.3321895>

1 INTRODUCTION

Search optimisation algorithms are popular methods to automatically explore a search space to find high-performing solutions. In such cases, the goal is traditionally to find the single best solution. However, in some problems it is useful for the search process to explore a range of both diverse and high-performing solutions. This approach is realised by a recent family of optimisation algorithms named Quality-Diversity (QD) algorithms or Illumination algorithms [3, 11], such as Map-Elites [9] and Novelty Search with Local Competition [7]. Algorithms like Map-Elites [9] are grid-based, and regroup the explored solutions in a grid of elites. This produces sets of high-performing solutions that vary according to features defined by the user, represented as axes of the grid.

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...\$15.00

<https://doi.org/10.1145/3319619.3321895>

These algorithms are particularly successful in evolutionary robotics problems [2, 6, 9], either by improving diversity to overcome deceptive search spaces [8], or by identifying and exploiting the generated repertoire of solutions [9].

QD algorithms are attractive methods to explore features landscapes of difficult and ill-defined problems, as the diversification of the explored solutions could help a traditional optimisation process to cope with highly multi-modal and deceptive target functions and prevent it from getting stuck in local-optima.

Here, we propose a method to compare grid-based QD algorithms by exploring artificial landscapes functions often used to compare global optimisation algorithms [5]. The goal is not just to find the global minimum as it is traditionally the case, but to exhaustively explore the landscape of this function. In this context, we compare the reliability [9] (*i.e.* how close the solutions found for each bin of the grid are to the oracle values) of two versions of the MAP-Elites algorithms on the Rastrigin function [12].

2 METHODS

2.1 Benchmark

We minimise the Rastrigin [12] function (Fig. 1). It is often used as a performance test problem for single-objective optimisation algorithms. On an n -dimensional domain [10], it is defined by:

$$f_n(\mathbf{x}) = An + \sum_{i=1}^n (x_i^2 - A \cos(2\pi x_i)) \quad (1)$$

where $A = 10$ and $x_i \in [-5.12, 5.12]$. Finding the minimum of this function is difficult due to its large search space and high multimodality. It is easy to increase the difficulty of this benchmark by increasing the dimensionality n of the target function.

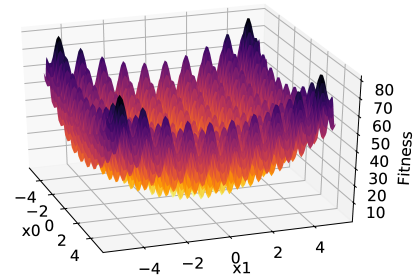


Figure 1: Surface plot of the Rastrigin function [12] with 2 dimensions. Global maxima are close to the corners and the single global minimum is at $x_i = 0$ where $f(x) = 0$.

We consider two features to illuminate, corresponding to the first two components of \mathbf{x} : x_0 and x_1 .

2.2 Reliability Metric

As defined in [9], the reliability of a QD algorithm corresponds to how close the solutions found for each bin of the grid are to the oracle values. Reliability can be measured locally (at the level of a single bin) or globally (covering the entire grid). We create a reference grid M of the oracle values by running the MAP-Elites [9] algorithm to illuminate the Rastrigin [12] function on a 2-dimensional domain.

The bin of M at position x, y is termed $M_{x,y}$. We set M_{\max} as the maximal quality values in M , and $n(M)$ as the number of filled bins in M . The local reliability $L(m_{x,y})$ of bin $m_{x,y}$ for grid m is:

$$L(m_{x,y}) = \begin{cases} 0, & \text{if either } m_{x,y} \text{ or } M_{x,y} \text{ is not filled.} \\ E(m_{x,y}), & \text{otherwise.} \end{cases} \quad (2)$$

$$E(m_{x,y}) = \max \left(\frac{M_{\max} - m_{x,y}}{M_{\max} - M_{x,y}}, 0 \right) \quad (3)$$

The global reliability $G(m)$ for grid m is defined as:

$$G(m) = \frac{1}{n(M)} \sum_{x,y} L(m_{x,y}) \quad (4)$$

All benchmarks were conducted using the QDpy [1] framework.

3 RESULTS

We use the proposed methodology to compare two versions of the Map-Elites [9] algorithm according to their global reliability. **ME1** uses polynomial bounded mutations, like the original Map-Elites algorithm and as described in [4] (mutation prob.: 0.5, eta: 10). **ME2** uses Gaussian mutations (mutation prob.: 0.5, mean: 0, stddev: 1.0). Both versions use 64 bins per feature (i.e. 4096 bins per grid). This number of bins is chosen arbitrarily to illustrate our methodology.

We define the **Reference** case as an oracle that describes the illuminated grid of solutions for the 2-dimension Rastrigin function. Other methods are then compared to the **Reference** case. It is possible to compare illuminations of f_k to illuminations of f_l with $k \geq l$ because the Rastrigin function [12] is separable.

Figure 2 compares the results obtained after 1000000 evaluations, for both the **ME1** and **ME2** methods used to illuminate the Rastrigin function with varying number of dimensions. Figures 2A and B show examples of the final grids for the reference case and for **ME1** applied on the 14-dimension Rastrigin function. Figure 2C presents the evolution of global reliability by evaluations, for all tested cases.

4 CONCLUSION

We presented a simple method to compare the reliability of grid-based QD algorithms with artificial landscapes functions. Namely, we compared two version of the Map-Elites [9] algorithm on the Rastrigin function [12] often used to compare global optimisers [5].

We plan to extend this study to other artificial landscapes functions for unconstrained or constrained optimisation. Our methodology could be extended to test algorithms with grids composed of several different numbers of bins.

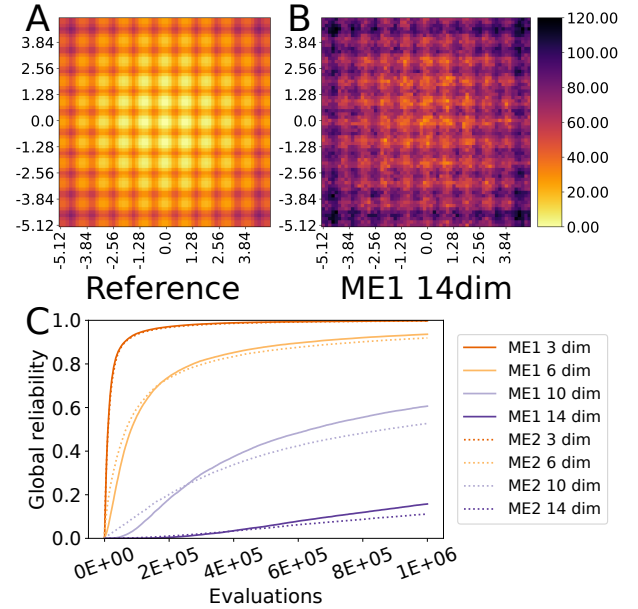


Figure 2: A, B: Quality of each bin of the final grids. C: Global reliability of the ME1 and ME2 algorithms on the Rastrigin function with $N = 3, 6, 10, 14$ dimensions compared to the reference case of Rastrigin with 2 dimensions.

REFERENCES

- [1] L. Cazenille. 2018. QDPY: A Python framework for Quality-Diversity. <https://gitlab.com/leo.cazenille/qdpy>. (2018).
- [2] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503.
- [3] Antoine Cully and Yiannis Demiris. 2018. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation* 22, 2 (2018), 245–259.
- [4] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and Tamaszt Meyarivan. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation* 6, 2 (2002), 182–197.
- [5] Johannes M Dieterich and Bernd Hartke. 2012. Empirical review of standard benchmark functions using evolutionary global optimization. *arXiv preprint arXiv:1207.4318* (2012).
- [6] Miguel Duarte, Jorge Gomes, Sancho Moura Oliveira, and Anders Lyhne Christensen. 2018. Evolution of repertoire-based control for robots with complex locomotor systems. *IEEE Transactions on Evolutionary Computation* 22, 2 (2018), 314–328.
- [7] Joel Lehman and Kenneth O Stanley. 2011. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the 13th annual conference on Genetic and evolutionary computation*. ACM, 211–218.
- [8] Joel Lehman, Kenneth O Stanley, and Risto Miikkulainen. 2013. Effective diversity maintenance in deceptive domains. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. ACM, 215–222.
- [9] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [10] Heinz Mühlenbein, M Schomisch, and Joachim Born. 1991. The parallel genetic algorithm as function optimizer. *Parallel computing* 17, 6-7 (1991), 619–632.
- [11] Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. 2016. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI* 3 (2016), 40.
- [12] Leonard Andrejevič Rastrigin. 1974. Systems of extremal control. *Nauka* (1974).