

Comparing multimodal optimization and illumination

Vassilis Vassiliades, Konstantinos Chatzilygeroudis, Jean-Baptiste Mouret

► **To cite this version:**

Vassilis Vassiliades, Konstantinos Chatzilygeroudis, Jean-Baptiste Mouret. Comparing multimodal optimization and illumination. Genetic and Evolutionary Computation Conference (GECCO 2017), 2017, Berlin, Germany. <hal-01518802>

HAL Id: hal-01518802

<https://hal.inria.fr/hal-01518802>

Submitted on 5 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Comparing multimodal optimization and illumination

Vassilis Vassiliades
Inria Nancy - Grand-Est
CNRS UMR 7503
Université de Lorraine
vassilis.vassiliades@inria.fr

Konstantinos Chatzilygeroudis
Inria Nancy - Grand-Est
CNRS UMR 7503
Université de Lorraine
konstantinos.chatzilygeroudis@inria.fr

Jean-Baptiste Mouret
Inria Nancy - Grand-Est
CNRS UMR 7503
Université de Lorraine
jean-baptiste.mouret@inria.fr

ABSTRACT

Illumination algorithms are a recent addition to the evolutionary computation toolbox that allows the generation of many diverse and high-performing solutions in a single run. Nevertheless, traditional multimodal optimization algorithms also search for diverse and high-performing solutions: could some multimodal optimization algorithms be better at illumination than illumination algorithms? In this study, we compare two illumination algorithms (Novelty Search with Local Competition (NSLC), MAP-Elites) with two multimodal optimization ones (Clearing, Restricted Tournament Selection) in a maze navigation task. The results show that Clearing can have comparable performance to MAP-Elites and NSLC.

KEYWORDS

illumination algorithms, multimodal optimization, MAP-Elites, quality diversity, behavioral diversity, novelty search

ACM Reference format:

Vassilis Vassiliades, Konstantinos Chatzilygeroudis, and Jean-Baptiste Mouret. 2017. Comparing multimodal optimization and illumination. In *Proceedings of The Genetic and Evolutionary Computation Conference, Berlin, Germany, July 2017 (GECCO'17)*, 2 pages.
DOI: 10.1145/nmnnnnn.nnnnnnn

1 INTRODUCTION

Illumination [8] or *quality diversity* [10] (QD) algorithms refer to a new type of evolutionary algorithms (EAs) capable of returning a large set of solutions that are as diverse and as high-performing as possible. These algorithms originated in the field of evolutionary robotics with the introduction of the novelty search algorithm (NS) [4] which suggests to look for individuals that are *behaviorally* different from previously encountered ones. That is, in NS there is an explicit distinction between the genotype space (in which the EA directly operates, e.g., a space of bit strings), the phenotype space (e.g., the space of neural networks derived from the genotype space) and the behavior space (e.g., the possible behaviors of individuals over their lifetimes, such as the end-locations of robots controlled by the neural networks of the phenotype space).

NS continually explores the behavior space, without considering the task performance. However, in many cases we are often interested in having pressure for performance (e.g., finding the

fastest controller that reaches a certain location and not just any controller). NS with Local Competition (NSLC) [5] addresses this issue using a multiobjective approach: it ranks individuals according to their novelty (as in NS) and their local performance (i.e., how many from the k closest neighbors the individual outperforms). The Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) [1, 8] algorithm offers a different solution: it discretizes the behavior space into a number of niches, storing in every niche only the *elite* individual throughout the evolutionary simulation.

Niching techniques have traditionally been used in EAs with the purpose of multimodal optimization (e.g., see [2, 3, 6, 9]), i.e., for discovering the multiple optima of the underlying genotype or phenotype space. In contrast to such approaches, the primary goal of illumination algorithms is *not* optimization but diversity [10]. Other differences between the way the two approaches are typically used are the following: (1) the number of solutions returned with multimodal optimization is in the order of tens or hundreds (e.g., see [3]), whereas with illumination algorithms it is in the order of hundreds or thousands [1, 10]; (2) illumination is performed in behavior space [1, 10], whereas multimodal optimization is performed either in genotype or phenotype space [2].

It is currently unknown how multimodal optimization algorithms behave when they are set to return as many solutions as illumination algorithms and to operate in behavior space (though there are works that use speciation in behavior space [12]). This is what we investigate in this short paper.

Multimodal optimization has a long history and the purpose of this study is not to provide a comprehensive evaluation of the different algorithms that exist. Instead, we select two representative algorithms that are simple to implement and can be applied in the behavior space: the “clearing” method [9] and “restricted tournament selection” (RTS) [3]. Clearing requires from the user to provide a clearing radius which defines the niche of an elite individual, whereas RTS restricts competition among a user-defined number of randomly selected individuals from the population.

2 EXPERIMENTAL SETUP

We compare 2 illumination (NSLC, MAP-Elites) and 2 multimodal optimization algorithms (RTS, Clearing) in a maze navigation task (Fig. 1 upper left) in which a simulated mobile robot (diameter: 20 units) is controlled by an artificial neural network, whose architecture and parameters are evolved [7]. The robot starts from the bottom of a maze (size: 1000×1000 sq. units) and needs to reach the goal point at the center. Thanks to the openings, this maze permits 16 families of trajectories towards the center (thus, at least 16 behaviorally distinct optima). The fitness function is the smallest Euclidean distance between the center of the robot and the goal

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'17, Berlin, Germany

© 2017 Copyright held by the owner/author(s). 978-x-xxxx-xxxx-x/YY/MM...\$15.00
DOI: 10.1145/nmnnnnn.nnnnnnn




Figure 1: Best distance to the goal (center) for each algorithm in the environment used during evolution (leftmost) and all 16 evaluation environments each of which permits a single path to the goal. The box plots show the median and the interquartile range over 30 solutions, apart from the last column which is calculated from the medians over all 17 environments. Overall, Clearing is comparable to MAP-Elites, which in turn is slightly better than NSLC. RTS has the worst overall performance.

over the robot’s lifetime [7], which is set to 3000 simulation steps. The behavior descriptor of each individual (i.e., a point in behavior space) is the end location (2D) of the robot [4, 7].

We evaluate the quality of the solutions produced by the illumination and multimodal optimization algorithms by assessing the performance of each individual found in the final archive/population in environments that are modified versions of the one used during evolution: if the archive is made of diverse and high-performing individuals, then it should contain individuals with all the kinds of trajectory, including some that work in the modified environments; in the extreme opposite, if all the individuals of an archive have the same behaviors, none of them will work in the modified environments. In our scenario, we assess whether the resulting sets contain controllers that find all 16 paths to the goal by evaluating all the solutions produced by each algorithm in 16 different environments (Fig. 1, top), each of which permitting only one path to the goal.

In addition, we measure the QD-score for each algorithm [10], which is calculated by mapping an archive’s behavior descriptors to a MAP-Elites grid, keeping the best performing one in a cell, and summing the fitness scores from all cells.

3 RESULTS AND CONCLUSION

We use 30 independent evolutionary runs of 200k evaluations. MAP-Elites ($71^2 = 5041$ cells), NSLC (max archive size = 5041) and Clearing (pop. size = 5041) return solutions with a median fitness of less than 10 units in all evaluation environments. RTS (pop. size = 5041) has the worst overall performance with a median distance of more than 10 units in 5 environments and more than 100 units in 2 environments. This indicates that RTS might have not found all optima, or it might have found them and eventually lost them. Interestingly, Clearing has comparable performance to MAP-Elites and NSLC and better performance in 2 environments.

The QD-scores for a typical archive/population of all algorithms, calculated in the *initial* environment using a 32×32 grid, are the following (lower is better): MAP-Elites: 129493.8; NSLC: 79396.1; RTS: 14504.6; Clearing: 1152.4. This shows that Clearing has a better QD-score, followed by RTS, then NSLC and finally MAP-Elites. However, this does not agree with the results of Fig. 1, according to which RTS archives are substantially less diverse than MAP-Elites archives, and illustrates that the QD-score does not capture *everything* about behavioral diversity: Fig.1 shows that in all the environments, MAP-Elites found at least a single, high-performing individual, whereas the QD-score shows that, in the

initial environment, on average, the fitness of the individual found by MAP-Elites for each bin is lower than the one found by RTS.

The similarity in performance between MAP-Elites and Clearing could be explained by the fact that they use a fixed niche size, despite their different niche shapes (MAP-Elites: rectangular; Clearing: spherical) or whether they operate in a bounded (MAP-Elites) or unbounded (Clearing) space. Clearing has better performance than RTS which is the opposite of the findings in multimodal optimization [11].

Overall, this study shows that when *some* multimodal optimization algorithms (e.g., Clearing, but not RTS) are provided with (1) large population sizes and (2) distance metrics that operate in behavior space, they can be as effective as illumination algorithms. More investigation is warranted to understand the links between the two approaches.

ACKNOWLEDGMENTS

This work was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Project: ResiBots, grant agreement No 637972).

REFERENCES

- [1] A. Cully, J. Clune, D. Tarapore, and J.-B. Mouret. 2015. Robots that can adapt like animals. *Nature* 521, 7553 (2015), 503–507.
- [2] S. Das, S. Maity, B.-Y. Qu, and P. N. Suganthan. 2011. Real-parameter evolutionary multimodal optimization – a survey of the state-of-the-art. *Swarm and Evolutionary Computation* 1 (2011), 71–88.
- [3] G. R. Harik. 1995. Finding Multimodal Solutions Using Restricted Tournament Selection. In *Proc. of the 6th International Conf. on Genetic Algorithms*. Morgan Kaufmann, San Francisco, CA, 24–31.
- [4] J. Lehman and K. O. Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evol. Comp.* 19 (2011), 189–223.
- [5] J. Lehman and K. O. Stanley. 2011. Evolving a Diversity of Virtual Creatures Through Novelty Search and Local Competition. In *GECCO ’11*. ACM, 211–218.
- [6] S. Mahfoud. 1995. *Niching methods for genetic algorithms*. Ph.D. Dissertation. University of Illinois at Urbana-Champaign, Urbana, IL.
- [7] J.-B. Mouret. 2011. Novelty-based multiobjective optimization. In *New Horizons in Evolutionary Robotics*. Springer, 139–154.
- [8] J.-B. Mouret and J. Clune. 2015. Illuminating search spaces by mapping elites. *arXiv preprint arXiv:1504.04909* (2015).
- [9] A. Pétrowski. 1996. A clearing procedure as a niching method for genetic algorithms. In *Proc. of IEEE Intern. Conf. on Evolut. Comp.* IEEE, 798–803.
- [10] J. K. Pugh, L. B. Soros, and K. O. Stanley. 2016. Quality Diversity: A New Frontier for Evolutionary Computation. *Frontiers in Robotics and AI* 3 (2016), 40.
- [11] G. Singh and K. Deb. 2006. Comparison of Multi-modal Optimization Algorithms Based on Evolutionary Algorithms. In *GECCO ’06*. ACM, 1305–1312.
- [12] L. Trujillo, G. Olague, E. Lutton, F. D. De Vega, L. Dozal, and E. Clemente. 2011. Speciation in behavioral space for evolutionary robotics. *J. Intell. Robot. Syst.* 64 (2011), 323–351.