An analysis of ϵ-lexicase selection for large-scale many-objective optimization

William La Cava
University of Pennsylvania
Philadelphia, PA
lacava@upenn.edu

Jason H. Moore
University of Pennsylvania
Philadelphia, PA
jhmoore@upenn.edu

ABSTRACT

In this paper we adapt ϵ-lexicase selection, a parent selection strategy designed for genetic programming, to solve many-objective optimization problems. ϵ-lexicase selection has been shown to perform well in regression due to its use of full program semantics for conducting selection. A recent theoretical analysis showed that this selection strategy preserves individuals located near the boundaries of the Pareto front in semantic space. We hypothesize that this strategy of biasing search to extreme positions in objective space may be beneficial for many-objective optimization as the number of objectives increases. Here, we replace program semantics with objective fitness to define ϵ-lexicase selection for many-objective optimization. We then compare this method to multi-objective optimization methods from literature on problems ranging from 3 to 100 objectives. We find that ϵ-lexicase selection outperforms state-of-the-art optimization algorithms in terms of convergence to the Pareto front, spread of solutions, and CPU time for problems with more than 3 objectives.

CCS CONCEPTS

• **Theory of computation → Bio-inspired optimization**;

KEYWORDS

many-objective optimization, selection

ACM Reference Format:
William La Cava and Jason H. Moore. 2018. An analysis of ϵ-lexicase selection for large-scale many-objective optimization. In *GECCO '18 Companion: Genetic and Evolutionary Computation Conference Companion, July 15–19, 2018, Kyoto, Japan.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3205651.3205656

1 SUMMARY

The multi-objective optimization (MO) community is increasingly interested in algorithms that can scale to large numbers of objectives. Dealing with large numbers of objectives affects the search process [4, 5] and as a result, different types of algorithms perform well [2]. As research has progressed, studies have analyzed the ability of evolutionary multi-objective optimization (EMO) algorithms to find or appapproximate Pareto-optimal solution sets for problems with up to 6 [12], 50 [1], and more recently, 100 objectives [9].

At the same time that EMO research has moved to larger sets of objectives, the genetic algorithm (GA) and genetic programming (GP) communities have shown strong interest in the so-called "multi-objectivization" of single-objective problems [6]. In GP, the idea of restructuring search drivers around program semantics, defined as the outputs or behavior of a GP program, has gained traction [10]. Rather than aggregating performance across training instances, semantic methods can re-define the problem as a set of smaller objectives for driving search. One such method is lexicase selection [11], which uses program semantics to filter the population via randomized orders of training cases at each selection event. By doing so it is able to adapt selection pressure to subsets of training cases that are harder to solve.

A variant of lexicase selection [11] known as ϵ-lexicase selection was introduced to apply lexicase selection to continuous error spaces for symbolic regression [8]. A recent theoretical analysis considered ϵ-lexicase selection through a multi-objective lens [7]. It showed that ϵ-lexicase selects individuals located near boundaries of the Pareto set defined by the population's error vectors. In this sense, ϵ-lexicase selection demonstrated an instance of a multi-objective treatment of regression with promising results.

In this work[1], we evaluate the performance of ϵ-lexicase selection as a many-objective optimization algorithm, shown in Algorithm 1. Our experimental study consists of a comparison of ϵ-lexicase selection to two EMO algorithms: NSGA-II and HypE. We compared these methods on the scalable DTLZ problems [3] using 3 to 100 objectives. Performance was assessed using the convergence measure (CM) and the inverted generational distance (IGD), as recommended in [2, 9]. The results of this experiment are shown in terms of rankings in Figures 1 and 2. We also compared wall clock runtimes in Figure 3.

The results make a compelling case for ϵ-lexicase selection, especially for 5 to 100 objectives. In this range, it finds solution populations with better convergence measures than HypE or NSGA-II (p<0.01) with no significant difference from HypE in terms of IGD, a measure that takes into account spread along the Pareto front. In addition to these promising results, ϵ-lexicase selection finishes in significantly less time than the other methods for 25 to 100 objectives.

Future work should consider more adaptations of ϵ-lexicase selection to incorporate concepts from other EMOs. In a broader sense, the success of ϵ-lexicase selection suggests that it is useful to promote solutions near Pareto-set boundaries that perform well on randomized subsets of other objectives.

[1] Code: http://github.com/lacava/emo-lex

GECCO '18 Companion, July 15–19, 2018, Kyoto, Japan

ACM ISBN 978-1-4503-5764-7/18/07…$15.00
https://doi.org/10.1145/3205651.3205656

Algorithm 1 : ϵ**-Lexicase Selection** applied to individuals $x \in \mathcal{N}$ with objective values $f_i(x)$, $f_i \in \mathcal{F}$.

Selection($\mathcal{N}$, $\mathcal{F}$) :
  $\mathcal{P} \leftarrow \emptyset$ ◊ parents
  for $f_i \in \mathcal{F}$: ◊ get ϵ for each f_i
    $\epsilon_i \leftarrow \lambda(\mathbf{f_i})$
  do N **times**:
    $\mathcal{P} \leftarrow \mathcal{P} \cup$ GetParent($\mathcal{N}$, $\mathcal{F}$, ϵ) ◊ add selection to $\mathcal{P}$

GetParent($\mathcal{N}$, $\mathcal{F}$, ϵ) :
  $\mathcal{F}' \leftarrow \mathcal{F}$ ◊ objectives
  $\mathcal{S} \leftarrow \mathcal{N}$ ◊ selection pool
  while $|\mathcal{F}'| > 0$ **and** $|\mathcal{S}| > 1$:
    $f_i \leftarrow$ random choice from $\mathcal{F}'$ ◊ pick random f_i
    $f_i^* \leftarrow \min f_i(x)$ **for** $x \in \mathcal{S}$ ◊ best score on f_i in pool
    for $x \in \mathcal{S}$: ◊ filter pool
      if $f_i(x) > f_i^* + \epsilon_m$ **then**
        $\mathcal{S} \leftarrow \mathcal{S} \setminus \{x\}$
    $\mathcal{F}' \leftarrow \mathcal{F}' \setminus \{f_i\}$ ◊ remove f_i
  return random choice from $\mathcal{S}$

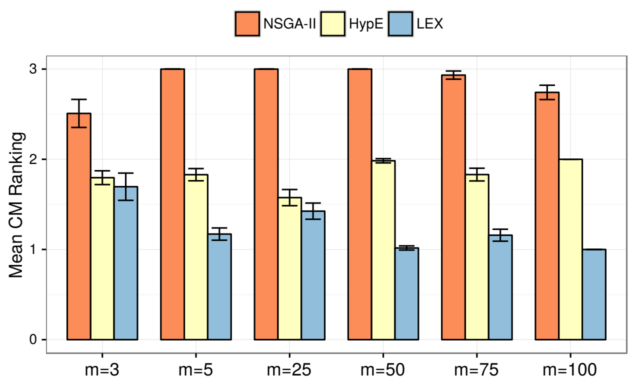

Figure 1: Average CM rankings of each algorithm as a function of number of objectives m**.**

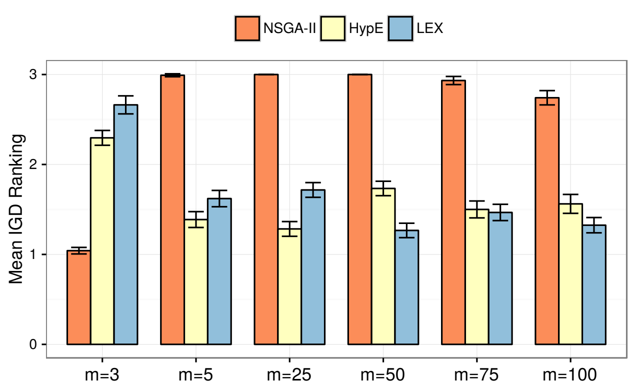

Figure 2: Average IGD rankings of each algorithm as a function of number of objectives m**.**

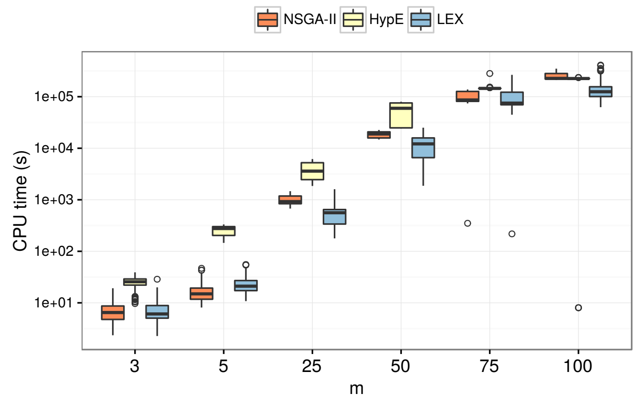

Figure 3: CPU time for each algorithm as a function of number of objectives m**.**

2 ACKNOWLEDGMENTS

This work was supported by NIH grants LM010098 and AI116794.

REFERENCES

[1] Johannes Bader and Eckart Zitzler. 2011. Hype: An Algorithm for Fast Hypervolume-based Many-objective Optimization. *Evol. Comput.* 19, 1 (March 2011), 45–76. https://doi.org/10.1162/EVCO_a_00009
[2] Shelvin Chand and Markus Wagner. 2015. Evolutionary many-objective optimization: A quick-start guide. *Surveys in Operations Research and Management Science* 20, 2 (Dec. 2015), 35–42. https://doi.org/10.1016/j.sorms.2015.08.001
[3] Kalyanmoy Deb, Lothar Thiele, Marco Laumanns, and Eckart Zitzler. 2005. *Scalable test problems for evolutionary multiobjective optimization.* Springer. http://link.springer.com/10.1007/1-84628-137-7_6
[4] Marco Farina and Paolo Amato. 2002. On the optimal solution definition for many-criteria optimization problems. In *Fuzzy Information Processing Society, 2002. Proceedings. NAFIPS. 2002 Annual Meeting of the North American.* IEEE, 233–238. http://ieeexplore.ieee.org/abstract/document/1018061/
[5] Hisao Ishibuchi, Noritaka Tsukamoto, and Yusuke Nojima. 2008. Evolutionary many-objective optimization: A short review.. In *IEEE congress on evolutionary computation.* Citeseer, 2419–2426. http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.139.2046&rep=rep1&type=pdf
[6] Joshua D. Knowles, Richard A. Watson, and David W. Corne. 2001. Reducing local optima in single-objective problems by multi-objectivization. In *International Conference on Evolutionary Multi-Criterion Optimization.* Springer, 269–283.
[7] William La Cava, Thomas Helmuth, Lee Spector, and Jason H. Moore. 2018. A probabilistic and multi-objective analysis of lexicase selection and epsilon-lexicase selection. *Evolutionary Computation Journal* (May 2018). http://arxiv.org/abs/1709.05394 arXiv: 1709.05394.
[8] William La Cava, Lee Spector, and Kourosh Danai. 2016. Epsilon-Lexicase Selection for Regression. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016 (GECCO '16).* ACM, New York, NY, USA, 741–748. https://doi.org/10.1145/2908812.2908898
[9] Ke Li, Kalyanmoy Deb, Tolga Altinoz, and Xin Yao. 2017. Empirical investigations of reference point based methods when facing a massively large number of objectives: First results. In *International Conference on Evolutionary Multi-Criterion Optimization.* Springer, 390–405.
[10] Pawel Liskowski, Krzysztof Krawiec, Thomas Helmuth, and Lee Spector. 2015. Comparison of Semantic-aware Selection Methods in Genetic Programming. In *Proceedings of the Companion Publication of the 2015 Annual Conference on Genetic and Evolutionary Computation (GECCO Companion '15).* ACM, New York, NY, USA, 1301–1307. https://doi.org/10.1145/2739482.2768505
[11] Lee Spector. 2012. Assessment of problem modality by differential performance of lexicase selection in genetic programming: a preliminary report. In *Proceedings of the fourteenth international conference on Genetic and evolutionary computation conference companion.* 401–408. http://dl.acm.org/citation.cfm?id=2330846
[12] Tobias Wagner, Nicola Beume, and Boris Naujoks. 2007. Pareto-, Aggregation-, and Indicator-Based Methods in Many-Objective Optimization. In *Evolutionary Multi-Criterion Optimization.* Springer, Berlin, Heidelberg, 742–756. http://link.springer.com/chapter/10.1007/978-3-540-70928-2_56 DOI: 10.1007/978-3-540-70928-2_56.