

Interactive Parameter Tuning of Bi-objective Optimisation Algorithms Using the Empirical Attainment Function

Manuel López-Ibáñez
Universidad de Málaga
Málaga, Spain
manuel.lopez-ibanez@uma.es

Juan Esteban Diaz
Universidad San Francisco de Quito
Quito, Ecuador
jediaz@usfq.edu.ec

ABSTRACT

We propose a visual approach of eliciting preferences from a Decision Maker (DM) in the context of comparing the stochastic outcomes of two alternative designs or parameter configurations of an optimization algorithm for bi-objective problems. Our proposal is based on visualizing the differences between the empirical attainment functions (EAFs) of the two alternative algorithmic configurations, and then ask the DM to choose their preferred side of the differences. Information about the regions preferred by the DM is translated into a weighted hypervolume indicator that will assign higher quality values to approximation fronts that result in EAF differences preferred by the DM. This indicator may be used to guide an automatic algorithm configuration method, such as irace, to search for parameter values that perform better in the objective space regions preferred by the DM. Experiments on the well-known bi-objective quadratic assignment problem and a real-world production planning problem arising in the manufacturing industry show the benefits of the proposal.

This manuscript for the Hot-off-the-Press track at GECCO 2021 is based on the paper: "Incorporating Decision-Maker's Preferences into the Automatic Configuration of Bi-Objective Optimisation Algorithms" published by European Journal of Operational Research [6].

CCS CONCEPTS

• **Theory of computation** → **Evolutionary algorithms**; • **Applied computing** → **Multi-criterion optimization and decision-making**.

KEYWORDS

Automatic algorithm design and configuration, Multi-objective optimisation, Decision maker's preferences, Empirical Attainment Function

ACM Reference Format:

Manuel López-Ibáñez and Juan Esteban Diaz. 2021. Interactive Parameter Tuning of Bi-objective Optimisation Algorithms Using the Empirical Attainment Function. 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/10.1145/3449726.3462727>

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

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8351-6/21/07.

<https://doi.org/10.1145/3449726.3462727>

1 INTRODUCTION

In order to compare different configurations of multiobjective optimisers, most automatic configuration (AC) methods use unary quality indicators to assign a single scalar value to each approximation of the Pareto front [2]. However, the preferences of the user of the AC method, i.e., the decision maker (DM), may be different to the ones implied by such quality indicators. The incorporation of DM's preferences into quality indicators can be achieved, for instance through the weighted hypervolume indicator (HV^w) [1], but expressing preferences in terms of a weight function may be a challenging task, especially when comparing the stochastic outcomes of alternative algorithm configurations.

2 FROM EAF DIFFERENCES TO HV^w

Given a multiset of approximation fronts, the EAF [8] can be understood as the number of fronts that attain, i.e., weakly dominate, each point in the objective space. By visualising the EAFs of two stochastic multi-objective algorithms, it is possible to assess differences in the expected location of their outcomes [9], which is an effective and intuitive way to assess their performance [10, 11].

Our approach starts with two multisets, each containing n approximation fronts generated by two different algorithms. First we compute a finite representation of the rectangular *regions* of the objective space with the same value of the EAF difference between the multisets. Next, the DM is shown a visualisation of the regions and is asked to choose their preferred side (e.g., Fig. 2). The *regions* with an EAF difference in favour of the DM's choice can be seen as weights on the objective space according to the DM's preferences.

Given these *regions* and a new approximation front, we compute the value of the corresponding HV^w as a combination of the unweighted hypervolume (HV) [12] and a weighted component, which is scaled by a factor and added to HV . The weighted component takes into account the area and the EAF value of the regions dominated by the points of the approximation front. The scaling factor is relative to the maximum possible HV in order to provide a strong bias towards the regions preferred by the DM, even if their total area is much smaller than the total area of the objective space.

The resulting HV^w assigns a better quality to objective vectors or approximation fronts that attain the areas where the EAF differences are in favour of the DM's preferences, while respecting the Pareto-optimality criterion. This HV^w may be used to rank approximation fronts that are incomparable in terms of Pareto optimality, to guide a hypervolume-based optimiser [3], or to guide the AC of multi-objective optimisers [6]. We believe that eliciting preferences through this approach is simpler and more intuitive than directly specifying a weight function, as the DM only needs to select between the two sides of a plot similar to Fig 2.

3 RESULTS

We demonstrate our approach on the configuration of the length parameter of Weighted Robust Taboo Search (W-RoTS) [9] for the bi-objective quadratic assignment problem (P1). We start from two initial configurations: P1C1 ($\ell = 1$) performs better along the front except for a few specific regions and P1C2 ($\ell = 100$) returns higher-quality solutions at the extremes of the front. From their EAF differences, we obtain two different HV^w indicators and use them to guide two different settings of irace, denoted as P1iraceC1 and P1iraceC2. Each configuration returned by irace is run 30 times. We simulate DM's preferences for the centre of the front and for extreme values of either objective by computing their weighted hypervolume according to a goal-oriented weight distribution w_{goal} [1] and an exponential weight distribution w_{exp} [1], respectively, which irace did not know about. Figure 1 (see [6] for complete details) shows that the configurations generated by irace follow the true preferences of the DM, and thus confirms that our elicitation method works as expected.

Moreover, we demonstrate that our approach also works on a real-world production planning problem under uncertainty (P2) [4, 5], where no prior knowledge is available about the algorithm configurations that would satisfy the DM's preferences. Figure 2 (see [6] for complete details) shows that irace returned configurations that are clearly specialised towards satisfying two alternative preferences of the DM, confirming once again the effectiveness of our approach.

4 CONCLUSIONS

Our results show that the algorithm configurations found by irace when guided by the HV^w computed from EAF differences not only outperformed the ones generated when guided by the unweighted HV , but also performed significantly better on preferred regions of the objective space, thus confirming that the AC method was indeed biased by the DM's preferences. To the best of our knowledge, this is the first method proposed to incorporate DM's preferences into the AC of multi-objective optimizers. Given its generality, it can be integrated into other AC methods different from irace.

REFERENCES

- [1] A. Auger, J. Bader, D. Brockhoff, and E. Zitzler. 2009. Articulating User Preferences in Many-Objective Problems by Sampling the Weighted Hypervolume. In *GECCO*, F. Rothlauf (Ed.). ACM Press, New York, NY, 555–562.
- [2] L. C. T. Bezerra, M. López-Ibáñez, and T. Stützle. 2020. Automatic Configuration of Multi-objective Optimizers and Multi-objective Configuration. In *High-Performance Simulation-Based Optimization*, T. Bartz-Beielstein et al. (Eds.). Springer International Publishing, Cham, Switzerland, 69–92. https://doi.org/10.1007/978-3-030-18764-4_4
- [3] D. Brockhoff, J. Bader, L. Thiele, and E. Zitzler. 2013. Directed Multiobjective Optimization Based on the Weighted Hypervolume Indicator. *Journal of Multi-Criteria Decision Analysis* 20, 5-6 (2013), 291–317. <https://doi.org/10.1002/mcda.1502>
- [4] J. E. Diaz, J. Handl, and D.-L. Xu. 2017. Evolutionary robust optimization in production planning: interactions between number of objectives, sample size and choice of robustness measure. *Comput. Oper. Res.* 79 (2017), 266–278. <https://doi.org/10.1016/j.cor.2016.06.020>
- [5] J. E. Diaz, J. Handl, and D.-L. Xu. 2018. Integrating meta-heuristics, simulation and exact techniques for production planning of a failure-prone manufacturing system. *Eur. J. Oper. Res.* 266, 3 (2018), 976–989. <https://doi.org/10.1016/j.ejor.2017.10.062>
- [6] J. E. Diaz and M. López-Ibáñez. 2021. Incorporating Decision-Maker's Preferences into the Automatic Configuration of Bi-Objective Optimisation Algorithms. *Eur. J. Oper. Res.* 289, 3 (2021), 1209–1222. <https://doi.org/10.1016/j.ejor.2020.07.059>
- [7] C. M. Fonseca, A. P. Guerreiro, M. López-Ibáñez, and L. Paquete. 2011. On the Computation of the Empirical Attainment Function. In *EMO*, R. H. C. Takahashi

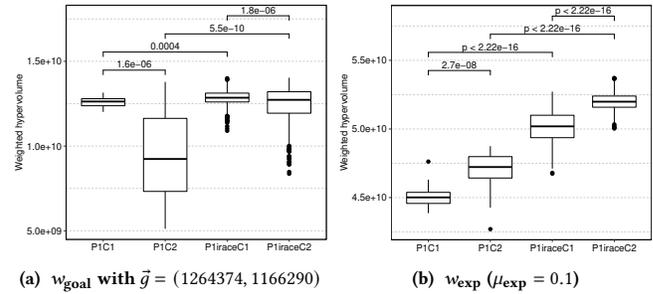


Figure 1: Boxplots of HV^w values according to two different weight distributions: w_{goal} (left) and w_{exp} (right). The same approximation fronts are evaluated on both subfigures. Sets P1C1 and P1C2 contain 30 fronts each (30 runs of one configuration of W-RoTS). Sets P1iraceC1 and P1iraceC2 contain 900 fronts each (30 runs of the best configuration of W-RoTS returned by each of the 30 independent runs of irace). Pairs joined by a line were compared with the Wilcoxon rank-sum test, and the resulting p-value is shown above the line.

(Reproduced from [6])

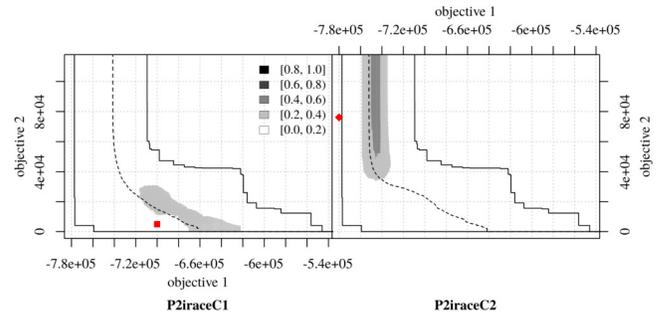


Figure 2: EAF differences between P2iraceC1 and P2iraceC2. ■ (left) and ◆ (right) show the location of the two goals that simulate the DM's true preferences. P2iraceC1 and P2iraceC2 denote the 30 configurations given by irace when the DM prefers each goal, respectively. Each side contains data from the approximation fronts generated by 30 runs of the best configuration returned by each of 30 runs of irace guided by the corresponding HV^w . (Reproduced from [6])

- et al. (Eds.). *LNCS*, Vol. 6576. Springer, 106–120. https://doi.org/10.1007/978-3-642-19893-9_8
- [8] V. Grunert da Fonseca and C. M. Fonseca. 2010. The Attainment-Function Approach to Stochastic Multiobjective Optimizer Assessment and Comparison. In *Experimental Methods for the Analysis of Optimization Algorithms*, T. Bartz-Beielstein et al. (Eds.). Springer, Berlin, Germany, 103–130.
- [9] M. López-Ibáñez, L. Paquete, and T. Stützle. 2006. Hybrid Population-based Algorithms for the Bi-objective Quadratic Assignment Problem. *Journal of Mathematical Modelling and Algorithms* 5, 1 (2006), 111–137. <https://doi.org/10.1007/s10852-005-9034-x>
- [10] M. López-Ibáñez, L. Paquete, and T. Stützle. 2010. Exploratory Analysis of Stochastic Local Search Algorithms in Biobjective Optimization. In *Experimental Methods for the Analysis of Optimization Algorithms*, T. Bartz-Beielstein et al. (Eds.). Springer, Berlin, Germany, 209–222. https://doi.org/10.1007/978-3-642-02538-9_9
- [11] E. Zitzler, J. D. Knowles, and L. Thiele. 2008. Quality Assessment of Pareto Set Approximations. In *Multi-objective Optimization: Interactive and Evolutionary Approaches*, J. Branke et al. (Eds.). LNCS, Vol. 5252. Springer, 373–404.
- [12] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. Grunert da Fonseca. 2003. Performance Assessment of Multiobjective Optimizers: an Analysis and Review. *IEEE Trans. Evol. Comput.* 7, 2 (2003), 117–132.