A decision tree-based interactive evolutionary multi-objective algorithm

Seyed Mahdi Shavarani[1], Mahmoud Golabi[2], Lhassane Idoumghar[2]

[1] Centre for Logistics and Sustainability Analytics, Kent Business School, University of Kent, UK
m.shavarani@kent.ac.uk

[2] Université de Haute-Alsace, IRIMAS UR 7499, F-68100 Mulhouse, France
{mahmoud.golabi, lhassane.idoumghar}@uha.fr

Mots-clés : *Interactive methods, evolutionary multi-objective algorithms, decision tree, preference learning, machine decision-maker*

1 Introduction

Interactive Evolutionary Multi-Objective Algorithms (iEMOAs) iteratively gather and utilize the Decision Maker's (DM) discriminations and preferences to enhance selection pressure and resolve ties among solutions within the same front to direct the search towards DM's Most Preferred Solution (MDM).

Despite their potentials, there is limited empirical evidence supporting the performance of iEMOAs under realistic conditions [1]. To address challenges encountered in testing with human decision makers (DMs), recent studies have opted to utilize a Machine DMs (MDM) for statistically measuring the performance of iEMOAs. [3]. Findings indicate that iEMOAs may underperform in realistic scenarios, particularly when faced with a higher number of objectives or when preference information is influenced by human-specific biases like inconsistent decisions and fatigue. This suggests a lack of robustness in algorithms that excel under ideal conditions when confronted with biases and complications typical in real-world situations. The main causes emanate from the core features of interactive methods, including interaction style and preference learning.

Considering the accuracy in reflecting the DM preferences as the core of any interactive methods, Shavarani et al. [2] proposed leveraging decision trees (DTs) to gain insights into advantageous trade-offs and predict the preferred solution from the DM's perspective in pairwise comparisons. They used the trained DT to give an score to each individual in the population in comparison with all other individuals. This study investigates the efficiency of decision tree regressor in scoring individuals directly based on preference information which is elicited in form of ranking a subset of solutions and avoid comprehensive comparisons to speed up the learning-to-rank problems.

2 Methodology

This study employs the Non-dominated Sorting Genetic Algorithm (NSGA-II) as the underlying evolutionary multi-objective optimizer. After several generations of NSGA-II, the algorithm pauses to initiate its first interaction, obtaining a ranking from a subset of solutions and subsequently constructing the training set T. A DT is then trained using T to forecast the favored solution in pairwise comparisons. In successive generations, DT computes scores for solutions, facilitating the ranking of solutions that share an identical non-dominated sorting rank. Essentially, solution scores replace the crowding distance of NSGA-II to distinguish between non-dominated solutions. Solutions with higher scores are more likely to survive,

TAB. 1: Accurany in preference learning

Interactions	Regression DT	Binary Classification DT
1	44.80	69.62
2	42.54	74.18
3	40.24	75.48
4	43.09	77.44
5	42.14	77.88
6	40.38	79.17

participate in mating, and contribute to the generation of new offspring. It is worth mentioning that each subsequent interaction supplements T with more training examples, progressively refining the accuracy of the model.

This study uses both the binary and the regression DTs for preference learning. The construction of the training set in binary DT is done using the pairwise transformation that computes the differences in objective vectors for pairwise comparisons and assigns binary ranks to them. The regression DT for learning to rank is employed to leverage the ranking provided by a DM on a set of solutions, enabling the algorithm to learn and predict solution rankings based on the acquired preferences. Comparing the obtained results using these approaches allows for a comprehensive assessment of the efficacy of distinct preference learning methods, offering insights into the strengths and limitations of each approach across different scenarios. An MDM using a version of the Sigmoid utility function is applied to analyze the performance of these methods [3]. Considering the obtained utility values, the accuracy rates of preference learning using binary and regression DT are presented in Table 1.

3 Conclusions and perspectives

This study assesses the effectiveness of binary classification and regression DTs in learning DM's preferences within an interactive multi-objective evolutionary optimization algorithm. The findings underscore the exceptional efficacy of binary classification decision trees (DTs) in preference learning. Noteworthy is the consistently high accuracy across diverse interactions, emphasizing the binary DT's proficiency in harnessing accumulated data. This stems from its adept handling of trade-offs among distinct objective values during pairwise solution comparisons, ensuring reliability even as the algorithm converges toward the true Pareto Front. In contrast, the DT regressor, trained on objective values, faces challenges in maintaining accuracy when the population's position shifts, rendering it less reliable for unseen data. As an extension, further investigation could be done to explore the efficacy of weighted DTs or alternative learning frameworks to enhance preference learning abilities.

References

[1] Afsar, B., Miettinen, K., Ruiz, F.: Assessing the performance of interactive multiobjective optimization methods: a survey. ACM Computing Surveys (CSUR) **54**(4), 1–27 (2021)

[2] Shavarani, S.M., López-Ibáñez, M., Allmendinger, R., Knowles, J.: An interactive decision tree-based evolutionary multi-objective algorithm. In: International Conference on Evolutionary Multi-Criterion Optimization. pp. 620–634. Springer (2023)

[3] Shavarani, S.M., López-Ibáñez, M., Knowles, J.: Realistic utility functions prove difficult for state-of-the-art interactive multiobjective optimization algorithms. In: Proceedings of the Genetic and Evolutionary Computation Conference. pp. 457–465 (2021)