

A General Interactive Approach for Solving Multi-Objective Combinatorial Optimization Problems with Imprecise Preferences

Nawal Benabbou, Thibaut Lust

Sorbonne Université, CNRS, Laboratoire d'Informatique de Paris 6, LIP6 F-75005 Paris, France
nawal.benabbou@lip6.fr, thibaut.lust@lip6.fr

Abstract

We consider multi-objective combinatorial optimization problems where preferences can be represented by a parameterized scalarizing function that is linear in its parameters. Assuming that the parameters are initially not known, we iteratively collect preference information from the decision maker until being able to identify her preferred solution. To obtain informative preference information, we ask the decision maker to compare two promising solutions generated using the extreme points of the polyhedron representing the admissible parameters. Moreover, our stopping criterion guarantees that the returned solution is optimal (or near-optimal) according to the decision maker's preferences. Finally, we provide numerical results to demonstrate the practical efficiency of our approach.

Introduction

Multi-objective optimization is concerned with optimization problems involving several (conflicting) objectives to be optimized simultaneously. Without preference information, we only know that the best solution for the decision maker (DM) is among the Pareto-optimal solutions. Since the number of these solutions can be exponential in the size of the problem, one may want to refine Pareto dominance with preferences to be able to determine the best solution for the DM. We assume here that the DM's preferences can be represented by a parameterized scalarizing function, the corresponding parameters being initially unknown, and we study the potential of incremental elicitation (White III, Sage, and Dozono 1984; Wang and Boutilier 2003) in this setting.

Preference elicitation on combinatorial domains is an active topic that has been recently studied in various contexts, e.g., (Korhonen 2005; Regan and Boutilier 2011; Weng and Zanuttini 2013; Drummond and Boutilier 2014; Benabbou and Perny 2016; Kaddani et al. 2017; Benabbou and Perny 2018). Our aim is to propose a general interactive approach for multi-objective optimization with imprecise preference parameters. Our approach is general in the sense that it can be applied to any multi-objective optimization problem, providing that the scalarizing function is linear in its parameters (e.g., weighted sums, Choquet integrals) and that there exists an efficient algorithm to solve the problem when the parameters are known.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Background and Notations

We consider a general multi-objective combinatorial optimization (MOCO) problem with n objective functions $y_i, i \in \{1, \dots, n\}$, to be minimized. Any solution $x \in \mathcal{X}$ is therefore associated with a cost vector $y(x) = (y_1(x), \dots, y_n(x)) \in \mathbb{R}^n$ where $\mathcal{X}$ is the feasible set in the decision space and $y_i(x)$ is the evaluation of x on the i -th objective. Solutions are here compared through their images in the objective space. We assume that the DM's preferences can be represented by a scalarizing function f_ω that is linear in its parameters ω ; hence $x \in \mathcal{X}$ is preferred to $x' \in \mathcal{X}$ iff $f_\omega(y(x)) \leq f_\omega(y(x'))$. Assuming that ω are initially not known, we consider the set Ω of all parameters satisfying $f_\omega(u) \leq f_\omega(v)$ for every collected pair $(u, v) \in \mathbb{R}^n \times \mathbb{R}^n$ such that u is known to be preferred to v . Since f_ω is linear in ω , we can assume that Ω is a convex polyhedron throughout the paper. To reduce parameter imprecision, we use the minimax regret criterion (MMR) which can be defined using pairwise max regrets (PMR) and max regrets (MR):

Definition 1. For any two solutions $x, x' \in \mathcal{X}$:
 $PMR(x, x', \Omega) = \max_{\omega \in \Omega} \{f_\omega(y(x)) - f_\omega(y(x'))\}$
 $MR(x, \mathcal{X}, \Omega) = \max_{x' \in \mathcal{X}} PMR(x, x', \Omega)$
 $MMR(\mathcal{X}, \Omega) = \min_{x \in \mathcal{X}} MR(x, \mathcal{X}, \Omega)$

The set $\arg \min_{x \in \mathcal{X}} MMR(x, \mathcal{X}, \Omega)$ is the set of MMR-optimal solutions, allowing to minimize the worst-case loss; note that if $MMR(\mathcal{X}, \Omega) = 0$, then these solutions are necessarily optimal according to the DM's preferences. As observed in previous works, $MMR(\mathcal{X}, \Omega)$ can only decrease when collecting new preference data. This observation has led to the following incremental elicitation approach: reduce parameter imprecision by asking queries to the DM in an iterative way, until $MMR(\mathcal{X}, \Omega)$ drops below a given threshold $\delta \geq 0$ representing the maximum allowable gap to optimality.

A General Interactive Method

At each step, $MMR(\mathcal{X}, \Omega)$ could be obtained by computing $PMR(x, x', \Omega)$ for all pairs $x, x' \in \mathcal{X}$. However this would not be efficient for MOCO problems due to the large size of $\mathcal{X}$. We propose instead to combine incremental elicitation and search as follows: at each step, we generate a (small) set of promising solutions using the extreme points of Ω , we ask the DM to compare two of these solutions, we update Ω according to her answer and we stop the process whenever

a (near-)optimal solution is detected (i.e. a solution $x \in \mathcal{X}$ s.t. $MR(x, \mathcal{X}, \Omega) \leq \delta$). More precisely, taking as input a MOCO problem P , a threshold $\delta \geq 0$, a scalarizing function f_ω with unknown parameters and an initial set Ω of feasible parameters, our algorithm iterates as follows:

1. We generate the set EP_Ω of all extreme points of polyhedron Ω ; its k th element is denoted by ω^k .
2. For all $\omega^k \in EP_\Omega$, we compute x^k the optimal solution of problem P for the precise function f_{ω^k} .
3. We compute $MMR(X_\Omega, \Omega)$, where $X_\Omega = \bigcup_{k=1}^{|EP_\Omega|} \{x^k\}$. If $MMR(X_\Omega, \Omega) > \delta$, then the DM is asked to compare two solutions $x, x' \in X_\Omega$ and we update Ω by imposing the linear constraint $f_\omega(x) \leq f_\omega(x')$ (or $f_\omega(x) \geq f_\omega(x')$ depending on her answer); otherwise, the algorithm returns a solution x^* in $\arg \min_{x \in X_\Omega} MR(x, X_\Omega, \Omega)$.

Our algorithm is called IEEP for Incremental Elicitation based on Extreme Points. Its validity is established by the following proposition (proof omitted due to space constraints):

Proposition 1. *For any positive threshold δ , algorithm IEEP returns a solution $x^* \in \mathcal{X}$ such that $MR(x^*, \mathcal{X}, \Omega) \leq \delta$.*

Experimental Results

We focus here on the multicriteria spanning tree (MST) problem¹. We generate instances of graph $G = (V, E)$ with a number of vertices $|V|$ varying between 50 and 100 and a number of objectives n ranging from 3 to 5. The edge costs are drawn within $\{1, \dots, 1000\}^n$ uniformly at random. The DM's preferences are here represented by a weighted sum and we start the execution with no preference information. In our experiments, we simulate answers to queries by generating a weighting vector ω uniformly at random before running the algorithm. We compare our interactive method with the state-of-the-art method proposed in (Benabbou and Perny 2015) called IE-Prim hereafter. For both procedures, we consider the following query generation strategies:

- **Random:** The DM is asked to compare two solutions that are randomly chosen in X_Ω .
- **Max-Dist:** We choose a pair of solutions in X_Ω that maximizes the Euclidean distance in the objective space.
- **CSS:** The Current Solution Strategy (Boutillier et al. 2006).

These strategies are evaluated in terms of running time and number of generated preference queries. Results obtained by averaging over 30 runs are given in Table 1 for $\delta = 0$. First, we see that IEEP outperforms IE-Prim in all settings. Then, we observe that Random and Max-Dist strategies are much faster than CSS strategy. Finally, we see that Max-Dist is the best strategy for minimizing the number of preference queries, which is quite surprising since CSS is intensively used in incremental elicitation (e.g., (Boutillier et al. 2006)).

Acknowledgements

This work was supported by the Paris Ile-de-France Region.

¹Tests were performed on a Intel Core i7-7700, at 3.60GHz; the extreme points of Ω are generated by `polymake` library; MST problems with precise weights are solved using Prim algorithm.

		IEEP - Random		IEEP - Max Dist		IEEP - CSS	
n	$ V $	time (s)	queries	time (s)	queries	time (s)	queries
3	50	16.91	16.17	16.46	15.23	17.91	16.90
3	100	18.66	17.57	19.00	17.37	18.97	17.73
4	50	27.51	25.73	26.40	24.60	30.69	28.93
4	100	31.99	29.90	30.11	28.43	34.83	32.47
5	50	37.73	35.03	36.24	34.33	42.29	39.83
5	100	41.84	39.77	42.07	39.20	55.89	51.70

		IE-Prim - Random		IE-Prim - Max-Dist		IE-Prim - CSS	
n	$ V $	time (s)	queries	time (s)	queries	time (s)	queries
3	50	28.56	26.67	26.05	24.50	31.90	29.57
3	100	34.60	32.37	33.61	31.07	36.93	35.27
4	50	45.02	42.13	42.48	39.70	55.55	50.83
4	100	55.56	51.60	54.73	51.17	66.60	61.63
5	50	59.71	55.47	56.94	53.23	80.41	73.40
5	100	75.36	70.70	76.39	71.73	103.69	95.33

Table 1: Comparison between IEEP and IE-Prim algorithms.

References

- Benabbou, N., and Perny, P. 2015. On possibly optimal trade-offs in multicriteria spanning tree problems. In *Proceedings of ADT'15*, 322–337.
- Benabbou, N., and Perny, P. 2016. Solving multi-agent knapsack problems using incremental approval voting. In *Proceedings of ECAI'16*, 1318–1326.
- Benabbou, N., and Perny, P. 2018. Interactive resolution of multiobjective combinatorial optimization problems by incremental elicitation of criteria weights. *EURO journal on decision processes* 6:283–319.
- Boutillier, C.; Patrascu, R.; Poupert, P.; and Schuurmans, D. 2006. Constraint-based optimization and utility elicitation using the minimax decision criterion. *AI* 170(8–9):686–713.
- Drummond, J., and Boutillier, C. 2014. Preference elicitation and interview minimization in stable matchings. In *Proceedings of AAAI'14*, 645–653.
- Kaddani, S.; Vanderpooten, D.; Vanpeperstraete, J.-M.; and Aissi, H. 2017. Weighted sum model with partial preference information: application to multi-objective optimization. *European Journal of Operational Research* 260:665–679.
- Korhonen, P. 2005. *Multiple criteria decision analysis*. Greco, Salvatore and Figueira, J and Ehrgott, M. Springer.
- White III, C. C.; Sage, A. P.; and Dozono, S. 1984. A model of multiattribute decisionmaking and trade-off weight determination under uncertainty. *IEEE Transactions on Systems, Man, and Cybernetics* 14(2):223–229.
- Regan, K., and Boutillier, C. 2011. Eliciting additive reward functions for markov decision processes. In *Proceedings of IJCAI'11*, 2159–2164.
- Wang, T., and Boutillier, C. 2003. Incremental Utility Elicitation with the Minimax Regret Decision Criterion. In *Proceedings of IJCAI'03*, 309–316.
- Weng, P., and Zanuttini, B. 2013. Interactive value iteration for markov decision processes with unknown rewards. In *Proceedings of IJCAI'13*, 2415–2421.