POSTER

Lower Confidence Bound for Preference Selection in Interactive Multi-Objective Optimization

ARASH HEIDARI, Ghent University, Ghent, VOV, Belgium

SEBASTIAN ROJAS-GONZALEZ, Ghent University, Ghent, VOV, Belgium

TOM DHAENE, Ghent University, Ghent, VOV, Belgium

IVO COUCKUYT, Ghent University, Ghent, VOV, Belgium

Published: 14 July 2024

Citation in BibTeX format

GECCO '24 Companion: Genetic and Evolutionary Computation Conference Companion
July 14 - 18, 2024
VIC, Melbourne, Australia

Conference Sponsors:
SIGEVO

GECCO '24 Companion: Proceedings of the Genetic and Evolutionary Computation Conference Companion (July 2024)
https://doi.org/10.1145/3638530.3654317
ISBN: 9798400704956

Lower Confidence Bound for Preference Selection in Interactive Multi-Objective Optimization

Arash Heidari
Arash.Heidari@UGent.be
Ghent University - imec
Ghent, Belgium

Sebastian Rojas Gonzalez
Sebastian.RojasGonzalez@UGent.be
Ghent University - imec
Ghent, Belgium

Tom Dhaene
Tom.Dhaene@UGent.be
Ghent University - imec
Ghent, Belgium

Ivo Couckuyt
Ivo.Couckuyt@UGent.be
Ghent University - imec
Ghent, Belgium

ABSTRACT

In multi-objective optimization, the goal is to find the *non-dominated* or *Pareto-optimal* set that reveals the optimal trade-offs among the conflicting objectives. Conventionally, the Decision-Maker (DM) selects their preferred solution from this set post-optimization. Evidently, this approach necessitates the computation of numerous non-dominated solutions, incurring high computational costs when the evaluation of the objectives is expensive. To address these challenges, interactive optimization offers a potentially more efficient approach: sequential interactions with the DM during the optimization process, enabling the judicious allocation of the optimization budget by guiding the process toward the most desirable regions of the Pareto front. In this study, we propose to exploit the well-known Lower Confidence Bound acquisition function in Bayesian optimization, to interactively estimate the DM's preferred solution in a data-efficient manner, even in scenarios where the DM experiences uncertainty in their decision-making process.

CCS CONCEPTS

• **Theory of computation → Continuous optimization**; *Gaussian processes*; • **Computing methodologies → Optimization algorithms**; • **Applied computing → Multi-criterion optimization and decision-making**.

KEYWORDS

Bayesian Optimization, Multi-Objective Optimization, Interactive Optimization, Preference Learning

ACM Reference Format:
Arash Heidari, Sebastian Rojas Gonzalez, Tom Dhaene, and Ivo Couckuyt. 2024. Lower Confidence Bound for Preference Selection in Interactive Multi-Objective Optimization. In *Genetic and Evolutionary Computation Conference (GECCO '24 Companion), July 14–18, 2024, Melbourne, VIC, Australia.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3638530.3654317

GECCO '24 Companion, July 14–18, 2024, Melbourne, VIC, Australia

ACM ISBN 979-8-4007-0495-6/24/07.
https://doi.org/10.1145/3638530.3654317

1 INTRODUCTION

In the realm of multi-objective optimization, the optimization process yields a set of non-dominated solutions, referred to as the *Pareto set* in input space, and the *Pareto front* in output space. A solution belonging to this set is not dominated by any other solution and reflects the optimal trade-offs between the objectives [14]. Without loss of generality, a multi-objective problem can be defined as follows:

$$\textit{minimize } f_1(\mathbf{x}), f_2(\mathbf{x}), ..., f_m(\mathbf{x}) \qquad \mathbf{x} \in \Omega \subseteq \mathbb{R}^n \tag{1}$$

where $\mathbf{x}$ is an n-dimensional vector of inputs evaluated in m objectives, and Ω indicates the design space.

Standard multi-objective optimization algorithms often aim to estimate a set of solutions on the Pareto front; note that such a front is often continuous and a discrete set is used as an approximation [8]. However, obtaining a comprehensive and diverse Pareto front may present computational challenges or, depending on the problem, financial constraints. Moreover, the Decision-Maker (DM) typically seeks to explore specific regions of the Pareto front, rather than having interest in the entire set. In practice the DM will likely implement one Pareto-optimal solution – the one which reflects the *most preferred* trade-off. As a result, typical a posteriori methods will likely incur high computational costs, since an approximation of the *entire* Pareto front needs to be obtained first.

An alternative strategy for addressing this challenge is through interactive optimization [15]. This method involves iterative interactions with the DM during the optimization process, which allows the optimization to be guided, and thus it mitigates the unnecessary waste of resources by allocating the budget efficiently towards the exploration of the preferred regions of the Pareto front. Interactive methods predominantly leverage two principal frameworks for addressing black-box optimization problems: Evolutionary Algorithms (EAs) [5] and Bayesian optimization (BO) [9]. EAs are renowned for their effectiveness in tackling complex black-box problems, but at the same time are acknowledged for being computationally intensive. In scenarios where the evaluation cost is high, the utilization of EAs becomes either computationally prohibitive or impractical. While surrogate models may be embedded within an evolutionary framework to mitigate the high evaluation costs, BO algorithms offer a powerful alternative to expensive black-box optimization by exploiting the surrogate information in a principled way. Gaussian Processes (GPs) are often used as the surrogate of choice; by explicitly requiring priors they make the design assumptions clear, and allow for the quantification of the prediction uncertainty.

While the existing literature on interactive BO algorithms is limited, these methods can be generally categorized into two distinct classes. The algorithms in the first class [1, 2, 7, 18] often seek to understand the DM's preference by learning a *utility* function, denoted as U. This function encapsulates the DM's preferences, and the algorithms aim to find the preferred solution by maximizing this learned utility function. While the aforementioned methods offer a promising alternative, the published research manifests significant limitations; in particular, the utility functions used are often assumed a priori, and/or these functions have unknown parameters that also have to be learned during the optimization. The algorithms in the second class [10–12], on the other hand, offer flexible mechanisms to guide the search towards interesting and well-performing solutions in a heuristic manner, without the need of assuming or learning the structure of U.

Ideally, an algorithm should invest resources efficiently from the start of the optimization to find solutions that are Pareto-optimal *and* congruent with the DM's preferences. Indeed, in practice DMs are interested in an unbiased and diverse but relatively small set of solutions to choose from, all of which offer an interesting trade-off of the objectives. In this work we exploit the well-known *lower confidence bound* (LCB) acquisition function [9] for interactively selecting the most preferred region(s) of the Pareto front, without explicitly modeling the DM's preference. The proposed method is designed to optimize black-box problems, especially when evaluations are resource-intensive, by leveraging Bayesian Optimization as a tool to balance the *exploration* of interesting regions of the objective space, while at the same time *exploiting* the information collected so far to the fullest extent. Additionally, we empirically demonstrate the framework's capability to handle uncertainty in the choices of the DM.

2 LOWER CONFIDENCE BOUND FOR PREFERENCE SELECTION

Balancing the quality and diversity of solutions presented to the DM poses a challenge in multiple ways. An approach to ensure a diverse Pareto front with numerous options involves employing multi-objective optimization algorithms specifically designed to extract the entire Pareto front. However, this method shifts the paradigm from interactive optimization to a posteriori preference selection, which may not be desirable. However, leveraging surrogate models to emulate the behavior of each objective independently allows for cost-effective implementation of evolutionary multi-objective optimization algorithms (EMOAs).

Here we use a surrogate-based EMOA (see [4] for a comprehensive survey) for extracting the Pareto front at each step of the optimization. We employ GPs as it is standard in the BO literature [16] to facilitate this information, but other surrogates may be used as well. We compute the entire Pareto front using the LCB *instead* of the predicted mean, such that the observed Pareto front displays the trade-offs of the lower confidence bound. This is useful since we aim at exploring the entire search space for the most preferred regions of the Pareto front. Note that in standard BO the goal of the LCB is often to increase the number of evaluations in promising regions in the design space that haven't been explored yet; here we also exploit it to direct the search towards preferred solutions:

$$\textit{minimize}\ \hat{g_1}(\mathbf{x}), \hat{g_2}(\mathbf{x}), ..., \hat{g_m}(\mathbf{x}) \qquad \mathbf{x} \in \Omega \subseteq \mathbb{R}^n \tag{2}$$

$$\hat{g}_i = \mu(g_i) - \omega \cdot \sigma(g_i) \tag{3}$$

where g_i denotes the surrogate model for the i^{th} objective, $\mu(g_i)$ is the posterior mean of the surrogate model, $\sigma(g_i)$ represents the posterior uncertainty, and $\omega = [0, 3]$. By varying the value of ω, the algorithm can focus on local areas or explore the search space more globally [17]. This allows to offer the DM a menu of diverse solutions, with the quantity being adjustable by manipulating the population size of the evolutionary algorithm. Another advantage of this approach is its capability to present the DM with solutions that have not been evaluated yet. Algorithm 1 outlines the proposed algorithm.

Algorithm 1 LCB for Preference Selection

Input $\mathfrak{D}$: Evaluated design of experiment
1: **while** *Budget* left **do**
2: $g_i \leftarrow$ Surrogate Model for the i^{th} objective
3: $\hat{g}_i = \mu(g_i) - \omega \cdot \sigma(g_i)$ ▷ Eq. 3
4: $\widehat{P} \leftarrow$ Pareto front of $\hat{g}_i$ for $i = 1, ..., m$ ▷ Eq. 2
5: $c \leftarrow$ DM's preference from $\widehat{P}$ ▷ Interaction with the DM
6: $\mathfrak{D} = \mathfrak{D} \cup \{c, F(c)\}$ ▷ Evaluation of the selected solution
7: Reduce *Budget*
8: **end while**

Algorithm 1 crucially differs from the conventional BO paradigm in Step 5. In standard BO the acquisition function quantifies the exploration-exploitation balance of each solution, and determines which solution should be evaluated next by selecting the maximizer of the acquisition function, all without human intervention. In contrast, in our algorithm the DM *partially* assumes the role of the acquisition function, actively *selecting* the most preferred solution for evaluation in the current step. It is partial because the exploration-exploitation balance is already conveyed by the LCB of the Pareto front, and the DM selects the most preferred solution (as opposed to the one maximizing the LCB). An important remark is that in the experimental results we present we consider up to 3 objectives, where visualization is still possible. Thus, we can provide an entire approximation of the *current* Pareto front based on LCB, and the DM will select or pinpoint the most preferred solution observed. Later in Section 4 we discuss the additional challenges w.r.t. scalability in objective space.

Evidently, the transition of the DM to act as an acquisition function introduces a high demand for frequent interactions, necessitating engagement at every optimization step, which may pose challenges or prove infeasible depending on the nature of the problem at hand. However, this shift in responsibility empowers the DM to guide the optimization process based on their preferences and insights w.r.t. the *entire* objective space and the observed trade-offs (as opposed to e.g., pairwise comparisons), such that the DM can not only learn from the performance trade-offs, but also about the location, geometry and *uncertainty* on the Pareto front.

3 EXPERIMENTS

3.1 Performance Evaluation

To assess the performance of the proposed algorithm, we conducted tests on various well-known benchmark multiobjective functions presented in Table 1.

Table 1: Summary of the benchmark Problems

Name	Input Dimension	# of Objectives
DEB2DK [3]	8	2
DTLZ7 [6]	8	3
ZDT1 [19]	8	2
ZDT3 [19]	8	2

Given that the Pareto front is known for these artificial test functions, the DM's *true preference* is simulated through scalarization and random weight assignment. In essence, after computing the Pareto front with any well-known EMOA, the different objective outcomes are combined using a Tchebycheff scalarization (Eq. 4) with random weights, referred to as *preferred weights* hereafter. The non-dominated solution with the minimum scalarized objective is then identified as the true preference of the DM during the optimization process, which is used solely for performance evaluation.

$$U = max_i[w_i \cdot f_i] + \rho \sum_{i=1}^{m}[w_i \cdot f_i], \qquad \sum_{i=1}^{m} w_i = 1 \tag{4}$$

We use the Opportunity Cost (OC) as the metric to evaluate the performance of each algorithm [18]. OC quantifies the distance between a pre-defined true preferred solution and any other evaluated solution, defined as follows:

$$OC = U(F(\mathbf{x})) - U(F(\mathbf{x}^*)) \tag{5}$$

where $\mathbf{x}^*$ denotes the true preference of the DM estimated using the Pareto front of each function, $\mathbf{x}$ is any solution selected for evaluation, and $U(\cdot)$ represents the scalarization of different objectives $F(\cdot)$. Evidently, Equation 5 is to be minimized. Therefore, to model the behavior of the DM, at each interaction, the solution which minimizes the utility of $\widehat{P}$ is chosen as the preference of the DM at each step (see Step 4 in Algorithm 1).

3.2 Uncertainty from the Decision-Maker

To simulate the uncertainty and exploration of the DM, the chosen solution is not precisely the one that minimizes the scalarized function with the preferred weights. Instead, we introduce a 10% chance that the DM selects a solution randomly, modelling their curiosity to explore other seemingly interesting regions on the current front. Conversely, there is a 90% chance that the DM will focus on the current preference. In this case, the preferred weights still undergo a modification according to Eq. 6 to simulate the uncertainty of the DM:

$$\mathbf{w_{modified}} = \frac{\mathbf{w_{preferred}} * \theta}{\sum_i w_{preferred,i} * \theta}, \quad \theta \sim U(1-\eta, 1+\eta)^m \tag{6}$$

where $\mathbf{w_{preferred}}$ and $\mathbf{w_{modified}}$ denote the preferred and modified weights, respectively. The modification of the preferred weights involves multiplying the weights by a random number sampled based on the uncertainty degree, η, followed by normalization to ensure that the elements of the modified weight sum up to 1. As η increases, the modified weight is more likely to deviate further from the preferred weights, simulating higher levels of uncertainty.

3.3 Results

We evaluate the performance against two state-of-the-art algorithms [11, 12], denoted as 'Interactive ParEGO' and 'WAPE' respectively. We used a value of $\omega = 2$ in Eq. 3, and NSGA-II with a population size of 100. We use an RBF kernel with Automatic Relevance Determination (ARD) [16] for the GPs. This property enhances the flexibility of the Gaussian Process by independently selecting lengthscales for different input dimensions. The initial points for training the GP are sampled using a Halton quasi-random sequence. The number of these points is determined by $X_{init} = \frac{(n+1)(n+2)}{2}$, where n represents the input dimension [13].

The boxplots in Figure 1 depict the final OC values. Each experiment was repeated 20 times to ensure statistical validity. It is noteworthy that each experiment adhered to a budget of 100 BO iterations, in addition to the initial sampling size X_{init}. Importantly, each interaction with the DM is considered a BO iteration.

Based on the final OC values, the newly proposed algorithm consistently succeeds in accurately identifying the preferences of the DM; in particular, the low variability observed is remarkable. The algorithm's effective accuracy indicates its capability to manage both exploration and uncertainty introduced by the DM. By introducing uncertainty to the decisions of the DM we would expect slower convergence and potentially worse performance. Figure 2 shows the algorithm's capacity of handling the uncertainty in the decisions of the DM. Figure 2 (top) shows the variability around the minimum OC, and Figure 2 (bottom) the convergence. These results clearly show that the relative degradation of the performance of the proposed approach is evident when the degree of uncertainty becomes exceptionally high, as we would expect in practice.

4 CONCLUSION

This study presents an algorithm for selecting preferred solutions in expensive multiobjective optimization, where the DM plays a key role in the acquisition process. Using a surrogate-based evolutionary algorithm and a well-known acquisition function in Bayesian optimization, experiments show our method to outperform existing approaches and is capable of handling decision uncertainty heuristically. The algorithm presents a set of Pareto-optimal solutions to the DM based on the lower confidence bound of the predicted means, allowing the DM to make informed choices based on predicted performance and uncertainty. Scalability in handling more than 3 objectives remains a challenge since the algorithm currently relies on the visualization of the Pareto front, highlighting a need for further research. Additionally, reducing DM interactions is a key future research direction, particularly for problems with limited engagement opportunities.

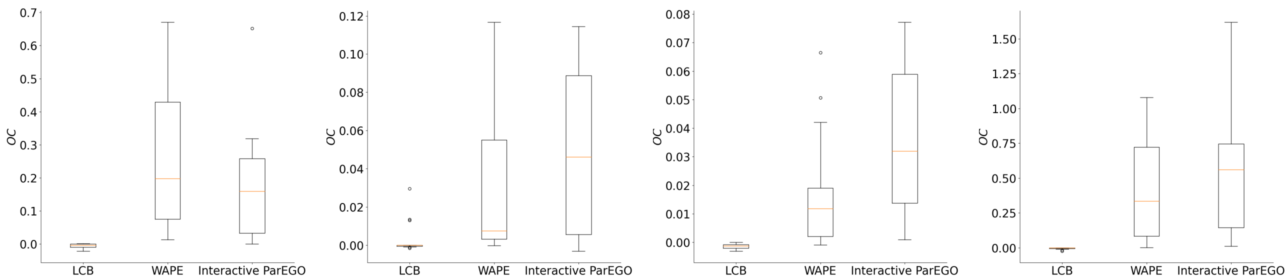

(a) Final OC Value for DEB2DK (b) Final OC Value for ZDT3 (c) Final OC Value for ZDT1 (d) Final OC Value for DTLZ7

Figure 1: Final Opportunity Cost (OC) values for the proposed algorithm and competing baselines.

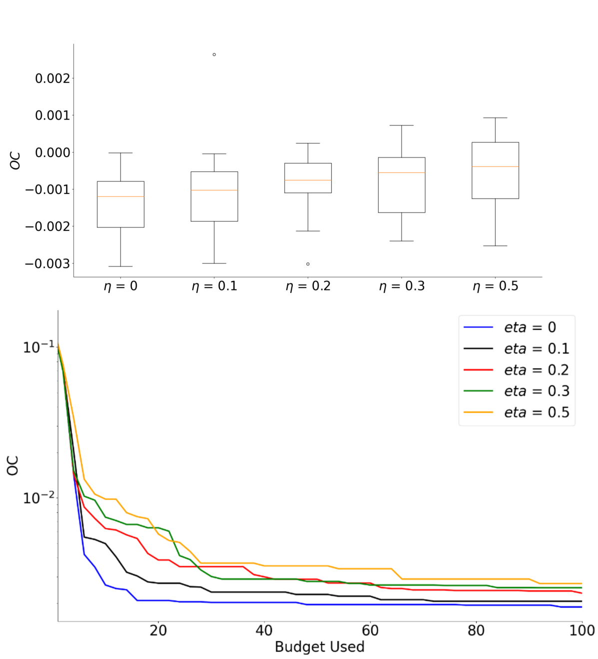

Figure 2: OC value during optimization for ZDT1 under uncertainty

ACKNOWLEDGMENTS

This work has been supported by the Flemish Government under the 'Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen' and the 'Fonds Wetenschappelijk Onderzoek (FWO)' programmes.

REFERENCES

[1] Benjamin Armbruster and Erick Delage. 2015. Decision making under uncertainty when preference information is incomplete. *Management science* 61, 1 (2015), 111–128.

[2] Raul Astudillo and Peter Frazier. 2020. Multi-attribute Bayesian optimization with interactive preference learning. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 108)*, Silvia Chiappa and Roberto Calandra (Eds.). PMLR, 4496–4507. https://proceedings.mlr.press/v108/astudillo20a.html

[3] Jürgen Branke, Kalyanmoy Deb, Henning Dierolf, and Matthias Osswald. 2004. Finding knees in multi-objective optimization. In *International conference on parallel problem solving from nature*. Springer, 722–731.

[4] Tinkle Chugh, Karthik Sindhya, Jussi Hakanen, and Kaisa Miettinen. 2019. A survey on handling computationally expensive multiobjective optimization problems with evolutionary algorithms. *Soft Computing* 23 (2019), 3137–3166.

[5] Carlos A. Coello Coello, Gary B. Lamont, and David A. Van Veldhuizen. 2007. *Evolutionary Algorithms for Solving Multi-Objective Problems*. Springer.

[6] Kalyanmoy Deb, Lothar Thiele, Marco Laumanns, and Eckart Zitzler. 2005. *Scalable Test Problems for Evolutionary Multiobjective Optimization*. Springer London, London, 105–145. https://doi.org/10.1007/1-84628-137-7_6

[7] Ian Dewancker, Michael McCourt, and Samuel Ainsworth. 2016. Interactive Preference Learning of Utility Functions for Multi-Objective Optimization. arXiv:1612.04453 [math.OC]

[8] Michael Emmerich and André H Deutz. 2018. A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Natural computing* 17, 3 (2018), 585–609.

[9] Peter I. Frazier. 2018. A Tutorial on Bayesian Optimization. arXiv:1807.02811 [stat.ML]

[10] Finley J. Gibson, Richard M. Everson, and Jonathan E. Fieldsend. 2022. Guiding surrogate-assisted multi-objective optimisation with decision maker preferences. In *Proceedings of the Genetic and Evolutionary Computation Conference* (Boston, Massachusetts) *(GECCO '22)*. Association for Computing Machinery, New York, NY, USA, 786–795. https://doi.org/10.1145/3512290.3528814

[11] Jussi Hakanen and Joshua D. Knowles. 2017. On Using Decision Maker Preferences with ParEGO. In *Evolutionary Multi-Criterion Optimization*, Heike Trautmann, Günter Rudolph, Kathrin Klamroth, Oliver Schütze, Margaret Wiecek, Yaochu Jin, and Christian Grimme (Eds.). Springer International Publishing, Cham, 282–297.

[12] Arash Heidari, Sebastian Rojas Gonzalez, Tom Dhaene, and Ivo Couckuyt. 2024. Data-Efficient Interactive Multi-Objective Optimization Using ParEGO. arXiv:2401.06649 [cs.NE]

[13] Jack Kleijnen, Inneke Nieuwenhuyse, and Wim Beers. 2022. Constrained optimization in simulation: efficient global optimization and Karush-Kuhn-Tucker conditions. *SSRN Electronic Journal* (08 2022). https://doi.org/10.2139/ssrn.3958881

[14] Kaisa Miettinen. 1999. *Nonlinear multiobjective optimization*. Vol. 12. Springer Science & Business Media.

[15] Kaisa Miettinen, Francisco Ruiz, and Andrzej P. Wierzbicki. 2008. *Introduction to Multiobjective Optimization: Interactive Approaches*. Springer Berlin Heidelberg, Berlin, Heidelberg, 27–57.

[16] Carl Edward Rasmussen. 2004. *Gaussian Processes in Machine Learning*. Springer Berlin Heidelberg, Berlin, Heidelberg, 63–71. https://doi.org/10.1007/978-3-540-28650-9_4

[17] Sebastian Rojas-Gonzalez and Inneke Van Nieuwenhuyse. 2020. A survey on kriging-based infill algorithms for multiobjective simulation optimization. *Computers & Operations Research* 116 (2020), 104869.

[18] Juan Ungredda, Juergen Branke, Mariapia Marchi, and Teresa Montrone. 2022. Single Interaction Multi-Objective Bayesian Optimization. In *Parallel Problem Solving from Nature – PPSN XVII*, Günter Rudolph, Anna V. Kononova, Hernán Aguirre, Pascal Kerschke, Gabriela Ochoa, and Tea Tušar (Eds.). Springer International Publishing, Cham, 132–145.

[19] Eckart Zitzler, Kalyanmoy Deb, and Lothar Thiele. 2000. Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. *Evolutionary Computation* 8, 2 (2000), 173–195. https://doi.org/10.1162/106365600568202