

# User-Preference Meets Pareto-Optimality: Multi-Objective Bayesian Optimization with Local Gradient Search

Ip, Joshua Hang Sai; Chakrabarty, Ankush; Mesbah, Ali; Romeres, Diego

TR2025-018 March 01, 2025

## Abstract

Incorporating user preferences into multi-objective Bayesian optimization (MOBO) allows for personalization of the optimization procedure. Preferences are often abstracted in the form of an unknown utility function, estimated through pair-wise comparisons of potential outcomes. However, utility-driven MOBO methods can yield solutions that are dominated by nearby solutions, as non-dominance is not enforced. Additionally, classical MOBO commonly relies on estimating the entire Pareto-front to identify the Pareto-optimal solutions, which can be expensive and ignore user preferences. Here, we present a new method, termed preference-utility-balanced MOBO (PUB-MOBO), that allows users to disambiguate between near-Pareto candidate solutions. PUB-MOBO combines utility-based MOBO with local multi-gradient descent to refine user-preferred solutions to be near-Pareto-optimal. To this end, we propose a novel preference-dominated utility function that concurrently preserves user-preferences and dominance amongst candidate solutions. A key advantage of PUB-MOBO is that the local search is restricted to a (small) region of the Pareto-front directed by user preferences, alleviating the need to estimate the entire Pareto-front. PUB-MOBO is tested on three synthetic benchmark problems: DTLZ1, DTLZ2 and DH1, as well as on three real-world problems: Vehicle Safety, Conceptual Marine Design, and Car Side Impact. PUB-MOBO consistently outperforms state-of-the-art competitors in terms of proximity to the Pareto-front and utility regret across all the problems.

*AAAI Conference on Artificial Intelligence 2025*



# User-Preference Meets Pareto-Optimality: Multi-Objective Bayesian Optimization with Local Gradient Search

Joshua Hang Sai Ip<sup>1</sup>, Ankush Chakrabarty<sup>2</sup>, Ali Mesbah<sup>1</sup>, Diego Romeres<sup>2</sup>

<sup>1</sup>University of California, Berkeley

<sup>2</sup>Mitsubishi Electric Research Laboratories

ipjoshua@berkeley.edu, chakrabarty@merl.com, mesbah@berkeley.edu, romeres@merl.com

## Abstract

Incorporating user preferences into multi-objective Bayesian optimization (MOBO) allows for personalization of the optimization procedure. Preferences are often abstracted in the form of an unknown utility function, estimated through pairwise comparisons of potential outcomes. However, utility-driven MOBO methods can yield solutions that are dominated by nearby solutions, as non-dominance is not enforced. Additionally, classical MOBO commonly relies on estimating the entire Pareto-front to identify the Pareto-optimal solutions, which can be expensive and ignore user preferences. Here, we present a new method, termed preference-utility-balanced MOBO (PUB-MOBO), that allows users to disambiguate between near-Pareto candidate solutions. PUB-MOBO combines utility-based MOBO with local multi-gradient descent to refine user-preferred solutions to be near-Pareto-optimal. To this end, we propose a novel preference-dominated utility function that concurrently preserves user-preferences and dominance amongst candidate solutions. A key advantage of PUB-MOBO is that the local search is restricted to a (small) region of the Pareto-front directed by user preferences, alleviating the need to estimate the entire Pareto-front. PUB-MOBO is tested on three synthetic benchmark problems: DTLZ1, DTLZ2 and DH1, as well as on three real-world problems: Vehicle Safety, Conceptual Marine Design, and Car Side Impact. PUB-MOBO consistently outperforms state-of-the-art competitors in terms of proximity to the Pareto-front and utility regret across all the problems.

## 1 Introduction

The challenge of identifying optimal trade-offs between multiple complex objective functions is pervasive in many real-world scientific and industrial applications. When the objectives are black-box functions constructed from noisy observations, multi-objective Bayesian optimization (MOBO) (Konakovic Lukovic, Tian, and Matusik 2020; Daulton, Balandat, and Bakshy 2020) is an effective multi-objective optimization (MOO) approach owing to its high sample efficiency, especially as compared to classic MOO methods such as CMA-ES (Hansen and Ostermeier 2001) and NSGA-II (Deb et al. 2002). Nonetheless, state-of-the-art MOBO methods generally seek to estimate the entire Pareto-front, which can become prohibitively expensive due

to sample-based evaluation of acquisition functions such as  $q$ -Expected Hypervolume Improvement (qEHVI) (Daulton, Balandat, and Bakshy 2020).

Recently, there has been a growing interest in preference-based MOBO (e.g., (Abdolshah et al. 2019; Ahmadian-shalchi, Belakaria, and Doppa 2024; Shao et al. 2023; Ozaki et al. 2024)) that aims to incorporate user preferences into MOO for selecting optimal points. In essence, preference-based MOBO guides the optimization process towards regions of interest within the Pareto-front by leveraging user feedback, typically in the form of pairwise comparisons between solutions generated by the optimization algorithm. These comparisons are used to estimate an underlying utility function that describes user preferences. However, while preference-based MOBO can effectively identify solutions with high utility as informed by user feedback, the resulting solutions may not be Pareto-optimal.

To address this challenge, we present a new method, termed Preference-Utility-Balanced MOBO (PUB-MOBO), that systematically approaches the user-informed regions of interest within the Pareto-front by synergizing global and local search strategies. PUB-MOBO begins with a global search driven by utility maximization to identify regions in the solution space that align with user preferences. Subsequently, a local search is conducted in the vicinity of these solutions to discover solutions that are closer to Pareto-optimality. Additionally, a new utility function, the Preference-Dominated Utility Function (PDUF), is proposed that encapsulates the concept of dominance within a single function. PDUF allows for consistently identifying dominating solutions, while providing a straightforward means for expressing all possible user preferences. This differs to existing utility functions for preference-based MOBO such as the negative  $l_1$  distance from an ideal solution irrespective of the solution being on the Pareto-front or an infeasible ideal solution (Miettinen 1999), or the weighted sum where not all Pareto-optimal points can be assigned with the highest utility value from any choice of weights (Chiandussi et al. 2012). Empirical demonstrations on several synthetic benchmark and real-world problems show that PUB-MOBO not only enhances the utility of the optimization solutions, but also yields near Pareto-optimal solutions.

In sum, the main contributions of this paper include:

(C1) We introduce a new preference-based MOBO method,

PUB-MOBO, that can effectively integrate feedback about user preferences and moves towards Pareto-optimality in MOO of black-box functions.

- (C2) We propose the use of gradient descent (GD) in the context of MOBO, along with a new utility function, PDUf, that seamlessly combine user preferences with the notion of dominance to identify user-preferred solutions that are approximately Pareto-optimal.
- (C3) We illustrate the importance of reducing gradient uncertainty in GD to aid PUB-MOBO in locating high-utility near Pareto-optimal solutions. Our approach is based on the Gradient Information acquisition function (Müller, von Rohr, and Trimpe 2021).
- (C4) Through numerical experiments on synthetic and real-world benchmark problems, we demonstrate that incorporating GD into utility-based MOBO can significantly enhance performance, yielding solutions that better align user preferences and Pareto-optimality.

## 2 Related Work

Traditional MOBO methods such as  $q$ -EHVI (Daulton, Balandat, and Bakshy 2020) assume that all Pareto-optimal solutions are equally desirable to the user, which might not be the case in practice. Instead, preference-based MOBO models user preferences with utility functions. Lin et al. (2022); Astudillo and Frazier (2020) propose the EUBO and qEIU acquisition functions respectively, which take advantage of user-preference when querying new points.

Various works have proposed different methods to incorporate gradients as additional information in single-objective BO to enhance global search. Wu et al. (2017) construct a joint GP to correlate zeroth and first order information, demonstrating that gradients aid the surrogate in approximating the posterior and change which points to query. Similarly, Makrygiorgos, Paulson, and Mesbah (2023) leverage gradients to establish stationarity conditions within a Karush-Kuhn-Tucker (KKT) formulation.

Other works instead use gradients for first-order optimization, also in the single-objective BO context. Bayesian and local optimisation sample-wise switching optimisation method (BLOSSOM) switches from global search to local search with BFGS when the posterior objective is close to the objective (McLeod, Roberts, and Osborne 2018). On the other hand, (Müller, von Rohr, and Trimpe 2021; Nguyen et al. 2022) abandon global search and construct local GPs for local optimization. To our knowledge, involving gradients for preference-based MOBO has not been studied and PUB-MOBO is the first algorithm to do this.

## 3 Preliminaries

### 3.1 Problem Formulation

We tackle a multi-objective optimization (MOO) problem for *minimizing*  $n_f$  expensive-to-evaluate objective functions, denoted by  $f_i(\mathbf{x})$  for  $i \in \{1, \dots, n_f\}$ . Consequently, the objective function vector is denoted  $\mathbf{f}(\mathbf{x})$ , where  $\mathbf{x} \in \mathbb{R}^{n_x}$  denote the decision variables. We assume that for a candidate  $\mathbf{x}$ , the function  $\mathbf{f}(\mathbf{x})$  can be evaluated, but no first- or

higher-order information about any component of  $\mathbf{f}$  is available. Furthermore, an analytical form of  $\mathbf{f}$  is not known.

For MOO problems without user-preferences, the objective is to attain Pareto-optimality, which is defined as follows (Deb and Gupta 2005). A point  $\bar{\mathbf{x}}$  in a feasible set  $\mathbb{X}$  is *Pareto-optimal* if it is not dominated<sup>1</sup> by any other solution in  $\mathbb{X}$ . That is, there exists no other feasible point  $\mathbf{x} \in \mathbb{X}$  such that  $f_i(\mathbf{x}) \leq f_i(\bar{\mathbf{x}})$  for all  $i \in \{1, \dots, n_f\}$  and  $f_j(\mathbf{x}) < f_j(\bar{\mathbf{x}})$  for at least one  $j \in \{1, \dots, n_f\}$ . Note that accurately computing the set of Pareto-optimal points, referred to as the Pareto-front  $\mathbb{X}_{\text{pareto}}$ , can often be computationally prohibitive, even for small  $n_f$ .

In the presence of a user, estimating the entire Pareto-front may become unnecessary, especially when only specific sub-regions of the feasible set  $\mathbb{X}$  is of interest. Mathematically, such user-preferences are often abstracted in the MOBO literature via *utility functions*. Specifically, the MOO problem is recast as a (scalar) utility maximization problem

$$\max_{\mathbf{x} \in \mathbb{X}} u(\mathbf{f}(\mathbf{x})), \quad (1)$$

where  $u : \mathbb{R}^{n_f} \rightarrow \mathbb{R}$  is the unknown utility function that dictates the behavior of the user. Note that the input to the utility is a noise-corrupted outcome vector  $\mathbf{y} = \mathbf{f}(\mathbf{x}) + \varepsilon$ , where  $\varepsilon$  is zero-mean noise with variance  $\sigma_\varepsilon^2 \mathbf{I}_{n_y}$  where  $\mathbf{I}_{n_y}$  is the  $n_f \times n_f$  identity matrix. For the sequel, let the highest utility Pareto-point be defined as

$$\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \mathbb{X}_{\text{pareto}}} u(\mathbf{f}(\mathbf{x})). \quad (2)$$

Following the preference-based BO literature, we assume that the utility function is not available to evaluate directly, and no functional form is known. Additionally, it is well-established that user preferences are difficult to be assigned to continuous numerical values; instead we suppose that users are more inclined to provide weak supervision in the form of pairwise comparisons (Chu and Ghahramani 2005; Lin et al. 2022). The following assumption asserts that a typical user will select dominating solutions when possible.

**Assumption 1** *If  $\mathbf{y}_1$  and  $\mathbf{y}_2$  are candidate outcomes presented to the user and  $\mathbf{y}_1 \succ \mathbf{y}_2$ , then the user will always select  $\mathbf{y}_1$ ; that is,  $u(\mathbf{y}_1) > u(\mathbf{y}_2)$ .*

This assumption is a constraint that must be respected when modeling preference-based MOBO problems to accurately reflect real user behavior.

### 3.2 Modeling with Gaussian processes

We first discuss the modeling choices considered to learn the outcome function  $\mathbf{f}$  and the utility function  $u$ : their respective approximations are denoted  $\hat{\mathbf{f}}$  and  $\hat{u}$ .

**Modeling outcomes.** Gaussian process (GP) regression is a popular choice for constructing the surrogate  $\hat{\mathbf{f}}$  for the true outcome function  $\mathbf{f}$ . We train an independent GP for each

<sup>1</sup>Generally, a candidate  $\mathbf{x}_1$  *dominates* a candidate  $\mathbf{x}_2$ , denoted by  $\mathbf{x}_1 \succ \mathbf{x}_2$ , if  $f_i(\mathbf{x}_1) \leq f_i(\mathbf{x}_2)$  for all  $i = 1, \dots, n_f$  and for one such element,  $f_i(\mathbf{x}_1) < f_i(\mathbf{x}_2)$ .

objective  $f_i$ , though a multi-output GP that models correlations between the objectives could also be considered (Alvarez et al. 2012). Each GP is defined *a priori* by a mean function  $m(\mathbf{x})$  and covariance function  $k_i(x, x')$  called kernel. For this work, any  $\mathcal{C}^2$  kernel is admissible.

Let  $\mathbf{X}_T = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$ ; we drop the subscript for brevity. Given a dataset  $D := (\mathbf{X}, \mathbf{Y})$ , comprising input-outcome pairs, the mean and variance of the posterior are:

$$\mu_i(\mathbf{x}) = m(\mathbf{x}) + k_i(\mathbf{x}, \mathbf{X})\mathcal{K}_\sigma^{-1}(\mathbf{X})(Y_i - m(\mathbf{X})), \quad (3a)$$

$$\Sigma_i(\mathbf{x}) = k_i(\mathbf{x}, \mathbf{x}) - k_i(\mathbf{x}, \mathbf{X})\mathcal{K}_\sigma^{-1}(\mathbf{X})k_i(\mathbf{X}, \mathbf{x}), \quad (3b)$$

where  $\mathcal{K}_\sigma(\mathbf{X}) := K(\mathbf{X}, \mathbf{X}) + \sigma^2 \mathbf{I}$  and  $m(\cdot)$  is the prior mean. Since the derivative is a linear operator, the derivative GP is another GP (Williams and Rasmussen 2006) characterized fully by the mean and covariance functions

$$\mu_i^\nabla(\mathbf{x}) = \nabla m(\mathbf{x}) + \nabla k_i(\mathbf{x}, \mathbf{X})\mathcal{K}_\sigma^{-1}(\mathbf{X})(Y_i - m(\mathbf{X})), \quad (4a)$$

$$\Sigma_i^\nabla(\mathbf{x}) = \nabla^2 k_i(\mathbf{x}, \mathbf{x}) - \nabla k_i(\mathbf{x}, \mathbf{X})\mathcal{K}_\sigma^{-1}(\mathbf{X})\nabla k_i(\mathbf{X}, \mathbf{x}), \quad (4b)$$

**Modeling preferences.** We assume the user is only capable of weak supervisions in the form of pairwise comparisons (PC). That is, if the user prefers  $\mathbf{y} := \mathbf{f}$  over  $\mathbf{y}' := \mathbf{f}'$ , the pairwise comparison function  $r(\mathbf{y}, \mathbf{y}') = 0$ . In the event that the user prefers  $\mathbf{y}'$  instead,  $r(\mathbf{y}, \mathbf{y}') = 1$ . Pairwise GPs c.f. (Chu and Ghahramani 2005) allow us to learn a latent functional representation  $\hat{u}$  of the true user utility based on this preference feedback. The latent function satisfies  $\hat{u}(\mathbf{y}) > \hat{u}(\mathbf{y}')$  if the user prefers  $\mathbf{y}$ , and vice versa.

## 4 Preference-Utility-Balanced (PUB) MOBO

This work is motivated by the practical consideration that most users expect to be shown promising candidate outcomes after very few interactions. Furthermore, they often require some assurance that the suggested candidates are not only high in utility, but also high in performance, i.e., close to the Pareto-front. To this end, we propose PUB-MOBO, which relies on utility maximization to ascertain candidate solutions that are preferred by the user, while promoting a local search towards the Pareto-front using estimated gradients. Empirically, we observe that the local search finds solutions near Pareto points, which subsequently accelerates the search for high-utility solutions. The rationale behind this is that the local search likely yields dominating solutions over those found by utility maximization alone. These dominating solutions, by Assumption 1, have higher utility. This implies that even with few user interactions, i.e. pairwise comparisons, we can obtain promising candidates that are unlikely to be dominated. As more user feedback is collected, we expect to obtain near-maximal utility solutions due to the utility maximization, from which following local gradients should result in a near-Pareto solution.

### 4.1 PUB-MOBO Algorithm

The proposed PUB-MOBO method operates in three stages. Following the work of Lin et al. (2022), the first two

---

### Algorithm 1: PUB-MOBO

---

- 1: Generate initial data:  $\mathbf{x}_{\text{INIT}}, \mathbf{y}_{\text{INIT}}, r(\mathbf{y}_{\text{INIT}})$
  - 2:  $D = (\mathbf{x}_{\text{INIT}}, \mathbf{y}_{\text{INIT}})$
  - 3:  $P \leftarrow$  comparisons on a subset of  $\mathbf{y}_{\text{INIT}} \times \mathbf{y}_{\text{INIT}}$
  - 4: Update outcome model  $\hat{\mathbf{f}}$  with  $D$
  - 5: Update preference model  $\hat{u}$  with  $P$
  - 6: **while** # outcome evaluations  $\leq$  budget **do**
  - 7:   **PE stage**
  - 8:    $\mathbf{x}_1, \mathbf{x}_2 \leftarrow \text{argmax}_{\mathbf{x}_1, \mathbf{x}_2}$  EUBO
  - 9:    $\mathbf{y}_1, \mathbf{y}_2 = \hat{\mathbf{f}}(\mathbf{x}_1), \hat{\mathbf{f}}(\mathbf{x}_2)$
  - 10:    $r(\mathbf{y}_1, \mathbf{y}_2) \leftarrow$  user provides a comparison
  - 11:   Append  $P$  with  $(\mathbf{y}_1, \mathbf{y}_2, r(\mathbf{y}_1, \mathbf{y}_2))$
  - 12:   Update pref. model  $\hat{u}$  with  $(\mathbf{y}_1, \mathbf{y}_2, r(\mathbf{y}_1, \mathbf{y}_2))$
  - 13:   **EXP stage**
  - 14:    $\mathbf{x}_{\text{EXP}} \leftarrow \text{argmax}_x$  qEIUU
  - 15:    $\mathbf{y}_{\text{EXP}} = \mathbf{f}(\mathbf{x}_{\text{EXP}})$
  - 16:   Append  $D$  with  $(\mathbf{x}_{\text{EXP}}, \mathbf{y}_{\text{EXP}})$
  - 17:   Update outcome model  $\hat{\mathbf{f}}$  with  $(\mathbf{x}_{\text{EXP}}, \mathbf{y}_{\text{EXP}})$
  - 18:   **GD stage**
  - 19:    $(\mathbf{X}_{\text{GD}}, \mathbf{Y}_{\text{GD}}) \leftarrow$  Local Gradient Descent( $\mathbf{x}_{\text{EXP}}$ )
  - 20:   Append  $D$  with  $(\mathbf{X}_{\text{GD}}, \mathbf{Y}_{\text{GD}})$
  - 21: **end while**
- 

stages of the framework are the preference exploration (PE) stage and the outcome evaluation via experimentation (EXP) stage. Our method extends these two stages with an additional stage based on local multi-gradient descent, denominated the *GD stage*.

In each PUB-MOBO iteration, these three stages are executed, and the process is repeated *ad infinitum*, or (more practically) until a pre-decided budget for total number of outcome/function evaluations is attained; see Algorithm 1.

**Preference Exploration.** In the PE stage, the user expresses their preferences over a query of two candidate solutions in a form of pairwise comparisons. The comparison is then used to update the estimate of the utility function  $\hat{u}$ . The candidate solutions in the query are obtained by optimizing the expected utility of both outcomes acquisition function, which is a look-ahead function that maximizes the difference in expected utility after a user query  $r(\hat{\mathbf{f}}(\mathbf{x}_1), \hat{\mathbf{f}}(\mathbf{x}_2))$ . Since this is expensive to compute, a simplified acquisition called EUBO, proposed in (Lin et al. 2022), is used instead:

$$\text{EUBO}(\mathbf{x}_1, \mathbf{x}_2) = \mathbb{E}[\max(\hat{u}(\hat{\mathbf{f}}(\mathbf{x}_1)), \hat{u}(\hat{\mathbf{f}}(\mathbf{x}_2)))]. \quad (5)$$

This has the advantage of estimating utility with outcome posteriors that are feasible under  $\mathbf{f}$ . Note that no evaluation of  $\mathbf{f}$  is required in this stage.

**Outcome evaluation via Experiments.** The EXP stage involves computing the optimal decision variables and collecting outcome evaluations to update the outcome model  $\hat{\mathbf{f}}$ . In this process, the well-established expected improvement under utility uncertainty (qEIUU) acquisition function, (Astudillo and Frazier 2020), widely used in preference-based contexts, is maximized to determine the optimal decision

variables,  $\mathbf{x}_{\text{EXP}}$ . Here,

$$\text{qEIUU}(\mathbf{x}) = \mathbb{E} \left[ \max(\hat{u}(\hat{\mathbf{f}}(\mathbf{x})) - \hat{u}(\mathbf{f}(\mathbf{x}_{\text{best}})), 0) \right], \quad (6)$$

where  $\mathbf{x}_{\text{best}} = \arg \max_{\mathbf{x} \in \mathcal{X}} \hat{u}(\hat{\mathbf{f}}(\mathbf{x}))$ , and

$$\mathbf{x}_{\text{EXP}} := \arg \max_{\mathbf{x}} \text{qEIUU}(\mathbf{x}).$$

Since the expectation in (6) is with respect to the outcome and utility models, the analytical expression is challenging to compute. Instead, this can be evaluated via the Monte Carlo approach, where the reparameterization trick is applied to both  $\hat{\mathbf{f}}$  and  $\hat{u}$  and the acquisition function is optimized using sample-averaging; c.f. Wilson, Hutter, and Deisenroth (2018); Balandat et al. (2020). Note that after  $\mathbf{x}_{\text{EXP}}$  is obtained, we append it along with its true outcome value  $\mathbf{f}(\mathbf{x}_{\text{EXP}})$  to the current dataset  $D$ .

**Local gradient descent.** This GD stage is motivated by the fact that  $\mathbf{x}_{\text{EXP}}$ , while expected to be high in utility, is not specifically designed to be near the Pareto-front. Analogous to single-objective optimization, we will pursue local gradients that are expected to generate a trajectory of  $\mathbf{x}$  candidates that evolves towards a nearby Pareto-optimal point. We will refer to these gradient-following decision variables as ‘ $\mathbf{x}_{\text{GD}}$ ’. We set the initial  $\mathbf{x}_{\text{GD}}$  to be  $\mathbf{x}_{\text{EXP}}$ .

Clearly, for a MOO problem, gradient descent must be adapted for multiple objectives. We propose the use of multiple gradient descent algorithm (MGDA) (Désidéri 2012), which was designed for smooth multi-outcome objective functions. While MGDA is provably convergent for white-box optimization problem settings i.e. when gradients of each outcome of  $\mathbf{f}$  is accessible, it has not been tested in the context of black-box MOO such as in this paper, where gradients are not accessible, and must be estimated. However, MGDA exhibits some theoretical properties that, we hypothesize, and demonstrate via experiments, are beneficial in the MOBO context. MGDA exploits the KKT conditions (Fliege and Svaiter 2000; Schäffler, Schultz, and Weinzierl 2002)

$$\boldsymbol{\alpha} \geq \mathbf{0}, \mathbf{1}^\top \boldsymbol{\alpha} = 1, \boldsymbol{\alpha}^\top \nabla \mathbf{f}(\mathbf{x}) = 0, \quad (7)$$

and recasts this for MOO as a quadratic cost constrained on the probability simplex, that is:

$$\min_{\boldsymbol{\alpha} \geq \mathbf{0}} \|\boldsymbol{\alpha}^\top \nabla \mathbf{f}(\mathbf{x})\|^2 \text{ subject to: } \mathbf{1}^\top \boldsymbol{\alpha} = 1. \quad (8)$$

It is well-known, c.f. Désidéri (2012), that a solution to (8) is either:  $\boldsymbol{\alpha}^\top \nabla \mathbf{f}(\mathbf{x}) = 0$ , in which case the current parameters  $\mathbf{x}$  are Pareto-optimal, or  $\boldsymbol{\alpha}^\top \nabla \mathbf{f}(\mathbf{x}) \neq 0$ , and  $\boldsymbol{\alpha}^\top \nabla \mathbf{f}(\mathbf{x})$  is a feasible descent direction. Given that (8) is a quadratic cost over linear constraints, we can use the Frank-Wolfe algorithm (Sener and Koltun 2018; Jaggi 2013) to efficiently compute optimal solutions.

Solving (8) yields an optimal  $\boldsymbol{\alpha}$  with which we can take a gradient step  $\mathbf{x}_{\text{GD}} \leftarrow \mathbf{x}_{\text{GD}} - \eta \boldsymbol{\alpha}^\top \nabla \mathbf{f}(\mathbf{x}_{\text{GD}})$ . However, there are two clear difficulties at this juncture. First, this update may yield an  $\mathbf{x}_{\text{GD}} \notin \mathbb{X}$ . To counter this, we stop updating when this happens, and stop the local gradient search phase, moving on to the next PUB-MOBO iterations with an updated dataset  $D$  that contains all the  $\mathbf{x}_{\text{GD}}$  and correspond  $\mathbf{y}_{\text{GD}}$  observed so far.

---

## Algorithm 2: LOCAL GRADIENT DESCENT

---

```

1: Initialize  $\mathbf{x}_{\text{GD}} \leftarrow \mathbf{x}_{\text{EXP}}$ 
2:  $(\mathbf{X}_{\text{GD}}, \mathbf{Y}_{\text{GD}}) = (\emptyset, \emptyset)$ 
3: # of multi-gradient steps,  $n_{\text{GD}}$ 
4: # of GI optimizations,  $n_{\text{GI}}$ 
5: Early stopping threshold,  $\varepsilon_{\text{GD}}$ 
6: for  $i \leq n_{\text{GD}}$  do
7:   Compute  $\boldsymbol{\mu}^\nabla(\mathbf{x}_{\text{GD}})$  using (4a)
8:   Compute  $\mathbf{M} = \boldsymbol{\mu}^\nabla(\mathbf{x}_{\text{GD}})^\top \boldsymbol{\mu}^\nabla(\mathbf{x}_{\text{GD}})$ 
9:    $\boldsymbol{\alpha} \leftarrow \text{Frank-Wolfe}(\mathbf{M})$ 
10:   $\mathbf{x}_{\text{GD}} \leftarrow \mathbf{x}_{\text{GD}} - \eta \boldsymbol{\alpha}^\top \boldsymbol{\mu}^\nabla(\mathbf{x}_{\text{GD}})$ 
11:  if  $\mathbf{x}_{\text{GD}} \in \mathbb{X}$  and  $\|\boldsymbol{\alpha}^\top \boldsymbol{\mu}^\nabla(\mathbf{x}_{\text{GD}})\|_2^2 > \varepsilon_{\text{GD}}$  then
12:    Evaluate the true objective:  $\mathbf{y}_{\text{GD}} = \mathbf{f}(\mathbf{x}_{\text{GD}})$ 
13:    Append  $(\mathbf{X}_{\text{GD}}, \mathbf{Y}_{\text{GD}})$  with  $(\mathbf{x}_{\text{GD}}, \mathbf{y}_{\text{GD}})$ 
14:    Update outcome model  $\hat{\mathbf{f}}$  with  $(\mathbf{x}_{\text{GD}}, \mathbf{y}_{\text{GD}})$ 
15:    for  $j \leq n_{\text{GI}}$  do
16:       $\mathbf{x}_{\text{GI}} \leftarrow \arg \max_{\mathbf{x}'} \text{GI}$ 
17:      Evaluate the true objective:  $\mathbf{y}_{\text{GI}} = \mathbf{f}(\mathbf{x}_{\text{GI}})$ 
18:      Append  $(\mathbf{X}_{\text{GD}}, \mathbf{Y}_{\text{GD}})$  with  $(\mathbf{x}_{\text{GI}}, \mathbf{y}_{\text{GI}})$ 
19:      Update outcome model  $\hat{\mathbf{f}}$  with  $(\mathbf{x}_{\text{GI}}, \mathbf{y}_{\text{GI}})$ 
20:    end for
21:  else
22:    break
23:  end if
24: end for
25: return  $(\mathbf{X}_{\text{GD}}, \mathbf{Y}_{\text{GD}})$ 

```

---

The second and more debilitating problem is that we do not have access to gradients of  $\mathbf{f}$ . Thankfully, we do have a surrogate model  $\hat{\mathbf{f}}$  with which we can obtain an estimate of the gradient at any  $\mathbf{x}$  with  $\boldsymbol{\mu}^\nabla := \mathbb{E}[\nabla \hat{\mathbf{f}}(\mathbf{x})]$  through (4a). The gradient step is then

$$\mathbf{x}_{\text{GD}} \leftarrow \mathbf{x}_{\text{GD}} - \eta \boldsymbol{\alpha}^\top \boldsymbol{\mu}^\nabla(\mathbf{x}_{\text{GD}}). \quad (9)$$

Unfortunately, there is no clear correlation between the uncertainties in  $\mathbf{f}$  and  $\nabla \mathbf{f}$ , so  $\boldsymbol{\mu}^\nabla$  could have large uncertainties even near previously observed points. Therefore, it is imperative to incorporate techniques that can reduce uncertainty in the posterior of the gradient estimate. To this end, we propose to use the gradient information (GI) acquisition function, described in Müller, von Rohr, and Trimpe (2021).

We explain briefly the mechanism of the GI acquisition. Suppose we select the best candidate from the EXP stage,  $\mathbf{x}_{\text{EXP}}$ , and set it as the initial candidate for the local gradient search:  $\mathbf{x}_{\text{GD}}$ . The GI acquisition tries to select a subsequent point  $\mathbf{x}'$  that will minimize the uncertainty of the gradient at  $\mathbf{x}_{\text{GD}}$  if  $\mathbf{x}'$  and its corresponding  $\mathbf{y}'$  were known.

By considering all  $n_f$  objective independently distributed, we can formulate the uncertainty information contained in the covariance matrix  $\boldsymbol{\Sigma}^\nabla$  in (4b) using A-optimal design (Chakrabarty, Buzzard, and Rundell 2013), and maximize:

$$\sum_{i=1}^{n_f} \mathbb{E} [\text{Tr}(\boldsymbol{\Sigma}_i^\nabla(\mathbf{x}_{\text{GD}}|D)) - \text{Tr}(\boldsymbol{\Sigma}_i^\nabla(\mathbf{x}_{\text{GD}}|D, (\mathbf{x}', \mathbf{y}')))], \quad (10)$$

which, for Gaussian distributions, is equivalent to

$$\text{GI}(\mathbf{x}') = \sum_{i=1}^{n_f} \text{Tr}(\nabla k_i(\mathbf{x}_{\text{GD}}, \mathbf{X}') \mathcal{K}_\sigma^{-1}(\mathbf{X}') \nabla k_i^\top(\mathbf{x}_{\text{GD}}, \mathbf{X}')), \quad (11)$$

where  $\mathbf{X}' = \{\mathbf{X} \cup \mathbf{x}'\}$ . For each gradient-step in  $n_{\text{GD}}$ , the GI acquisition function is optimized  $n_{\text{GI}}$  times to reduce gradient uncertainty. Upon each optimization, we evaluate the outcome function to obtain a corresponding  $\mathbf{y}_{\text{GD}}$ , which is appended to the dataset  $D$  for subsequent PUB-MOBO iterations.

## 5 Preference-Dominated Utility Function

The utility function represents the user preference in preference-based MOBO algorithms and is used to simulate user responses. Practically, the utility function is employed to respond to user queries, such as providing pairwise comparisons between two outcomes (Lin et al. 2022). A utility function used to test preference-based MOBO algorithms should satisfy two key properties:

- (P1) *Dominance Preservation*: When evaluating a query, the true utility function should satisfy Assumption 1.
- (P2) *Preference Integration*: The utility function should have parameters  $\theta_u$  that allow unique strictly maximal-utility Pareto-optimal solutions. That is, for any  $\mathbf{x} \in \mathbb{X}_{\text{pareto}}$ , there exists an easily computable  $\theta_u \in \mathbb{R}^{n_u}$  such that  $u(\mathbf{f}(\mathbf{x})|\theta_u) > u(\mathbf{f}(\{\mathbb{X}_{\text{pareto}} \setminus \mathbf{x}\}|\theta_u))$ .

For instance, the commonly used  $\ell_1$  distance (a) fails to satisfy the *Preference Integration* property when calculated from the utopia point, and violates *Dominance Preservation* when calculated from any other point. This is illustrated in Fig. 1a, where the contours of an  $\ell_1$  distance utility function is shown with an example Pareto-front. Here, the two red points are indistinguishable according to the utility function, demonstrating the limitations of  $\ell_1$  distance in distinguishing between Pareto-optimal solutions.

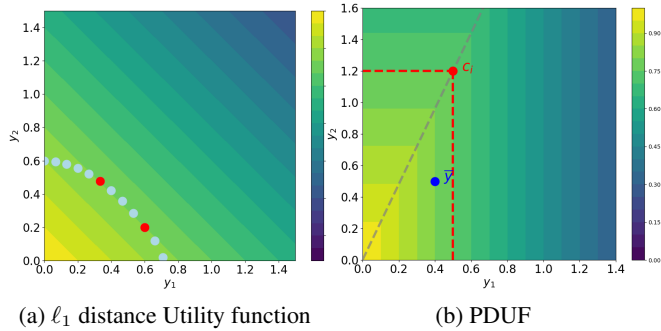


Figure 1: Contour plots of (a) the commonly used negative  $\ell_1$  distance Utility function (b) the proposed PDUF.

Therefore, we propose the preference-dominated utility function (PDUF) which merges the concept of dominance with user preferences. An illustration of the contours in a 2D case is shown in Fig. 1b. The PDUF integrates the concept of dominance with user preferences by combining multiple

logistic functions centered around different points in the objective function space. The PDUF is defined as:

$$u(\mathbf{y}) = \frac{1}{n_c} \sum_{i=1}^{n_c} \prod_{j=1}^{n_y} L_\beta(y_j, c_{i,j}), \quad (12)$$

$$L_\beta(y_j, c_{i,j}) = \frac{1}{1 + \exp(\beta \cdot (y_j - c_{i,j}))}, \quad (13)$$

where  $c_i = (c_{i,1}, c_{i,2}, \dots, c_{i,n_y})$  denotes the  $i^{\text{th}}$  center for one logistic function,  $\beta$  denotes a parameter that controls the steepness of the logistic function, and  $n_c$  denotes the number of centers. The logistic function  $L_\beta(y_j, c_{i,j})$  approximates the step function and enforces dominance for each objective  $y_j$ , as seen in the red dashed lines in Fig. 1b, and the product aggregates this approximation for all objectives. Furthermore, the sum of logistic function products preserve dominance in the objective space. Indeed, for every  $\bar{\mathbf{y}}$  that dominates user query  $c_i$ , PDUF will express user preference with  $u(\bar{\mathbf{y}}) > u(c_i)$ . Finally, the centers define the parameters  $\theta_u$  that ensure the utility function adheres to the *preference integration* property by aligning them along an arbitrary line (the grey line in Fig. 1b).

## 6 Experiments

We empirically validate the proposed PUB-MOBO algorithm on 6 benchmark problems and report the performance of utility regret,  $R$ , and distance to Pareto-front,  $d_{\text{Pareto}}$ , against outcome evaluations and user queries

$$R = \frac{u(\mathbf{f}(\mathbf{x}^*)) - u(\mathbf{f}(\mathbf{x}))}{u(\mathbf{f}(\mathbf{x}^*))}, \quad (14)$$

$$d_{\text{Pareto}} = \min_{\mathbf{x}_{\text{Pareto}} \in \mathbb{X}_{\text{Pareto}}} \|\mathbf{f}(\mathbf{x}) - \mathbf{f}(\mathbf{x}_{\text{Pareto}})\|_2^2. \quad (15)$$

The problems are chosen from synthetic and real-world problems where the Pareto-front is known, so we can evaluate PUB-MOBO in the performance metrics  $R$  and  $d_{\text{Pareto}}$ . All experimental results are obtained using a 13th Gen Intel-Core i7-13620H repeated across 20 seeds with hyperparameters  $n_{\text{GD}} = 10$ ,  $n_{\text{GI}} = 1$ ,  $\varepsilon = 0.1$ . In the experiments, we investigate two aspects of PUB-MOBO: (a) the relevance of using gradient information and (b) the trade-off between cost of outcome evaluations and gradient uncertainty reduction in the GD stage. We compare against the s.o.t.a. preference-based MOBO method (Lin et al. 2022) and two variations of PUB-MOBO:

- EUBO+qEIUU is a 2-stage algorithm proposed in (Lin et al. 2022) which only contains the PE and EXP stages. The acquisition functions used are EUBO and qEIUU in each stage, exactly like in PUB-MOBO. This serves as a baseline of the performance when no gradients are used.
- PUB-MOBO-PG is a PUB-MOBO ablation that relies solely on predicted gradients (PG) in the GD stage, omitting both outcome evaluations and GI optimizations. This makes it relatively inexpensive, but it ignores the fact that additional samples can yield useful derivative information. This ascertains whether the gradient uncertainty needs to be accounted for at all, and how important the GI step is.

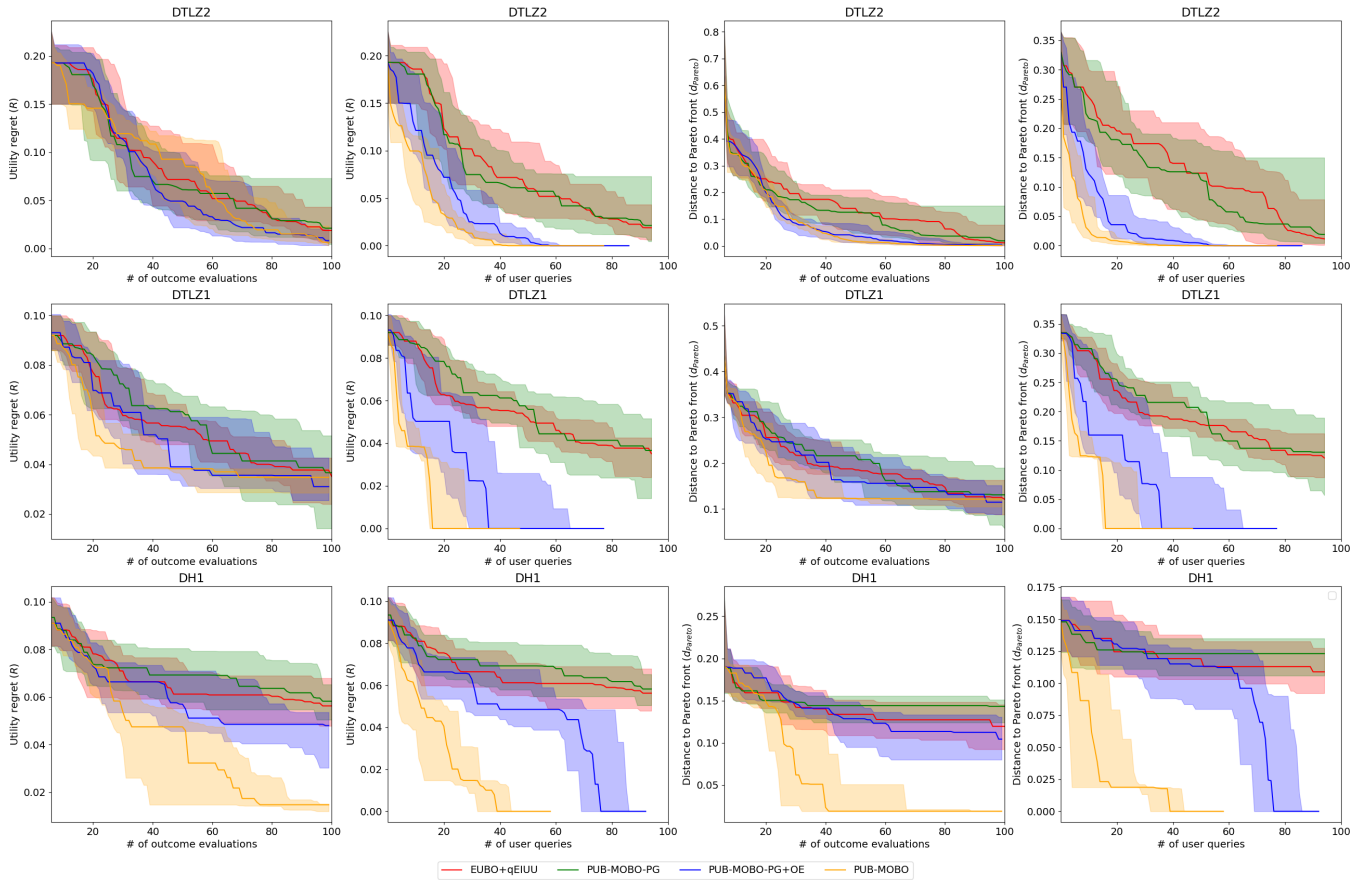


Figure 2: Median and 25-75 percentiles performance comparison on synthetic benchmarks on 20 random seeds with a 100 outcome evaluation budget. Plot titles indicate MOO benchmark; horizontal axis is either number of outcome evaluations or number of user queries. Vertical axis indicates simple regret of utility (user satisfaction) and distance from the Pareto-front (optimality).

- PUB-MOBO-PG+OE is a PUB-MOBO ablation that uses the predicted gradients as in PUB-MOBO-PG, but an Outcome Evaluation (OE) is performed at every gradient descent step in an effort to lower gradient uncertainty around observed points. This further checks whether the expensive GI optimization is needed to improve convergence, or if outcome evaluations are sufficient.
- PUB-MOBO is the proposed method from Alg.1-2 with up to  $n_{GD}(n_{GI} + 1)$  outcome evaluations in the GD stage.

### 6.1 Synthetic benchmarks

We examine 3 synthetic problems that are commonly found in MOO literature: DTLZ1, DTLZ2 (Deb et al. 2005), and DH1 (Deb and Gupta 2005). The results are displayed in Fig. 2 in order of increasing  $n_x$ : DTLZ2 ( $n_x = 8, n_f = 2$ ), DTLZ1 ( $n_x = 9, n_f = 2$ ), DH1 ( $n_x = 10, n_f = 2$ ).

EUBO+qEIIU is the poorest-performing algorithm across all the synthetic benchmarks, affirming the effectiveness of the additional stage based on local gradient search. The strategy of exploring locally dominating solutions off of the solution proposed by the utility maximization stages, not only results in solutions that are closer to

the Pareto-front but also achieves lower utility regret. However, PUB-MOBO-PG performs equally poor in most of the synthetic benchmarks, largely due to inaccurate gradient estimation obtained with the surrogate model  $\hat{f}$  by (4b). We frequently observe that the evolution of  $\mathbf{x}_{GD}$  in the GD stage is prematurely terminated either due to constraint violations in  $\mathbf{x}$  or because the Frank-Wolfe algorithm results in  $\|\alpha^\top \mu^\nabla(\mathbf{x}_{GD})\|_2^2 \leq \varepsilon_{GD}$ . This occurs not because the algorithm has reached a Pareto-optimal point, but rather due to erroneous gradient estimates. These findings underscore the critical importance of accurate gradient estimation and minimizing gradient uncertainty for the success of PUB-MOBO. The PG+OE variant significantly outperforms the PG variant in most experiments. Although evaluating the true outcome function at each  $\mathbf{x}_{GD}$  incurs additional computational cost, the increased accuracy in gradient estimation from updating the outcome model more than justifies the expense. This leads to a more efficient algorithm overall, as reflected by lower utility regret and closer proximity to the Pareto-front, both in terms of outcome evaluations and user queries. In a similar trend, PUB-MOBO further improves upon PG+OE



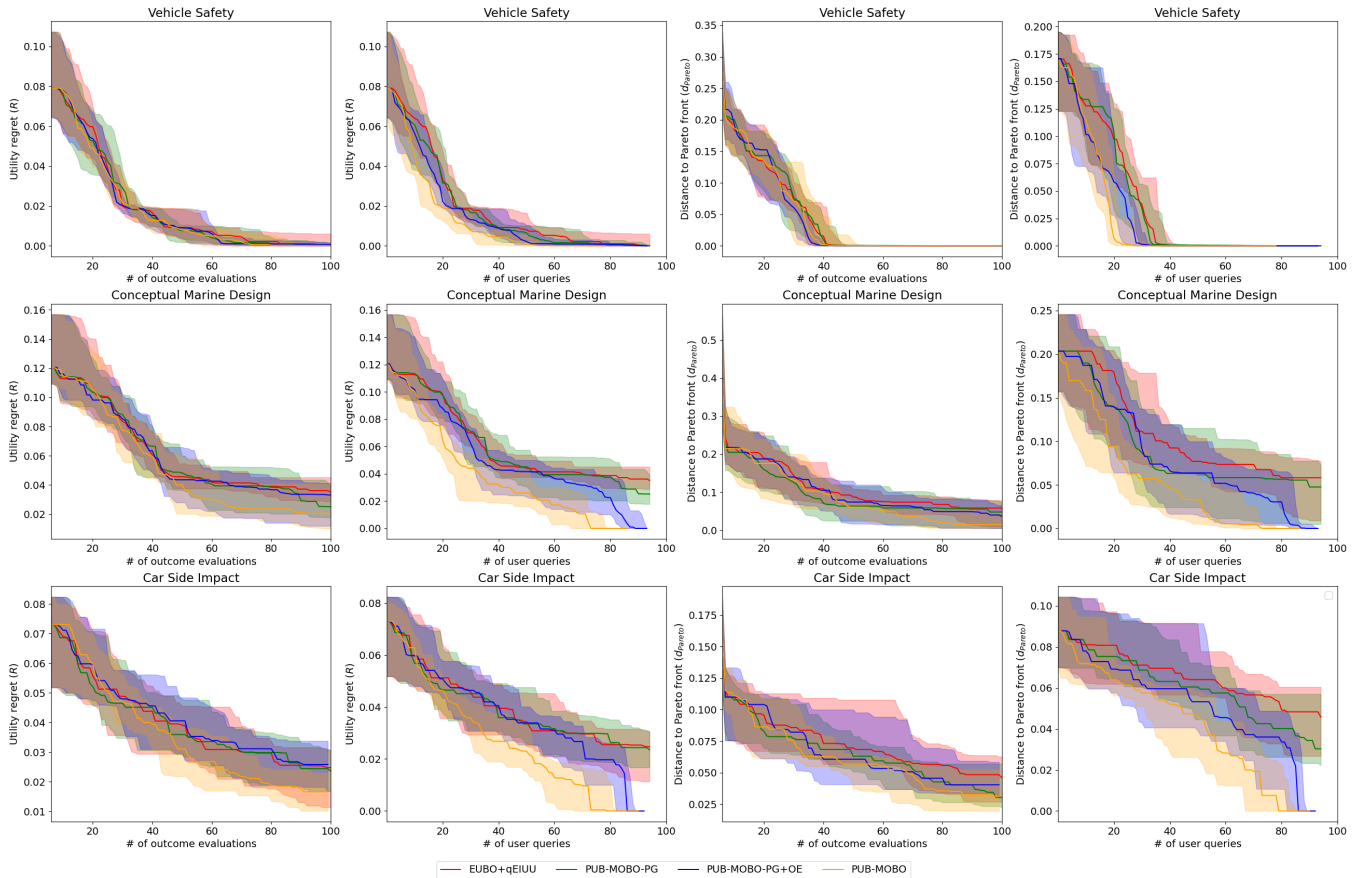


Figure 3: Median and 25-75 percentiles performance comparison on real-world benchmarks. Analogous description to Fig. 2.

variant. The reduction of the gradient uncertainty obtained by evaluating the  $x_{GD}$  suggested by the GI acquisition function, further improve the gradient estimate yielding lower  $R$  and  $d_{Pareto}$ . This empirically demonstrates that the additional outcome evaluations in the GD stage are justified.

## 6.2 Real-world benchmarks

We examine 3 problems based on real MOO problems: Vehicle Safety (Liao et al. 2008), Conceptual Marine Design (Parsons and Scott 2004), and Car Side Impact (Jain and Deb 2013). The implementations of these problems are taken from (Tanabe and Ishibuchi 2020). The results are shown in Fig. 3 in order of increasing  $n_x$ : Vehicle Safety ( $n_x = 5$ ,  $n_f = 3$ ), Conceptual Marine Design ( $n_x = 6$ ,  $n_f = 4$ ), Car Side Impact ( $n_x = 7$ ,  $n_f = 4$ ).

Similar to the results on the synthetic benchmarks, the EUBO and PUB-MOBO-PG methods are the worst-performing. The real-world benchmarks present a more complex optimization landscape, which increases the difficulty of accurate gradient estimation. Notably, PUB-MOBO-PG+OE, which performed better on synthetic problems, exhibits similarly poor performance to EUBO and PUB-MOBO-PG on these more challenging benchmarks. In contrast, PUB-MOBO continues to outperform all other methods, consistently delivering near-optimal solutions with

respect to both utility regret and proximity to the Pareto-front, due to its effective gradient uncertainty minimization.

The results from both synthetic and real-world experiments demonstrate that incorporating a gradient descent stage in utility-based MOBO leads to solutions with lower utility regret and closer proximity to the Pareto-front, all while requiring fewer outcome evaluations and user queries. Empirical evidence suggests that it is more effective to conduct additional outcome evaluations during each GD stage to achieve accurate gradient steps, as seen in PUB-MOBO, rather than opting for cheaper but less accurate steps.

## 7 Conclusion

We introduce PUB-MOBO, a sample-efficient multi-objective Bayesian optimization method that integrates user preferences with gradient-based search to find near-Pareto-optimal solutions. Across synthetic and real problems, it achieves high utility and reduced distance to Pareto-front solutions, highlighting the importance of reducing gradient uncertainty in the gradient-based search. Moreover, we introduce a new utility function, PDUF, that respects dominance and models different user preferences. An avenue of future research is to apply the Local Gradient Descent method to other MOBO algorithms.

## References

- Abdolshah, M.; Shilton, A.; Rana, S.; Gupta, S.; and Venkatesh, S. 2019. Multi-objective Bayesian optimisation with preferences over objectives. *Advances in neural information processing systems*, 32.
- Ahmadianshalchi, A.; Belakaria, S.; and Doppa, J. R. 2024. Preference-Aware Constrained Multi-Objective Bayesian Optimization. In *Proceedings of the 7th Joint International Conference on Data Science & Management of Data (11th ACM IKDD CODS and 29th COMAD)*, 182–191.
- Alvarez, M. A.; Rosasco, L.; Lawrence, N. D.; et al. 2012. Kernels for vector-valued functions: A review. *Foundations and Trends® in Machine Learning*, 4(3): 195–266.
- Astudillo, R.; and Frazier, P. 2020. Multi-attribute Bayesian optimization with interactive preference learning. In Chippa, S.; and Calandra, R., eds., *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, 4496–4507. PMLR.
- Balandat, M.; Karrer, B.; Jiang, D.; Daulton, S.; Letham, B.; Wilson, A. G.; and Bakshy, E. 2020. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. *Advances in neural information processing systems*, 33: 21524–21538.
- Chakrabarty, A.; Buzzard, G. T.; and Rundell, A. E. 2013. Model-based design of experiments for cellular processes. *Wiley Interdisciplinary Reviews: Systems Biology and Medicine*, 5(2): 181–203.
- Chiandussi, G.; Codegone, M.; Ferrero, S.; and Varesio, F. E. 2012. Comparison of multi-objective optimization methodologies for engineering applications. *Computers & Mathematics with Applications*, 63(5): 912–942.
- Chu, W.; and Ghahramani, Z. 2005. Preference learning with Gaussian processes. In *Proceedings of the 22nd international conference on Machine learning*, 137–144.
- Daulton, S.; Balandat, M.; and Bakshy, E. 2020. Differentiable expected hypervolume improvement for parallel multi-objective Bayesian optimization. *Advances in Neural Information Processing Systems*, 33: 9851–9864.
- Deb, K.; and Gupta, H. 2005. Searching for robust Pareto-optimal solutions in multi-objective optimization. In *International conference on evolutionary multi-criterion optimization*, 150–164. Springer.
- Deb, K.; Pratap, A.; Agarwal, S.; and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation*, 6(2): 182–197.
- Deb, K.; Thiele, L.; Laumanns, M.; and Zitzler, E. 2005. Scalable test problems for evolutionary multiobjective optimization. In *Evolutionary multiobjective optimization: theoretical advances and applications*, 105–145. Springer.
- Désidéri, J.-A. 2012. Multiple-gradient descent algorithm (MGDA) for multiobjective optimization. *Comptes Rendus Mathématique*, 350(5-6): 313–318.
- Fliege, J.; and Svaiter, B. F. 2000. Steepest descent methods for multicriteria optimization. *Mathematical methods of operations research*, 51: 479–494.
- Hansen, N.; and Ostermeier, A. 2001. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2): 159–195.
- Jaggi, M. 2013. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. In *International conference on machine learning*, 427–435. PMLR.
- Jain, H.; and Deb, K. 2013. An evolutionary many-objective optimization algorithm using reference-point based non-dominated sorting approach, part II: Handling constraints and extending to an adaptive approach. *IEEE Transactions on evolutionary computation*, 18(4): 602–622.
- Konakovic Lukovic, M.; Tian, Y.; and Matusik, W. 2020. Diversity-guided multi-objective bayesian optimization with batch evaluations. *Advances in Neural Information Processing Systems*, 33: 17708–17720.
- Liao, X.; Li, Q.; Yang, X.; Zhang, W.; and Li, W. 2008. Multiobjective optimization for crash safety design of vehicles using stepwise regression model. *Structural and multidisciplinary optimization*, 35: 561–569.
- Lin, Z. J.; Astudillo, R.; Frazier, P.; and Bakshy, E. 2022. Preference exploration for efficient bayesian optimization with multiple outcomes. In *International Conference on Artificial Intelligence and Statistics*, 4235–4258. PMLR.
- Makrygiorgos, G.; Paulson, J. A.; and Mesbah, A. 2023. No-Regret Bayesian Optimization with Gradients Using Local Optimality-Based Constraints: Application to Closed-Loop Policy Search. In *Proceedings of the 62nd IEEE Conference on Decision and Control*, 20–25.
- McLeod, M.; Roberts, S.; and Osborne, M. A. 2018. Optimization, fast and slow: optimally switching between local and Bayesian optimization. In *International Conference on Machine Learning*, 3443–3452. PMLR.
- Miettinen, K. 1999. *Nonlinear multiobjective optimization*, volume 12. Springer Science & Business Media.
- Müller, S.; von Rohr, A.; and Trimpe, S. 2021. Local policy search with Bayesian optimization. *Advances in Neural Information Processing Systems*, 34: 20708–20720.
- Nguyen, Q.; Wu, K.; Gardner, J.; and Garnett, R. 2022. Local Bayesian optimization via maximizing probability of descent. *Advances in neural information processing systems*, 35: 13190–13202.
- Ozaki, R.; Ishikawa, K.; Kanzaki, Y.; Takeno, S.; Takeuchi, I.; and Karasuyama, M. 2024. Multi-Objective Bayesian Optimization with Active Preference Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 14490–14498.
- Parsons, M. G.; and Scott, R. L. 2004. Formulation of multicriterion design optimization problems for solution with scalar numerical optimization methods. *Journal of Ship Research*, 48(01): 61–76.
- Schäffler, S.; Schultz, R.; and Weinzierl, K. 2002. Stochastic method for the solution of unconstrained vector optimization problems. *Journal of Optimization Theory and Applications*, 114: 209–222.
- Sener, O.; and Koltun, V. 2018. Multi-task learning as multi-objective optimization. *Advances in neural information processing systems*, 31.

Shao, K.; Romeres, D.; Chakrabarty, A.; and Mesbah, A. 2023. Preference-Guided Bayesian Optimization for Control Policy Learning: Application to Personalized Plasma Medicine. In *NeurIPS 2023 Workshop on Adaptive Experimental Design and Active Learning in the Real World*.

Tanabe, R.; and Ishibuchi, H. 2020. An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89: 106078.

Williams, C. K.; and Rasmussen, C. E. 2006. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA.

Wilson, J.; Hutter, F.; and Deisenroth, M. 2018. Maximizing acquisition functions for Bayesian optimization. *Advances in neural information processing systems*, 31.

Wu, J.; Poloczek, M.; Wilson, A. G.; and Frazier, P. I. 2017. Gradient-augmented Bayesian optimization. In *Advances in Neural Information Processing Systems*, 1033–1043. NeurIPS.