



# Vector Optimization with Gaussian Process Bandits

İlter Onat Korkmaz<sup>1</sup> · Yaşar Cahit Yıldırım<sup>1</sup> · Çağın Ararat<sup>1</sup> · Cem Tekin<sup>1</sup>

Received: 24 July 2024 / Revised: 10 November 2025 / Accepted: 2 March 2026 /  
Published online: 23 March 2026  
© The Author(s) 2026

## Abstract

We study black-box vector optimization with Gaussian process bandits, where there is an incomplete order relation on objective vectors described by a polyhedral convex cone. Existing black-box vector optimization approaches either suffer from high sample complexity or lack theoretical guarantees. We propose *Vector Optimization with Gaussian Process* (VOGP), an adaptive elimination algorithm that identifies Pareto optimal solutions sample efficiently by exploiting the smoothness of the objective function. We establish theoretical guarantees, deriving information gain-based and kernel-specific sample complexity bounds. Finally, we conduct a thorough empirical evaluation of VOGP and compare it with the state-of-the-art multi-objective and vector optimization algorithms on several real-world and synthetic datasets, emphasizing VOGP's efficiency (e.g.,  $\sim 18\times$  lower sample complexity on average). We also provide heuristic adaptations of VOGP for cases where the design space is continuous and where the Gaussian process model lacks access to the true kernel hyperparameters. This work opens a new frontier in sample-efficient multi-objective black-box optimization by incorporating preference structures while maintaining theoretical guarantees and practical efficiency.

**Keywords** Vector optimization · Ordering cones · Gaussian process bandits · Bayesian optimization · Sample complexity bounds

---

Editors: Kwang-Sung Jun and Michelangelo Ceci.

---

✉ Cem Tekin  
cemtekin@ee.bilkent.edu.tr  
İlter Onat Korkmaz  
onat.korkmaz@bilkent.edu.tr  
Yaşar Cahit Yıldırım  
cahit.yildirim@bilkent.edu.tr  
Çağın Ararat  
cararat@bilkent.edu.tr


<sup>1</sup> Bilkent University, Ankara, Türkiye

## 1 Introduction

In diverse fields such as engineering, economics, and computational biology, the problem of identifying the best set of designs across multiple black-box objectives is a recurrent challenge. Often, the evaluation of a particular design is both expensive and noisy, leading to the need for efficient and reliable optimization methods. This issue is exemplified in fields like aerospace engineering and drug discovery, where each evaluation can involve costly experiments or simulations (Obayashi et al., 1997; Mukaidaisi et al., 2022). A common framework for dealing with such multi-objective optimization problems is to focus on identifying Pareto optimal designs, for which there are no other designs that are better in all objectives. However, this approach can be restrictive, as it only permits a certain set of trade-offs between objectives where a disadvantage on one objective cannot be compensated by advantages on other objectives, i.e., the domination relation requires superiority in all dimensions.

A more comprehensive solution to this issue lies in vector optimization, where the partial ordering relations induced by convex cones are used. The core idea is to define a “cone” in the objective space that represents preferred directions of improvement. A solution is considered better than another if the vector difference between their objectives falls within this preference cone. The use of ordering cones in this way gives users a framework that can model a wide range of trade-off preferences (Jahn, 2010).

Consider the challenge of optimizing the reactor parameters for nucleophilic aromatic substitution (SnAr) reaction, an important process in synthetic chemistry (Takao, 2022; Grant et al., 2018). In this scenario, the goal is to maximize the space-time yield and minimize the waste-to-product ratio, i.e., the E-factor.<sup>1</sup> Consider two designs where one of the designs has slightly less space-time yield and a much lower E-factor. A reasonable preference would be to opt for the design that significantly reduces environmental impact, provided that the loss in space-time yield is not too substantial. Similarly, choosing designs that offer considerably higher space-time yield, albeit with a slightly increased environmental footprint, might also be justified. In vector optimization, this preference can be encoded



**Fig. 1** Vector optimization on nucleophilic aromatic substitution (SnAr) reaction data (Hone et al., 2017). (Left) The ordering cone. (Right) Cone-specific Pareto front ( $P^*$ ) and its comparison to Pareto front of the componentwise order ( $P_{\pi/2}^*$ )

<sup>1</sup>We provide experiments on this optimization task in Section 6.


with an obtuse cone (Figure 1). The regular Pareto optimality concept falls short in conveying such preferences as it requires superiority in all objectives for domination relations. In Figure 1, it can be seen that some of the designs (shown in red) in the Pareto front of traditional multi-objective order are not in the obtuse cone-based Pareto front, i.e., they are dominated by designs that are much better in one objective and slightly worse in the other.

The example we just discussed demonstrates vector optimization within a two-dimensional objective space, as this lower dimensionality allows for straightforward visualization. However, when extending to higher dimensions, the structure of the polyhedral cone may be more intricate with high number of faces. This is exemplified by Figure 2, which shows a polyhedral cone with six faces in three dimensions. Such cones are not uncommon in the literature. For example, Hunt et al. (2010) provide a model of relative importance where allowable trade-offs between pairs of (possibly many) objectives are given. They provide a complete algebraic characterization of the cones corresponding to these kinds of preferences and show that they result in highly faceted polyhedral cones.

While there exists previous work addressing Pareto set identification with respect to an ordering cone in noisy environments with theoretical guarantees (Ararat, 2023; Karagözlü et al., 2024), they tend to require vast amounts of samples from the objective function, making them impractical in real-world scenarios. In this paper, we address the problem of identifying the Pareto set with respect to an ordering cone  $C$  from a given set of designs with a minimum number of noisy black-box function queries. We propose *Vector Optimization with Gaussian Process* (VOGP), an  $(\epsilon, \delta)$ -probably approximately correct (PAC) adaptive elimination algorithm that performs black-box vector optimization by utilizing the modeling power of Gaussian processes (GPs). In particular, VOGP leverages confidence regions formed by GPs to perform cone-dependent, probabilistic elimination and Pareto identification steps, which adaptively explore the Pareto front with Bayesian optimization (Garnett, 2023). Similar to Pareto Active Learning (PAL) and  $\epsilon$ -PAL (Zuluaga et al., 2013, 2016) for the multi-objective case (see next subsection), VOGP operates in a fully exploratory setting, where the least certain design is queried at each round. When VOGP is considered with the special case of componentwise order (positive orthant cone in vector optimization), it can be seen as a variant of PAL and  $\epsilon$ -PAL that can handle correlated objectives. We prove strong theoretical guarantees on the convergence of VOGP and derive information gain-based sample complexity bounds. Our main contributions can be listed as follows:

- We propose VOGP, an  $(\epsilon, \delta)$ -PAC adaptive elimination algorithm that performs vector optimization by utilizing GPs. This is the first work to provide a theoretical foundation

Fig. 2 A 3D cone



for vector optimization within the framework of GP bandits.

- When the design set is finite, we theoretically prove that VOGP returns an  $(\epsilon, \delta)$ -PAC Pareto set with a cone-dependent sample complexity bound in terms of the maximum information gain. We also provide explicit kernel-specific sample complexity bounds.
- We empirically show that VOGP satisfies the theoretical guarantees across different datasets and ordering cones. We also demonstrate that VOGP outperforms existing methods on vector optimization and its special case, multi-objective optimization.
- We extend VOGP to work on continuous design sets by using tree-based adaptive discretization techniques. We empirically evaluate this extension in continuous design sets. We also study a heuristic variant of VOGP to handle the case of unknown GP kernel hyperparameters.

## 2 Related Work

There is considerable existing work on Pareto set identification that utilizes GPs (Lyu et al., 2018; Suzuki et al., 2020; Shu et al., 2020; Mathern et al., 2021; Picheny, 2013; Emmerich et al., 2011; Svenson & Santner, 2016; Candelieri et al., 2024). Some of these works use information-theoretic acquisition functions. For instance, PESMO tries to sample designs that optimize the mutual information between the observation and the Pareto set, given the dataset (Hernandez-Lobato et al., 2016). MESMO tries to sample designs that maximize the mutual information between observations and the Pareto front (Belakaria et al., 2019). JES tries to sample designs that maximize the mutual information between observations and the joint distribution of the Pareto set and Pareto front (Tu et al., 2022). Different from entropy search-based methods, some other methods utilize high-probability confidence regions formed by the GP posterior. For instance, PALS (Barracosa et al., 2022), PAL and  $\epsilon$ -PAL are confidence region-based adaptive elimination algorithms that aim to categorize the input data points into three groups using the models they have learned: those that are Pareto optimal, those that are not Pareto optimal, and those whose statuses are uncertain (Zuluaga et al., 2013, 2016). In every iteration, these algorithms choose a potential design for evaluation with the aim of reducing the quantity of data points in the uncertain category. Thanks to GP-based modeling, a substantial number of designs can be explored by a single query. Assuming independent objectives, PAL and  $\epsilon$ -PAL work in the fixed confidence setting and have a variable sampling budget. On the other hand, entropy search-based methods are often investigated in fixed-budget setting without any theoretical guarantees on the accuracy of Pareto set identification. However, all the algorithms mentioned above try to identify the Pareto set with respect to the componentwise order. They are not able to incorporate user preferences encoded by ordering cones.

In the context of multi-objective Bayesian optimization, there is a limited amount of work that incorporates user preferences. Abdolshah et al. (2019) propose MOBO-PC, a multi-objective Bayesian optimization algorithm which incorporates user preferences in the form of preference-order constraints. MOBO-PC uses expected Pareto hypervolume improvement weighted by the probability of obeying preference-order constraints as its acquisition function. Yang et al. (2016a) propose a method that uses truncated expected hypervolume improvement and considers the predictive mean, variance, and the preferred region in the objective space. Some works employ preference learning (Chu, a; Hakanen, 2017; Taylor et

al., 2021; Ungredda, 2023; Ignatenko et al., 2021; Bemporad & Piga, 2021), where the user interacts sequentially with the algorithm to learn the user preference (Astudillo, 2019; Lin et al., 2022). Astudillo (2019) employ Bayesian preference learning where the user's preferences are modeled as a utility function and they propose two novel acquisition functions that are robust to utility uncertainty. Lin et al. (2022) consider various preference exploration methods while also representing user preferences with a utility function which is modeled with a GP. Ahmadianshalchi et al. (2024) propose PAC-MOO, a constrained multi-objective Bayesian optimization algorithm which incorporates preferences in the form of weights that add up to one. PAC-MOO uses the information gained about the optimal constrained Pareto front weighted by preference weights as its acquisition function. Khan et al. (2022) propose a method that learns a utility function offline by using expert knowledge to avoid the repeated and expensive expert involvement.

In the broader context of gradient-free optimization, zeroth-order optimization techniques have been explored to handle optimization problems where gradient information is unavailable (Dang et al., 2024). In the multi-objective case, numerous studies employ evolutionary algorithms in order to estimate the Pareto front, thereby accomplishing this task through the iterative development of a population of evaluated designs (Seada, 2015; Deb et al., 2002; Knowles, 2006; Zhang & Li, 2007; Dang et al., 2023, 2024). Some of these methods use hypervolume calculation to guide their method (Yang et al., 2016b; Beume et al., 2007; Bader & Zitzler, 2011). Dang et al. (2024) incorporates generative adversarial networks into evolutionary algorithms to improve offspring diversity in the context of multimodal multi-objective optimization. Dang et al. (2023) uses multi-objective evolutionary algorithms in federated learning for device selection, where multi-objective formulation helps balance resource efficiency and test accuracy. There is a significant body of research that deals with the integration of user preferences into multi-objective evolutionary algorithms (Branke, 2005; Coello, 2000; Li, 2008; Phelps & Köksalan, 2003; Sinha et al., 2010; Zhou et al., 2011). Using evolutionary algorithms, preferences can be articulated *a priori*, *a posteriori*, or interactively. Various methods for expressing preferences have been suggested, primarily encompassing techniques based on objective comparisons, solution ranking, and expectation-based approaches (Zhou et al., 2011). There are also works that employ preference cones in multi-objective evolutionary algorithms. Batista et al. (2011) utilize polyhedral cones as a method of managing the resolution of the estimated Pareto front. Ferreira et al. (2020) apply preference cone-based multi-objective evolutionary algorithm to optimize distributed energy resources in microgrids.

There are techniques that transform a multi-objective optimization problem into a single-objective problem by assigning weights to the objectives (Knowles, 2006; Ponweiser et al., 2008). The transformed problem can be solved using standard single-objective optimization methods. ParEGO introduced in Knowles (2006) transforms the multi-objective optimization problem into a single-objective one by using Tchebycheff scalarization and solves it by Efficient Global Optimization (EGO) algorithm, which was designed for single-objective expensive optimization problems.

There are existing works that incorporate polyhedral structures to guide the design identification task. As an example, Katz-Samuels (2018) introduce the concept of feasible arm identification. Their objective is to identify arms whose average rewards fall within a specified polyhedron, using evaluations that are subject to noise. Ararat (2023) propose a Naïve Elimination (NE) algorithm for Pareto set identification with polyhedral ordering cones.

They provide sample complexity bounds on the  $(\epsilon, \delta)$ -PAC Pareto set identification performance of this algorithm. However, their algorithm assumes independent arms and does not perform adaptive elimination. As a result, they only have worst-case sample complexity bounds for this algorithm. Experiments show that identification requires a large sampling budget which renders NE impractical in real-world problems of interest. Karagözlü et al. (2024) introduce the PaVeBa algorithm, designed for identifying Pareto sets using polyhedral ordering cones. They also establish sample complexity bounds for the  $(\epsilon, \delta)$ -PAC Pareto set identification task within bandit settings. However, their analysis lacks theoretical guarantees when Gaussian processes are employed.

Table 1 compares our work with the most closely related prior works in terms of key aspects. An important distinction is that our method supports cone-based preferences while also offering theoretical sample complexity guarantees in PAC framework. While Karagözlü et al. (2024) and Ararat (2023) also incorporate cone-based preferences and provide sample complexity bounds, they do not utilize Gaussian processes and are restricted to discrete domains. In contrast, information theoretic methods such as PESMO (Hernandez-Lobato et al., 2016), MESMO (Belakaria et al., 2019), and JES (Tu et al., 2022) use Gaussian processes and operate in continuous domains but do not offer sample complexity guarantees, nor do they incorporate cone-based preferences. Zuluaga et al. (2013) employs Gaussian processes but does not incorporate cone-based preferences, and is also restricted to discrete domains.

### 3 Background and Problem Definition

The aim of this work is to incorporate polyhedral cones as preference structures to guide the multi-objective optimization, while maintaining theoretical guarantees and practical efficiency. Formally, we consider the problem of sequentially maximizing a vector-valued function  $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^M$  over a finite set of designs<sup>2</sup>  $\mathcal{X} \subset \mathbb{R}^D$  with respect to a polyhedral ordering cone  $C \subset \mathbb{R}^M$ . In each iteration  $t$ , a design point  $x_t$  is selected, and a noisy observation is recorded as  $\mathbf{y}_t = \mathbf{f}(x_t) + \boldsymbol{\nu}_t$ . Here, the noise component  $\boldsymbol{\nu}_t$  has the multivariate Gaussian distribution  $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{M \times M})$ , where  $\sigma^2$  denotes the variance of the noise. To define the optimality of designs, we use the partial order induced by  $C$ . A detailed table of notations can be found in Appendix A.

**Table 1** Comparison with Related Works

| Paper                          | Cone-based preferences | Sample complexity bounds | $(\epsilon, \delta)$ -PAC | Utilizes GPs | Continuous domain |
|--------------------------------|------------------------|--------------------------|---------------------------|--------------|-------------------|
| This work                      | Yes                    | Yes                      | Yes                       | Yes          | Yes               |
| Karagözlü et al. (2024)        | Yes                    | Yes                      | Yes                       | No           | No                |
| Ararat (2023)                  | Yes                    | Yes                      | Yes                       | No           | No                |
| Tu et al. (2022)               | No                     | No                       | No                        | Yes          | Yes               |
| Belakaria et al. (2019)        | No                     | No                       | No                        | Yes          | Yes               |
| Hernandez-Lobato et al. (2016) | No                     | No                       | No                        | Yes          | Yes               |
| Zuluaga et al. (2016)          | No                     | Yes                      | Yes                       | Yes          | No                |

<sup>2</sup>Extension to continuous case is discussed in Section 6.

### 3.1 Partial Order Induced by a Cone

The partial order among designs is induced by a known polyhedral ordering cone  $C \subseteq \mathbb{R}^M$ , that is, a polyhedral closed convex cone whose interior, denoted by  $\text{int}(C)$ , is nonempty with  $C \cap -C = \{\mathbf{0}\}$ . Such a cone can be written as  $C = \{\mathbf{x} \in \mathbb{R}^M \mid \mathbf{W}\mathbf{x} \geq \mathbf{0}\}$ , where  $\mathbf{W}$  is a  $N \times M$  matrix,  $N$  is the number of halfspaces that define  $C$ , and  $N \geq M$ . Without loss of generality, the rows  $\mathbf{w}_1^\top, \dots, \mathbf{w}_N^\top$  are assumed to be unit normal vectors of these halfspaces with respect to the Euclidean norm  $\|\cdot\|_2$ . For  $r \geq 0$ , we write  $\mathbb{B}(r) := \{\mathbf{y} \in \mathbb{R}^M \mid \|\mathbf{y}\|_2 \leq r\}$ . We say that  $\mathbf{y} \in \mathbb{R}^M$  weakly dominates  $\mathbf{y}' \in \mathbb{R}^M$  with respect to  $C$  if their difference lies in  $C$ :  $\mathbf{y}' \preceq_C \mathbf{y} \Leftrightarrow \mathbf{y} - \mathbf{y}' \in C$ . By the structure of  $C$ , an equivalent expression is:

$$\mathbf{y}' \preceq_C \mathbf{y} \Leftrightarrow \mathbf{w}_n^\top (\mathbf{y} - \mathbf{y}') \geq 0, \forall n \in [N] := \{1, \dots, N\} . \tag{1}$$

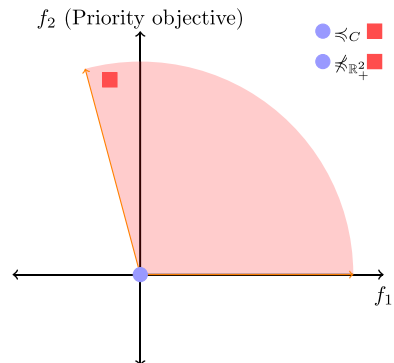
Moreover, we say that  $\mathbf{y}$  strictly dominates  $\mathbf{y}'$  with respect to  $C$  if their difference lies in  $\text{int}(C)$ :  $\mathbf{y}' \prec_C \mathbf{y} \Leftrightarrow \mathbf{y} - \mathbf{y}' \in \text{int}(C)$ . This is equivalent to:  $\mathbf{y}' \prec_C \mathbf{y} \Leftrightarrow \mathbf{w}_n^\top (\mathbf{y} - \mathbf{y}') > 0, \forall n \in [N]$ . Similarly, we write  $\mathbf{y}' \preceq_{C \setminus \{0\}} \mathbf{y} \Leftrightarrow \mathbf{y} - \mathbf{y}' \in C \setminus \{0\}$  to exclude the case where the two vectors coincide. The Pareto set with respect to cone  $C$  is the set of designs that are not weakly dominated by any other design with respect to  $C$  given as

$$P^* := P_C^* := \{x \in \mathcal{X} \mid \nexists x' \in \mathcal{X} : \mathbf{f}(x) \preceq_{C \setminus \{0\}} \mathbf{f}(x')\}.$$

**Example 1** Consider a scenario with two objectives ( $M = 2$ ), denoted by  $f_1$  and  $f_2$ . In this example, we assume that  $f_2$  is prioritized over  $f_1$ , meaning the user is willing to trade off some amount of  $f_1$  for a sufficient gain in  $f_2$ . Figure 3 illustrates the corresponding preference cone, where the red square is preferred over the blue circle. This preference structure differs from the componentwise order, under which these designs would be incomparable.

**Remark 1** The relation  $\preceq_C$  induced by the cone  $C$  is a vector preorder on  $\mathbb{R}^N$  that is non-total (incomplete). The choice  $C = \mathbb{R}_+^M$  corresponds to the usual componentwise order in multi-objective optimization. In general, using a larger cones results in a smaller Pareto set; we refer the reader to Ararat and Tekin (2023, Remark 2.3) for possible uses of larger cones in small molecule drug discovery problem. In multi-asset markets with proportional transaction costs in finance,  $C$  can be chosen as the so-called solvency cone determined by

**Fig. 3** Visualization of a two-dimensional cone that prioritizes the  $f_2$  objective. Unlike the componentwise order, the cone order establishes a domination relation by prioritizing  $f_2$ . The dominance relation under the preference cone is evident, as the difference of red square from blue circle lies inside the cone, confirming that the red square is preferred according to this ordering



the transaction costs rates; see Kabanov (1999). Cone-induced preorders form an important special class of preorders. The underlying scaling property, encoded by  $C$  being a cone, is a strong assumption whose suitability should be justified in the particular application. For instance, in the case of solvency cones, the assumption that the transaction costs are incurred in proportion to the volume at a fixed rate justifies the use of a cone naturally.

**Remark 2** In this paper, we assume that the non-total order relation  $\preceq_C$  is given a priori and the optimization task is performed with respect to it. A different strand of literature focuses on the problem of *learning* the underlying order relation from data; see, e.g., Chu (b); Chau et al. (2022); Audiffren (2017); Yue et al. (2012). Since the focus of the current paper is optimization with respect to a known vector preorder, the approach of these references is orthogonal to our work.

### 3.2 $(\epsilon, \delta)$ -PAC Pareto Set

We use the cone-dependent, direction-free suboptimality gaps defined in Ararat (2023). The gap between designs  $x, x' \in \mathcal{X}$  is given as

$$m(x, x') := \inf\{s \geq 0 \mid \exists \mathbf{u} \in \mathbb{B}(1) \cap C : \mathbf{f}(x) + s\mathbf{u} \notin \mathbf{f}(x') - \text{int}(C)\},$$

the  $-$  operation on sets is the Minkowski difference (vectors being treated as singletons). The suboptimality gap of design  $x$  is defined as  $\Delta_x^* := \max_{x' \in P^*} m(x, x')$ .

**Definition 1** Let  $\epsilon > 0, \delta \in (0, 1)$ . A random set  $P \subseteq \mathcal{X}$  is called an  $(\epsilon, \delta)$ -PAC Pareto set if the subsequent conditions are satisfied at least with probability  $1 - \delta$ :

$$(i) \bigcup_{x \in P} (\mathbf{f}(x) + \mathbb{B}(\epsilon) \cap C - C) \supseteq \bigcup_{x \in P^*} (\mathbf{f}(x) - C); (ii) \forall x \in P \setminus P^* : \Delta_x^* \leq 2\epsilon.$$

**Remark 3** Condition (i) of Definition 1 is equivalent to the following: For every  $x^* \in P^*$ , there exist  $x \in P$  and  $\mathbf{u} \in \mathbb{B}(\epsilon) \cap C$  such that  $\mathbf{f}(x^*) \preceq_C \mathbf{f}(x) + \mathbf{u}$ . Even though certain designs in  $P$  may not be optimal, condition (ii) of Definition 1 limits the extent of the inadequacies in these designs. Consequently, it ensures the overall quality of all the designs produced.

Our aim is to design an algorithm that returns an  $(\epsilon, \delta)$ -PAC Pareto set  $\hat{P} \subseteq \mathcal{X}$  with as little sampling from the expensive objective function  $\mathbf{f}$  as possible.

### 3.3 Modeling of the Objective Function

We model the objective function  $\mathbf{f}$  as a realization of an  $M$ -output GP with zero mean and a positive definite covariance function  $\mathbf{k}$  with bounded variance:  $k^{jj}(x, x) \leq 1$  for every  $x \in \mathcal{X}$  and  $j \in [M]$ . Let  $\tilde{x}_i$  be the  $i$ th design observed and  $\mathbf{y}_i$  the corresponding observation. The posterior distribution of  $\mathbf{f}$  conditioned on the first  $t$  observations is that of an  $M$ -output GP with mean  $(\boldsymbol{\mu}_t)$  and covariance  $(\mathbf{k}_t)$  functions given below, where  $\mathbf{k}_t(x) = [\mathbf{k}(x, x_1), \dots, \mathbf{k}(x, x_t)] \in \mathbb{R}^{M \times Mt}$ ,  $\mathbf{y}_{[t]} = [\mathbf{y}_1^\top, \dots, \mathbf{y}_t^\top]^\top$ ,  $\mathbf{K}_t = (\mathbf{k}(x_i, x_j))_{i, j \in [t]}$ , and  $\mathbf{I}_{Mt}$  is the  $Mt \times Mt$  identity matrix:

$$\boldsymbol{\mu}_t(x) = \mathbf{k}_t(x) (\mathbf{K}_t + \sigma^2 \mathbf{I}_{Mt})^{-1} \mathbf{y}_{[t]}^\top, \quad (2)$$

$$\mathbf{k}_t(x, x') = \mathbf{k}(x, x') - \mathbf{k}_t(x) (\mathbf{K}_t + \sigma^2 \mathbf{I}_{Mt})^{-1} \mathbf{k}_t(x')^\top. \quad (3)$$

Our sample complexity results will depend on the difficulty of learning  $f$ , which is characterized by its maximum information gain.

**Definition 2** The maximum information gain at round  $t$  is defined as  $\gamma_t := \max_{A \subseteq \mathcal{X}: |A|=t} \mathbb{I}(\mathbf{y}_A; \mathbf{f}_A)$ , where  $\mathbf{y}_A$  is the collection of observations corresponding to the designs in  $A$ ,  $\mathbf{f}_A$  is the collection of the corresponding function values, and  $\mathbb{I}(\mathbf{y}_A; \mathbf{f}_A)$  is the mutual information between the two.

## 4 Method

We propose *Vector Optimization with Gaussian Process* (VOGP), an  $(\epsilon, \delta)$ -PAC adaptive elimination algorithm that performs vector optimization by utilizing the modeling power of GPs (see Algorithm 1). VOGP takes as input the design set  $\mathcal{X}$ , cone  $C$ , accuracy parameter  $\epsilon$ , and confidence parameter  $\delta$ . VOGP adaptively classifies designs as suboptimal, Pareto optimal, and uncertain. Initially, all designs are put in the set of uncertain designs, denoted by  $\mathcal{S}_t$ . Designs that are classified as suboptimal are discarded and never considered again in comparisons, making it sample-efficient. Designs that are classified as Pareto optimal are moved to the predicted Pareto set  $\hat{\mathcal{P}}_t$ . VOGP tracks the set of active designs by assigning them to  $\mathcal{A}_t$  before discarding and  $\mathcal{W}_t$  afterwards. VOGP terminates and returns  $\hat{\mathcal{P}} = \hat{\mathcal{P}}_t$  when the set of uncertain designs becomes empty. Decision-making in each round is divided into four phases named *modeling*, *discarding*, *Pareto identification*, *evaluating*. Below, we describe each phase in detail.

---

**Input:** Design set  $\mathcal{X}$ ,  $\epsilon \geq 0$ ,  $\delta \in (0, 1)$ , polyhedral ordering cone  $C$ , kernel  $\mathbf{k}$   
**Initialize:** Predicted Pareto set  $\mathcal{P}_1 = \emptyset$ , set of uncertain designs  $\mathcal{S}_1 = \mathcal{X}$ , sets of active designs  $\mathcal{A}_1 = \emptyset$  and  $\mathcal{W}_1 = \emptyset$ ,  $\mathbf{R}_0(x) = \mathbb{R}^M$  for each  $x \in \mathcal{S}_1$ ,  $t = 1$   
**Compute:** Accuracy vector  $\mathbf{u}^* \in C$  given in Definition 4  
**while**  $\mathcal{S}_t \neq \emptyset$  **do**  
    Run MODELING( $\mathcal{S}_t, \mathcal{P}_t$ ) in Algorithm 2  
    Run DISCARDING( $\mathcal{S}_t, \mathcal{A}_t$ ) in Algorithm 3  
    Run PARETOIDENTIFICATION( $\mathcal{S}_t, \mathcal{P}_t$ ) in Algorithm 4  
    Run EVALUATING( $\mathcal{S}_t, \mathcal{W}_t$ ) in Algorithm 5  
     $\mathcal{P}_{t+1} = \mathcal{P}_t$ ,  $\mathcal{S}_{t+1} = \mathcal{S}_t$   
     $t = t + 1$   
**end while**  
**return** Predicted Pareto set  $\hat{\mathcal{P}} = \mathcal{P}_t$

---

### Algorithm 1 VOGP

**Remark 4** Our theoretical results, similar to other theoretical works related to ours (Srinivas et al., 2012; Zuluaga et al., 2013), assume that the kernel is known and fixed. In practice,

however, the kernel may be unknown and must be estimated from data concurrently with the decision-making process; see, e.g., Berkenkamp et al. (2019); Ziomek et al. (2024). We discuss this practical case in Section 6.5, where we provide an empirical adaptation of Algorithm 1 for unknown hyperparameters.

### 4.1 Modeling

This phase is summarized in Algorithm 2. VOGP uses the GP posterior means and variances of designs to define confidence regions in the form of  $M$ -dimensional hyperrectangles which are scaled by  $\beta_t$ , a function of the confidence parameter  $\delta$ . The probability that these hyperrectangles include the true objective values of the designs is at least  $1 - \delta$ . Using these confidence hyperrectangles allows for the formulation of probabilistic success criteria and the quantification of uncertainty with designs. At each round  $t$ , the hyperrectangle

$$Q_t(x) = \{y \in \mathbb{R}^M \mid \mu_t^j(x) - \beta_t^{1/2} \sigma_t^j(x) \leq y^j \leq \mu_t^j(x) + \beta_t^{1/2} \sigma_t^j(x), \forall j \in [M]\} \tag{4}$$

is computed for design  $x$ , where  $\beta_t = 2 \ln(M\pi^2 |\mathcal{X}| t^2 / (3\delta))$ .

---

```

Compute: Set of active designs  $\mathcal{A}_t = \mathcal{S}_t \cup \mathcal{P}_t$ ;
for  $x \in \mathcal{A}_t$  do
    Obtain GP posterior for  $x$ :  $\mu_t(x)$  and  $\sigma_t(x)$  given in Equation 2;
    Construct confidence hyperrectangle  $Q_t(x)$  given in Equation 4;
     $R_t(x) = R_{t-1}(x) \cap Q_t(x)$ ;
end for
    
```

---

**Algorithm 2** Modeling Subroutine

Moreover, all the calculated hyperrectangles up to round  $t$  are intersected cumulatively to obtain the cumulative confidence hyperrectangle  $R_t(x)$  of design  $x$ .

### 4.2 Discarding

This phase is summarized in Algorithm 3. VOGP discards undecided designs that are dominated with respect to  $C$  with high probability. The designs that are to be discarded with low probability are identified and are used to form the pessimistic Pareto set  $\mathcal{P}_{\text{pess},t}^{(C)}(\mathcal{A}_t)$  as defined below. The illustration of discarding phase can be seen in Figure 4(Left, Middle).

**Definition 3 (Pessimistic Pareto Set)** Let  $t \geq 1$  and let  $D \subseteq \mathcal{A}_t$  be a nonempty set of nodes. The *pessimistic Pareto set* of  $D$  with respect to  $C$  at round  $t$ , denoted by  $\mathcal{P}_{\text{pess},t}^{(C)}(D)$ , is the set of all nodes  $x \in D$  for which there is no node  $x' \in D$  such that  $R_t(x') + C \subsetneq R_t(x) + C$ .


---

```

Compute:  $\mathcal{P}_{\text{pess},t} = \mathcal{P}_{\text{pess},t}^{(C)}(\mathcal{A}_t)$  as in Definition 3
for  $x \in \mathcal{S}_t \setminus \mathcal{P}_{\text{pess},t}$  do
    if  $\exists x' \in \mathcal{P}_{\text{pess},t} \forall y \in R_t(x) : R_t(x') + \epsilon u^* \subseteq y + C$  then
         $\mathcal{S}_t = \mathcal{S}_t \setminus \{x\}$ ;
    end if
end for
    
```

---

**Algorithm 3** Discarding Subroutine



**Fig. 4** (Left, Middle) Visualization of discarding, and (Right) Pareto identification phases of VOGP in two dimensional objective space. Note that all boxes correspond to the  $R_t(x_i)$  for corresponding true objective values  $f(x_i)$  for  $i = 1, 2, 3$ . In this simple case, in the middle plot, the intersection of the two cones are equivalent to the intersection of all cones coming out of  $R_t(x_1)$

**Remark 5** In the setting of Definition 3,  $\mathcal{P}_{\text{pess},t}^{(C)}(D)$  can be seen as the set of all nodes that yield maximal elements of a certain partial order relation on sets. Indeed, for  $A, A' \subseteq \mathbb{R}^M$ , let us write  $A' \preceq_C A$  if and only if  $A' \subseteq A + C$ , which holds if and only if  $A' + C \subseteq A + C$ . Then, it is easy to check that the set relation  $\preceq_C$  is a preorder, i.e., a reflexive and transitive binary relation, on the power set of  $\mathbb{R}^M$ . Let  $\sim_C$  denote the symmetric part of this relation, i.e.,  $A' \sim_C A$  if and only if  $A' \preceq_C A$  and  $A \preceq_C A'$ , which holds if and only if  $A' + C = A + C$ . Then,  $\sim_C$  is an equivalence relation on the power set of  $\mathbb{R}^M$  and the set of equivalence classes can be identified by the collection  $\mathcal{M}_C := \{A \subseteq \mathbb{R}^M \mid A = A + C\}$ . Then,  $\mathcal{M}_C$  is a partially ordered set with respect to  $\preceq_C$  and the relation  $\preceq_C$  coincides with the usual  $\subseteq$  relation on  $\mathcal{M}_C$ . It then follows that  $\mathcal{P}_{\text{pess},t}^{(C)}(D)$  is the set of all nodes  $x \in D$  for which  $R_t(x) + C$  is a maximal element of  $\{R_t(x') + C \mid x' \in \mathcal{A}_t\}$  with respect to  $\preceq_C$ . In particular, since every finite nonempty partially ordered set has a maximal element, we have  $\mathcal{P}_{\text{pess},t}^{(C)}(D) \neq \emptyset$ . The reader is referred to Hamel, Heyde, Löhne, Rudloff, and Schrage (2015, Section 2.1) for more details on set relations.

VOGP computes pessimistic Pareto set of  $\mathcal{A}_t$  by checking for each design in  $\mathcal{A}_t$  whether there is another design that prevents it from being in the pessimistic Pareto set. To see if design  $x$  prevents design  $x'$  from being in the pessimistic Pareto set, by solving a convex optimization problem, VOGP checks if every vertex of  $R_t(x')$  can be expressed as the sum of a vector in  $C$  and a vector in  $R_t(x)$ . VOGP refines the undecided set by utilizing the pessimistic Pareto set. It eliminates designs based on the following criteria: for a given design  $x$ , if all values inside its confidence hyperrectangle, when added to an "accuracy vector"  $\epsilon u^*$ , are dominated with respect to  $C$  by all values in the confidence hyperrectangle of another design  $x'$ , then the dominated design  $x$  is discarded. The dominating design  $x'$  should be a member of the pessimistic Pareto set.

**Definition 4** Let  $A_C := \bigcap_{u \in \mathbb{S}^{M-1}} (u + C)$  and  $d_C := \inf\{\|z\|_2 \mid z \in A_C\}$ . Then, there exists a unique vector  $z^* \in A_C$  such that  $d_C = \|z^*\|_2$  and we define  $u^* := \frac{z^*}{d_C}$ .

**Lemma 1** Let  $y, z \in \mathbb{R}^M$  and  $\tilde{p} \in \mathbb{B}(\frac{\epsilon}{d_C})$ .

1. If  $y + \tilde{p} \preceq_C z$ , then  $y \preceq_C z + \epsilon u^*$ .
2. If  $\tilde{p} \in \text{int}(\mathbb{B}(\frac{\epsilon}{d_C}))$  and  $y + \tilde{p} \preceq_C z$ , then  $y \prec_C z + \epsilon u^*$ .

Due to its technical nature, the proof of Lemma 1 is given in Section 5.1 along with the proofs of several other technical lemmas.

**Remark 6**  $d_C$  corresponds to the magnitude of the smallest translation of a unit sphere so that the whole unit sphere is inside the cone  $C$  (i.e., so that the origin dominates the unit sphere).  $\mathbf{u}^*$  is the direction of this smallest translation.  $d_C$  quantifies how hard it is to reach to a domination relation when using a specific ordering cone. We call this term the *ordering hardness* of cone  $C$ .

### 4.2.1 Computation of the Pessimistic Pareto Set

Let  $t \in \mathbb{N}$ . According to Definition 3, to check if a design  $x \in \mathcal{A}_t$  is included in  $\mathcal{P}_{\text{pess},t}$ , the condition  $\mathbf{R}_t(x') + C \subsetneq \mathbf{R}_t(x) + C$  needs to be checked for every  $x' \in \mathcal{A}_t$ . We first formulate this condition in terms of convex feasibility problems. These problems can also be expressed as convex optimization problems whose objective functions are zero and we solve them numerically using the CVXPY library (Agrawal et al., 2018; Diamond & Boyd, 2016).

Let  $x, x' \in \mathcal{A}_t$ . Since  $\mathbf{R}_t(x')$  is a convex polytope, we have

$$\begin{aligned} \mathbf{R}_t(x') + C \subseteq \mathbf{R}_t(x) + C &\iff \mathbf{R}_t(x') \subseteq \mathbf{R}_t(x) + C \\ &\iff \forall \mathbf{v}' \in \text{vert}(\mathbf{R}_t(x')) : \mathbf{v}' \in \mathbf{R}_t(x) + C \\ &\iff \forall \mathbf{v}' \in \text{vert}(\mathbf{R}_t(x')) \exists \mathbf{y} \in \mathbf{R}_t(x) : \mathbf{y} \preceq_C \mathbf{v}'. \end{aligned}$$

Hence, for each  $\mathbf{v}' \in \text{vert}(\mathbf{R}_t(x'))$ , we check the condition  $\exists \mathbf{y} \in \mathbf{R}_t(x) : \mathbf{y} \preceq_C \mathbf{v}'$  by solving the following convex feasibility problem:

$$\begin{aligned} &\underset{\mathbf{y}}{\text{minimize}} && 0 \\ &\text{subject to} && \mathbf{y} \in \mathbb{R}^M, \\ & && y^j \geq \mu_s^j(x) - \beta_s^{1/2} \sigma_s^j(x), \quad \forall j \in [M], \forall s \in [t] \quad (F_t(x, \mathbf{v}')) \\ & && y^j \leq \mu_s^j(x) + \beta_s^{1/2} \sigma_s^j(x), \quad \forall j \in [M], \forall s \in [t] \\ & && \mathbf{w}_n^T \mathbf{y} \leq \mathbf{w}_n^T \mathbf{v}', \quad \forall n \in [N]. \end{aligned}$$

Moreover, under the condition  $\mathbf{R}_t(x') + C \subseteq \mathbf{R}_t(x) + C$ , we have

$$\begin{aligned} \mathbf{R}_t(x') + C \neq \mathbf{R}_t(x) + C &\iff \mathbf{R}_t(x) \not\subseteq \mathbf{R}_t(x') + C \\ &\iff \exists \mathbf{v} \in \text{vert}(\mathbf{R}_t(x)) : \mathbf{v} \notin \mathbf{R}_t(x') + C \\ &\iff \exists \mathbf{v} \in \text{vert}(\mathbf{R}_t(x)) \nexists \mathbf{y}' \in \mathbf{R}_t(x') : \mathbf{y}' \preceq_C \mathbf{v}. \end{aligned}$$

Hence, for each  $\mathbf{v} \in \text{vert}(\mathbf{R}_t(x))$ , we check the condition  $\exists \mathbf{y}' \in \mathbf{R}_t(x') : \mathbf{y}' \preceq_C \mathbf{v}$  by solving  $(F_t(x', \mathbf{v}))$  defined in the same way:

$$\begin{aligned} &\underset{\mathbf{y}'}{\text{minimize}} && 0 \\ &\text{subject to} && \mathbf{y}' \in \mathbb{R}^M, \\ & && (y')^j \geq \mu_s^j(x') - \beta_s^{1/2} \sigma_s^j(x'), \quad \forall j \in [M], \forall s \in [t] \quad (F_t(x', \mathbf{v})) \\ & && (y')^j \leq \mu_s^j(x') + \beta_s^{1/2} \sigma_s^j(x'), \quad \forall j \in [M], \forall s \in [t] \\ & && \mathbf{w}_n^T \mathbf{y}' \leq \mathbf{w}_n^T \mathbf{v}, \quad \forall n \in [N]. \end{aligned}$$

Hence, if there exists  $x' \in \mathcal{A}_t$  such that  $(F_t(x, v'))$  is feasible for every  $v' \in \text{vert}(\mathbf{R}_t(x'))$  and  $(F_t(x', v))$  is infeasible for at least one  $v \in \text{vert}(\mathbf{R}_t(x))$ , then  $x$  is not added to  $\mathcal{P}_{\text{pess},t}$ . Otherwise,  $x$  is added to  $\mathcal{P}_{\text{pess},t}$ .

### 4.2.2 Computation of Discarded Designs

In its original form, the discarding criteria of VOGP (see Algorithm 3) requires checking domination relations between infinitely many point pairs from two confidence hyperrectangles. In this section, we show that the sufficient number of checks can be reduced significantly by considering only the vertex pairs of the hyperrectangles, allowing for an efficient implementation. This is achieved thanks to the next lemma. For a nonempty convex polytope  $A \subseteq \mathbb{R}^M$ , we denote by  $\text{vert}(A)$  its set of all vertices.

**Lemma 2** *Let  $x, x' \in \mathcal{X}$  and  $t \in \mathbb{N}$ . For the cumulative confidence hyperrectangles  $\mathbf{R}_t(x)$ ,  $\mathbf{R}_t(x')$ , the following statements are equivalent.*

- (i)  $\forall \mathbf{y} \in \mathbf{R}_t(x) \forall \mathbf{y}' \in \mathbf{R}_t(x') : \mathbf{y} \preceq_C \mathbf{y}'$ .
- (ii)  $\forall \mathbf{v} \in \text{vert}(\mathbf{R}_t(x)) \forall \mathbf{v}' \in \text{vert}(\mathbf{R}_t(x')) : \mathbf{v} \preceq_C \mathbf{v}'$ .

**Proof** It is clear that (i) implies (ii) since  $\text{vert}(\mathbf{R}_t(x)) \subseteq \mathbf{R}_t(x)$  and  $\text{vert}(\mathbf{R}_t(x')) \subseteq \mathbf{R}_t(x')$ . To show the reverse implication, suppose that (ii) holds; let  $\mathbf{y} \in \mathbf{R}_t(x)$  and  $\mathbf{y}' \in \mathbf{R}_t(x')$ . We proceed in two steps.

First, let us further assume that  $\mathbf{y} \in \text{vert}(\mathbf{R}_t(x))$ . Moreover, since  $\mathbf{y}' \in \mathbf{R}_t(x')$  and  $\mathbf{R}_t(x')$  is a convex polytope, we may write

$$\mathbf{y}' = \sum_{\mathbf{v}' \in \text{vert}(\mathbf{R}_t(x'))} \lambda_{\mathbf{v}'} \mathbf{v}',$$

where  $\lambda_{\mathbf{v}'} \geq 0$  for each  $\mathbf{v}' \in \text{vert}(\mathbf{R}_t(x'))$  such that

$$\sum_{\mathbf{v}' \in \text{vert}(\mathbf{R}_t(x'))} \lambda_{\mathbf{v}'} = 1.$$

Note that (ii) implies that

$$\forall \mathbf{v}' \in \text{vert}(\mathbf{R}_t(x')) : \mathbf{y} \preceq_C \mathbf{v}',$$

which implies that

$$\forall \mathbf{v}' \in \text{vert}(\mathbf{R}_t(x')) : \lambda_{\mathbf{v}'} \mathbf{y} \preceq_C \lambda_{\mathbf{v}'} \mathbf{v}'.$$

Then, summing over all  $\mathbf{v}' \in \text{vert}(\mathbf{R}_t(x'))$  yields

$$\mathbf{y} = \sum_{\mathbf{v}' \in \text{vert}(\mathbf{R}_t(x'))} \lambda_{\mathbf{v}'} \mathbf{y} \preceq_C \sum_{\mathbf{v}' \in \text{vert}(\mathbf{R}_t(x'))} \lambda_{\mathbf{v}'} \mathbf{v}' = \mathbf{y}'.$$

Hence, (i) holds in this step.

Next, we consider the general case, where  $\mathbf{y} \in \mathbf{R}_t(x)$ . Since  $\mathbf{R}_t(x)$  is a convex polytope, we may write

$$\mathbf{y} = \sum_{\mathbf{v} \in \text{vert}(\mathbf{R}_t(x))} \alpha_{\mathbf{v}} \mathbf{v},$$

where  $\alpha_{\mathbf{v}} \geq 0$  for each  $\mathbf{v} \in \text{vert}(\mathbf{R}_t(x))$  such that

$$\sum_{\mathbf{v} \in \text{vert}(\mathbf{R}_t(x))} \alpha_{\mathbf{v}} = 1.$$

By the first step, we have

$$\forall \mathbf{v} \in \text{vert}(\mathbf{R}_t(x)) : \mathbf{v} \preceq_C \mathbf{y}'.$$

Then, we obtain

$$\forall \mathbf{v} \in \text{vert}(\mathbf{R}_t(x)) : \alpha_{\mathbf{v}} \mathbf{v} \preceq_C \alpha_{\mathbf{v}} \mathbf{y}'.$$

Finally, summing over all  $\mathbf{v} \in \text{vert}(\mathbf{R}_t(x))$  yields

$$\mathbf{y} = \sum_{\mathbf{v} \in \text{vert}(\mathbf{R}_t(x))} \alpha_{\mathbf{v}} \mathbf{v} \preceq_C \sum_{\mathbf{v} \in \text{vert}(\mathbf{R}_t(x))} \alpha_{\mathbf{v}} \mathbf{y}' = \mathbf{y}'.$$

Hence, (i) holds in the general case too.  $\square$

Let  $t \in \mathbb{N}$ . In Algorithm 3, to decide if a design  $x \in \mathcal{S}_t$  is going to be discarded, the condition

$$\exists x' \in \mathcal{P}_{\text{pess},t} \forall \mathbf{y} \in \mathbf{R}_t(x) : \mathbf{R}_t(x') + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C$$

needs to be checked, which requires enumerating all designs in  $\mathcal{P}_{\text{pess},t}$ . Given  $x' \in \mathcal{P}_{\text{pess},t}$ , observe that

$$\begin{aligned} \forall \mathbf{y} \in \mathbf{R}_t(x) : \mathbf{R}_t(x') + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C &\iff \forall \mathbf{y} \in \mathbf{R}_t(x) \forall \mathbf{y}' \in \mathbf{R}_t(x') : \mathbf{y} \preceq_C \mathbf{y}' + \epsilon \mathbf{u}^* \\ &\iff \forall \mathbf{y} \in \mathbf{R}_t(x) \forall \bar{\mathbf{y}} \in \mathbf{R}_t(x') + \epsilon \mathbf{u}^* : \bar{\mathbf{y}} \preceq_C \mathbf{y}. \end{aligned}$$

Hence, the last condition can be checked equivalently instead of the original condition. Notice that  $\mathbf{R}_t(x') + \epsilon \mathbf{u}^*$  is a shifted hyperrectangle, which is also a hyperrectangle. Therefore, by Lemma 2, the above conditions are also equivalent to the following:

$$\forall \mathbf{v} \in \text{vert}(\mathbf{R}_t(x)), \forall \mathbf{v}' \in \text{vert}(\mathbf{R}_t(x')) : \mathbf{v} \preceq_C \mathbf{v}' + \epsilon \mathbf{u}^*. \tag{5}$$

Hence, in the implementation of VOGP, to determine if  $x \in \mathcal{S}_t$  is going to be discarded, we check if there exists  $x' \in \mathcal{P}_{\text{pess},t}$  such that Equation (5) holds.

### 4.3 Pareto Identification

VOGP aims to identify the designs that are not dominated by any other design with high probability with respect to the ordering cone  $C$ . It does so by pinpointing designs that, after adding the accuracy vector to the values in their confidence hyperrectangles, remain non-dominated with respect to  $C$  when compared to the values in the confidence hyperrectangle of any other design. The identified designs are moved from the set of undecided designs to the predicted Pareto set. It is important to note that, once a design becomes a member of the predicted Pareto set, it remains a permanent member of that set. The illustration of Pareto identification phase can be seen in Fig. 4 (Right).

---

**Compute:** Set of active designs  $\mathcal{W}_t = \mathcal{S}_t \cup \mathcal{P}_t$ ;  
**for**  $x \in \mathcal{S}_t$  **do**  
    **if**  $\nexists x' \in \mathcal{W}_t : (\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap \mathbf{R}_t(x') \neq \emptyset$  **then**  
         $\mathcal{S}_t = \mathcal{S}_t \setminus \{x\}, \mathcal{P}_t = \mathcal{P}_t \cup \{x\}$ ;  
    **end if**  
**end for**

---

**Algorithm 4** Pareto Identification Subroutine

#### 4.3.1 Computation of the Pareto Identification Subroutine

Let  $t \in \mathbb{N}$ . According to Algorithm 4, to determine if a design  $x \in \mathcal{S}_t$  is going to be added to  $\mathcal{P}_t$ , the condition

$$(\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap \mathbf{R}_t(x') \neq \emptyset$$

needs to be checked for every  $x' \in \mathcal{W}_t$ , which can be rewritten as

$$\exists \mathbf{y}' \in \mathbf{R}_t(x') \exists \mathbf{y} \in \mathbf{R}_t(x) : \mathbf{y} + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}' . \tag{6}$$

We first restate (6) as a convex feasibility problem. Then, in the implementation of VOGP, this convex feasibility problem is solved using the CVXPY library (Agrawal et al., 2018; Diamond & Boyd, 2016).

Note that (6) can be checked by solving the following feasibility problem:

$$\begin{aligned} & \underset{\mathbf{y}, \mathbf{y}'}{\text{minimize}} && 0 \\ & \text{subject to} && \mathbf{y}, \mathbf{y}' \in \mathbb{R}^M, \\ & && y^j \geq \mu_s^j(x) - \beta_s^{1/2} \sigma_s^j(x), \quad \forall j \in [M], \forall s \in [t] \\ & && y^j \leq \mu_s^j(x) + \beta_s^{1/2} \sigma_s^j(x), \quad \forall j \in [M], \forall s \in [t] \\ & && (y')^j \geq \mu_s^j(x') - \beta_s^{1/2} \sigma_s^j(x'), \quad \forall j \in [M], \forall s \in [t] \\ & && (y')^j \leq \mu_s^j(x') + \beta_s^{1/2} \sigma_s^j(x'), \quad \forall j \in [M], \forall s \in [t] \\ & && \mathbf{w}_n^T(\mathbf{y}' - \mathbf{y} - \epsilon \mathbf{u}^*) \geq 0, \quad \forall n \in [N]. \end{aligned}$$

If this problem turns out to be infeasible for all  $x' \in \mathcal{W}_t$ , then  $x$  is added to  $\mathcal{P}_t$ . Otherwise,  $x$  is not added to  $\mathcal{P}_t$ .

### 4.4 Evaluating

VOGP selects the design whose confidence hyperrectangle has the widest diagonal. The diagonal of the hyperrectangle  $R_t(x)$  is given by  $\omega_t(x) = \max_{\mathbf{y}, \mathbf{y}' \in R_t(x)} \|\mathbf{y} - \mathbf{y}'\|_2^2$ . The motivation behind this step is to acquire as much information about the objective space as possible, so that the distinction between Pareto and non-Pareto designs can be made fast with high probability.

---

```

if  $S_t \neq \emptyset$  then
    Select design  $x_t \in \arg \max_{x \in \mathcal{W}_t} \omega_t(x)$ ;
    Observe  $\mathbf{y}_t = \mathbf{f}(x_t) + \boldsymbol{\nu}_t$ ;
end if
    
```

---

**Algorithm 5** Evaluating Subroutine

## 5 Main Theoretical Results and Analysis

In this section, we provide theoretical guarantees for VOGP. We first state the main theorems.

**Theorem 1** *Let  $\eta := \sigma^{-2} / \ln(1 + \sigma^{-2})$ . When VOGP is run, at least with probability  $1 - \delta$ , an  $(\epsilon, \delta)$ -Pareto set can be identified with no more than  $T$  function evaluations, where*

$$T := \min \left\{ t \in \mathbb{N} : \sqrt{\frac{8\beta_t \sigma^2 \eta M \gamma_t}{t}} < \frac{\epsilon}{d_C} \right\}. \tag{7}$$

This theorem provides an implicit bound on the sample complexity in terms of the maximum information gain  $\gamma_t$ . Given a specific kernel, one can numerically solve (7) to compute  $T$ . It is obvious that when the slackness term  $\epsilon$  gets smaller, sample complexity increases. Similarly, as the cone-specific ordering hardness  $d_C$  gets larger, sample complexity increases.

Here, we sketch the idea behind the proof of Theorem 1. We first define  $E$  to be the event that the confidence hyperrectangles of designs include their true objective values and prove that  $\mathbb{P}(E) \geq 1 - \delta$ . Then, we define  $d_C(\cdot)$  and  $A_C(\cdot)$  and show that they have the homogeneity property (see Lemma 4). Using this result, we establish the stopping criterion for VOGP by demonstrating that the algorithm concludes when the diagonals of the confidence hyperrectangles are less than  $\frac{\epsilon}{d_C}$  (see Lemmata 7, 8). Finally, by bounding the diagonals using the maximum information gain and noting their monotonic decrease, we determine the sample complexity. The resulting sample complexity is the smallest number of samples needed for the upper bound of the diagonals to drop below  $\frac{\epsilon}{d_C}$  (see Lemma 9).

The next theorem provides a more explicit bound on the sample complexity for a specific choice of the kernel.

**Theorem 2** *Let the GP kernel  $\mathbf{k}$  have the multiplicative form  $\mathbf{k}(x, x') = [\tilde{k}(x, x') k^*(p, q)]_{p, q \in [M]}$ , where  $\tilde{k} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  is a kernel for the design space and  $k^* : [M] \times [M] \rightarrow \mathbb{R}$  is a kernel for the objective space. Assume that  $\tilde{k}$  is a squared*

exponential kernel. Then, with at least  $1 - \delta$  probability, the sample complexity of VOGP is given by  $T = \tilde{O}\left(\frac{d_C^2}{\epsilon^2}\right)$ .

To prove Theorem 2, we use the bound on maximum information gain for single output GPs given in Vakili et al. (2021) to bound the maximum information gain of an  $M$ -output GP (see Lemma 11). Then, by algebraic manipulations, we establish that the proposed sampling budget satisfies the inequality in (7). The sample complexity of VOGP when using Matérn kernel was also established using a similar approach and can be found in Appendix B. Theorems 1 and 2 provide sample complexity analysis of VOGP in terms of ordering hardness  $d_C$ .

Theorem 1 provides a bound that is applicable in a very general context where the kernel  $k$  can be any bounded positive definite kernel. Theorem 2 provides a more explicit and practical formulation of sample complexity. The following remark compares the established bounds with the bounds in prior work.

**Remark 7** The sample complexity bounds of VOGP established in Theorems 1 and 2 exactly match the sample complexity bounds of previous works in multi-objective Bayesian optimization literature, namely PAL (Zuluaga et al., 2013),  $\epsilon$ -PAL (Zuluaga et al., 2016) in terms of  $\epsilon$ . These works present bounds of  $\mathcal{O}\left(\frac{1}{\epsilon^{2+\rho}}\right)$  for any  $\rho > 0$ . Additionally, some studies focusing on Pareto identification in the multi-objective multi-armed bandit setting, which do not utilize GPs, also provide sample complexity bounds (Ararat, 2023; Auer et al., 2016; Drugan, 2017) that matches those of VOGP up to logarithmic terms with  $\tilde{O}\left(\frac{1}{\epsilon^2}\right)$ . In contrast, other baseline methods in this study that employ GPs—such as JES, MESMO, and EHVI—do not provide sample complexity bounds.

The following two subsections are dedicated to the proofs of Theorems 1 and 2.

## 5.1 Proof of Theorem 1

Let accuracy parameter  $\epsilon > 0$  and confidence parameter  $\delta \in (0, 1)$  be given. Let  $\mathbb{N} := \{1, 2, \dots\}$ .

**Lemma 3** *Let us define the event that confidence hyperrectangles include the true objective values of designs as*

$$E := \left\{ \forall j \in [M] \forall t \in \mathbb{N} \forall x \in \mathcal{X} : |f^j(x) - \mu_t^j(x)| \leq \beta_t^{1/2} \sigma_t^j(x) \right\},$$

where

$$\beta_t := \ln\left(\frac{M\pi^2 |\mathcal{X}| t^2}{3\delta}\right).$$

Then,  $\mathbb{P}(E) \geq 1 - \delta$ .

**Proof** For an event  $E'$ , we denote by  $\mathbb{I}(E')$  its probabilistic indicator function, i.e.,  $\{\mathbb{I}(E') = 1\} = E'$  and  $\{\mathbb{I}(E') = 0\} = \Omega \setminus E'$ , where  $\Omega$  is the underlying sample space. Note that

$$\begin{aligned}
 1 - \mathbb{P}(E) &= \mathbb{E} \left[ \mathbb{I} \left( \left\{ \exists j \in [M] \exists t \in \mathbb{N} \exists x \in \mathcal{X} : |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \right) \right] \\
 &\leq \mathbb{E} \left[ \sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} \mathbb{I} \left( \left\{ |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \right) \right] \tag{8} \\
 &= \sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} \mathbb{E} \left[ \mathbb{E} \left[ \mathbb{I} \left( \left\{ |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \right) \mid \mathbf{y}_{[t-1]} \right] \right] \\
 &= \sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} \mathbb{E} \left[ \mathbb{P} \left( \left\{ |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \mid \mathbf{y}_{[t-1]} \right) \right] \\
 &\leq \sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} 2e^{-\beta_t/2} \\
 &= 2M|\mathcal{X}| \sum_{t=1}^{\infty} e^{-\beta_t/2} \tag{9} \\
 &= 2M|\mathcal{X}| \sum_{t=1}^{\infty} \left( \frac{M\pi^2|\mathcal{X}|t^2}{3\delta} \right)^{-1} \\
 &= \frac{6\delta}{\pi^2} \sum_{t=1}^{\infty} \frac{1}{t^2} = \delta,
 \end{aligned}$$

where (8) uses the tower rule and linearity of expectation and (9) uses Gaussian tail bound; here note that, given  $\mathbf{y}_{[t-1]}$ , the conditional distribution of  $f^j(x)$  is  $\mathcal{N}(\mu_t^j(x), \sigma_t^j(x))$ .  $\square$

For each  $s > 0$ , let us introduce the set  $A_C(s)$ , which consists of vectors that dominate all points on a sphere of radius  $s$  centered at the origin. We define  $d_C(s)$  as the infimum of the norms of the elements in this set. These concepts are formally expressed as follows:

$$A_C(s) := \bigcap_{\mathbf{u} \in \mathbb{S}^{M-1}} (s\mathbf{u} + C), \quad d_C(s) := \inf\{\|\mathbf{z}\|_2 \mid \mathbf{z} \in A_C(s)\}.$$

When  $s = 1$ , we recover the quantities in Definition 4, since  $d_C(1) = d_C$ ,  $A_C(1) = A_C$ .

**Lemma 4** *Let  $s > 0$ . Then,*

$$A_C(s) = \left\{ \mathbf{z} \in \mathbb{R}^M \mid \mathbf{w}_n^\top \mathbf{z} \geq s, \forall n \in [N] \right\}.$$

*In particular,  $A_C(s) = sA_C(1) := \{s\mathbf{y} \mid \mathbf{y} \in A_C(1)\}$  and  $d_C(s) = sd_C(1)$ . Moreover, there exists a unique  $\mathbf{z}^s \in A_C(s)$  such that  $d_C(s) = \|\mathbf{z}^s\|_2$ , where the superscript denotes dependence on  $s$  (not exponentiation).*

**Proof** Let  $s > 0$ . Then, it can be shown that

$$\begin{aligned}
 \mathbf{y} \in A_C(s) &\Leftrightarrow \mathbf{y} \in s\mathbf{u} + C, \forall \mathbf{u} \in \mathbb{S}^{M-1} \\
 &\Leftrightarrow \mathbf{y} - s\mathbf{u} \in C, \forall \mathbf{u} \in \mathbb{S}^{M-1} \\
 &\Leftrightarrow \mathbf{w}_n^\top (\mathbf{y} - s\mathbf{u}) \geq 0, \forall n \in [N], \forall \mathbf{u} \in \mathbb{S}^{M-1}
 \end{aligned}
 \tag{10}$$

$$\begin{aligned}
 &\Leftrightarrow \mathbf{w}_n^\top \mathbf{y} \geq s\mathbf{w}_n^\top \mathbf{u}, \forall n \in [N], \forall \mathbf{u} \in \mathbb{S}^{M-1} \\
 &\Leftrightarrow \mathbf{w}_n^\top \mathbf{y} \geq s \sup_{\mathbf{u} \in \mathbb{S}^{M-1}} (\mathbf{w}_n^\top \mathbf{u}), \forall n \in [N],
 \end{aligned}
 \tag{11}$$

where (10) follows from (1). Let  $n \in [N]$ . By the definition of  $\mathbf{w}_n$ , we have

$$\sup_{\mathbf{u} \in \mathbb{S}^{M-1}} \mathbf{w}_n^\top \mathbf{u} = \sup_{\mathbf{u} \in \mathbb{B}(1)} \mathbf{w}_n^\top \mathbf{u} = \|\mathbf{w}_n\|_2 = 1.
 \tag{12}$$

Combining (11) and (12), we get

$$\mathbf{y} \in A_C(s) \Leftrightarrow \mathbf{w}_n^\top \mathbf{y} \geq s, \forall n \in [N].$$

Therefore,  $A_C(s) = \{ \mathbf{z} \in \mathbb{R}^M \mid \mathbf{w}_n^\top \mathbf{z} \geq s, \forall n \in [N] \}$ , which implies that  $A_C(s) = sA_C(1)$  and hence  $d_C(s) = sd_C(1)$ .

The existence and uniqueness of  $\mathbf{z}^s$  is a direct consequence of the strict convexity and continuity of the  $\ell_2$ -norm  $\|\cdot\|_2$  together with the closedness and convexity of  $A_C(s)$ .  $\square$

**Remark 8** The last part of Lemma 4 justifies the definition of  $\mathbf{u}^*$  (see Definition 4).

In what follows, we denote by  $\mathbb{S}^{M-1}(r)$  the boundary of  $\mathbb{B}(r)$ , where  $r > 0$ ; in particular,  $\mathbb{S}^{M-1}(1) = \mathbb{S}^{M-1}$  is the unit sphere. We next prove Lemma 1 as a consequence of Lemma 4.

**Proof of Lemma 1 (i)** Suppose that  $\mathbf{y} + \tilde{\mathbf{p}} \preccurlyeq_C \mathbf{z}$ . Equivalently, we have  $\mathbf{z} \in \mathbf{y} + \tilde{\mathbf{p}} + C$ . Then, by (1), we have

$$\forall n \in [N] : \mathbf{w}_n^\top (\mathbf{z} - \mathbf{y} - \tilde{\mathbf{p}}) \geq 0.
 \tag{13}$$

Since  $\tilde{\mathbf{p}} \in \mathbb{B}(\frac{\epsilon}{d_C})$ , by Cauchy-Schwarz inequality, we have

$$|\mathbf{w}_n^\top \tilde{\mathbf{p}}| \leq \|\mathbf{w}_n\|_2 \|\tilde{\mathbf{p}}\|_2 \leq \frac{\epsilon}{d_C}
 \tag{14}$$

for every  $n \in [N]$ . Moreover, by Lemma 4 and Definition 4, we have

$$\mathbf{w}_n^\top \mathbf{u}^* = \frac{\mathbf{w}_n^\top \mathbf{z}^*}{d_C} \geq \frac{1}{d_C}
 \tag{15}$$

for every  $n \in [N]$ . Note that  $\mathbf{w}_n^T \mathbf{z}^* \geq 1$  follows from the definition of  $A_C$  and that  $\mathbf{z}^* \in A_C$ . Hence, by (13), (14), and (15), we get

$$\mathbf{w}_n^T (\mathbf{z} + \epsilon \mathbf{u}^* - \mathbf{y}) \geq \mathbf{w}_n^T \tilde{\mathbf{p}} + \epsilon \mathbf{w}_n^T \mathbf{u}^* \geq -\frac{\epsilon}{d_C} + \frac{\epsilon}{d_C} \geq 0$$

for every  $n \in [N]$ . This shows that  $\mathbf{z} + \epsilon \mathbf{u}^* - \mathbf{y} \in C$ , i.e.,  $\mathbf{y} \preceq_C \mathbf{z} + \epsilon \mathbf{u}^*$ .

(ii) Suppose that  $\tilde{\mathbf{p}} \in \text{int}(\mathbb{B}(\frac{\epsilon}{d_C}))$  and  $\mathbf{y} + \tilde{\mathbf{p}} \preceq_C \mathbf{z}$ . In particular, (13) and (15) still hold. Since  $\tilde{\mathbf{p}} \in \text{int}(\mathbb{B}(\frac{\epsilon}{d_C}))$ , by Cauchy-Schwarz inequality, we have

$$|\mathbf{w}_n^T \tilde{\mathbf{p}}| \leq \|\mathbf{w}_n\|_2 \|\tilde{\mathbf{p}}\|_2 < \frac{\epsilon}{d_C} \tag{16}$$

for every  $n \in [N]$ . Hence, by (13), (15), and (16), we get

$$\mathbf{w}_n^T (\mathbf{z} + \epsilon \mathbf{u}^* - \mathbf{y}) \geq \mathbf{w}_n^T \tilde{\mathbf{p}} + \epsilon \mathbf{w}_n^T \mathbf{u}^* > -\frac{\epsilon}{d_C} + \frac{\epsilon}{d_C} \geq 0$$

for every  $n \in [N]$ . This shows that  $\mathbf{z} + \epsilon \mathbf{u}^* - \mathbf{y} \in \text{int}(C)$ , i.e.,  $\mathbf{y} \prec_C \mathbf{z} + \epsilon \mathbf{u}^*$ .  $\square$

In the following lemma, we show that the predicted Pareto set returned by VOGP satisfies the first success condition. The proof establishes that under event  $E$ , every Pareto optimal design  $x \in P^*$  has a corresponding design  $z \in \hat{\mathcal{P}}$  such that  $\mathbf{f}(x) \preceq_C \mathbf{f}(z) + \epsilon \mathbf{u}^*$ . If  $x$  is not directly in  $\hat{\mathcal{P}}$ , it must have been discarded at some round. In that case, the discarding subroutine ensures the existence of a design  $z_1$  in the pessimistic Pareto set that caused  $x$  to be discarded. If  $z_1$  is in  $\hat{\mathcal{P}}$ , the success condition holds. Otherwise,  $z_1$  must have been excluded from the pessimistic Pareto set by another design  $z_2$ . If  $z_2 \in \hat{\mathcal{P}}$ , the success condition holds and if it does not, some design  $z_3$  must have excluded  $z_2$  from the pessimistic Pareto set, continuing this process iteratively. The chain eventually leads to a design that is included in  $\hat{\mathcal{P}}$ , ensuring that the success condition holds.

**Lemma 5** *Under event  $E$ , the set  $\hat{\mathcal{P}}$  returned by VOGP satisfies condition (i) in Definition 1.*

**Proof** We assume that event  $E$  occurs and claim that for every  $x \in P^*$ , there exists  $z \in \hat{\mathcal{P}}$  such that  $\mathbf{f}(x) \preceq_C \mathbf{f}(z) + \epsilon \mathbf{u}^*$ . If  $x \in \hat{\mathcal{P}}$ , then the claim holds with  $z = x$ , i.e.,  $\mathbf{f}(x) \preceq_C \mathbf{f}(x) + \epsilon \mathbf{u}^*$ , since  $\epsilon \mathbf{u}^* \in C$ . Suppose that  $x \notin \hat{\mathcal{P}}$ . Then,  $x$  must have been discarded at some round  $s_1$ . By the discarding rule, there exists  $z_1 \in \mathcal{P}_{\text{pess}, s_1}$  such that

$$\mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C \tag{17}$$

holds for every  $\mathbf{y} \in \mathbf{R}_{s_1}(x)$ .

At each round, the initial confidence hyperrectangle of an arbitrary design  $x' \in \mathcal{X}$  given by (4) can be expressed as

$$\mathbf{Q}_t(x') = \left\{ \mathbf{y} \in \mathbb{R}^M \mid \boldsymbol{\mu}_t(x') - \beta_t^{1/2} \boldsymbol{\sigma}_t(x') \preceq_{\mathbb{R}^M} \mathbf{y} \preceq_{\mathbb{R}^M} \boldsymbol{\mu}_t(x') + \beta_t^{1/2} \boldsymbol{\sigma}_t(x') \right\} \tag{18}$$

and the initial confidence hyperrectangle is intersected with previous hyperrectangles to obtain the confidence hyperrectangle of the current round, that is,

$$R_t(x') = R_{t-1}(x') \cap Q_t(x') . \tag{19}$$

It can be checked that, due to Lemma 3, we have  $f(x') \in R_{s_1}(x')$  for every  $x' \in \mathcal{X}$  under event  $E$ . In particular, by (17), we have

$$R_{s_1}(z_1) + \epsilon u^* \subseteq f(x) + C , \tag{20}$$

where  $z_1 \in \mathcal{P}_{\text{pess},s_1}$ .

Since  $f(z_1) \in R_{s_1}(z_1)$ , (20) implies that  $f(z_1) + \epsilon u^* \in f(x) + C$ , which is equivalent to  $f(x) \preceq_C f(z_1) + \epsilon u^*$ . Therefore, if  $z_1 \in \hat{\mathcal{P}}$ , then the claim holds by choosing  $z = z_1$ .

If  $z_1 \notin \hat{\mathcal{P}}$ , then it must have been discarded at some round  $s_2 \geq s_1$ . Because VOGP discards from the set  $\mathcal{S}_t \setminus \mathcal{P}_{\text{pess},t}$  at any round  $t$ , we have  $z_1 \notin \mathcal{P}_{\text{pess},s_2}$ . Then, using the definition of the pessimistic Pareto set (see Definition 3), there exists  $z_2 \in \mathcal{A}_{s_2}$  such that  $R_{s_2}(z_2) + C \subsetneq R_{s_2}(z_1) + C$ , which implies that

$$R_{s_2}(z_2) \subseteq R_{s_2}(z_2) + C \subseteq R_{s_2}(z_1) + C . \tag{21}$$

To proceed, we use (20) and the fact that  $C + C = C$  to obtain

$$R_{s_1}(z_1) + \epsilon u^* + C \subseteq f(x) + C + C = f(x) + C . \tag{22}$$

In addition, (21) implies that

$$R_{s_2}(z_2) + \epsilon u^* \subseteq R_{s_1}(z_1) + \epsilon u^* + C .$$

Combining the above display with (22) yields

$$R_{s_2}(z_2) + \epsilon u^* \subseteq R_{s_1}(z_1) + \epsilon u^* + C \subseteq f(x) + C .$$

According to Lemma 3, under event  $E$ , we have  $f(z_2) \in R_{s_2}(z_2)$ . Hence, it holds that

$$f(z_2) + \epsilon u^* \in f(x) + C .$$

So, if  $z_2 \in \hat{\mathcal{P}}$ , then the claim holds with  $z = z_2$ .

If  $z_2 \notin \hat{\mathcal{P}}$ , then  $z_2$  must have been discarded at some round  $s_3 \geq s_2$ . Because VOGP discards from the set  $\mathcal{S}_t \setminus \mathcal{P}_{\text{pess},t}$  at any round  $t$ , we have  $z_2 \notin \mathcal{P}_{\text{pess},s_3}$ . Then, using the definition of the pessimistic set, there exists  $z_3 \in \mathcal{A}_{s_3}$  such that  $R_{s_3}(z_3) + C \subsetneq R_{s_3}(z_2) + C$ , which implies that

$$R_{s_3}(z_3) \subseteq R_{s_3}(z_3) + C \subseteq R_{s_3}(z_2) + C \subseteq R_{s_2}(z_2) + C , \tag{23}$$

where the last inclusion follows from the fact that  $R_{s_3}(z_2) \subseteq R_{s_2}(z_2)$  as  $s_3 \geq s_2$ . Then, combining (21) and (23) and using the fact that  $R_{s_2}(z_1) \subseteq R_{s_1}(z_1)$  yield

$$\mathbf{R}_{s_3}(z_3) \subseteq \mathbf{R}_{s_2}(z_2) + C \subseteq \mathbf{R}_{s_2}(z_1) + C + C = \mathbf{R}_{s_2}(z_1) + C \subseteq \mathbf{R}_{s_1}(z_1) + C.$$

Then, by (22), we get

$$\mathbf{R}_{s_3}(z_3) + \epsilon \mathbf{u}^* \subseteq \mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* + C \subseteq \mathbf{f}(x) + C.$$

Since  $\mathbf{f}(z_3) \in \mathbf{R}_{s_3}(z_3)$  under event  $E$ , we have

$$\mathbf{f}(z_3) + \epsilon \mathbf{u}^* \in \mathbf{f}(x) + C.$$

So, if  $z_3 \in \hat{\mathcal{P}}$ , then the claim holds with  $z = z_3$ .

If  $z_3 \notin \hat{\mathcal{P}}$ , then a similar argument can be made until  $z_n \in \hat{\mathcal{P}}$ , where  $n \leq |\mathcal{X}|$ . Indeed, in the worst case, there exists some  $z_n \in \mathcal{P}_{\text{pess}, t_s}$  and it must be the case that  $z_n \in \hat{\mathcal{P}}$  since the algorithm terminates at round  $t_s$ . Then, the claim holds with  $z = z_n$ .  $\square$

The next lemma shows that if a design  $x$  has a gap  $\Delta_x^* > 2\epsilon$ , then it cannot be in the predicted Pareto set  $\hat{\mathcal{P}}$ . To show this by contradiction, suppose such an  $x$  is included in  $\hat{\mathcal{P}}$ . By the definition of  $\Delta_{x'}^*$ , there must exist a Pareto optimal design  $x' \in P^*$  such that  $m(x, x') > 2\epsilon$ . If  $x'$  is active when  $x$  is added to  $\hat{\mathcal{P}}$ , Pareto identification rule (Equation 6) leads to a contradiction. If  $x'$  is not active, we consider two further cases: (i) if the design that discarded  $x'$  was active when  $x$  was added to  $\hat{\mathcal{P}}$ , the Pareto identification rule leads to a contradiction again; (ii) if the mentioned design was not active, it must have been discarded and thus removed from the pessimistic Pareto set by another design. In this case, this process continues iteratively, until the final design in this sequence is active at the round at which  $x$  is added to  $\hat{\mathcal{P}}$ , again leading to a contradiction via the Pareto identification rule.

**Lemma 6** *Under event  $E$ , the set  $\hat{\mathcal{P}}$  returned by VOGP satisfies condition (ii) in Definition 1.*

**Proof** We will show that if  $\Delta_x^* > 2\epsilon$ , then  $x \notin \hat{\mathcal{P}} \setminus P^*$ . To prove this by contradiction, suppose that  $\Delta_x^* > 2\epsilon$  for a design  $x \in \hat{\mathcal{P}} \setminus P^*$ . By definition of  $m(x, x')$ , this means that there exists  $x' \in P^*$  such that  $m(x, x') > 2\epsilon$ . Since  $x \in \hat{\mathcal{P}}$ , it must have been added to  $\hat{\mathcal{P}}$  at some round  $t$ . By the  $\epsilon$ -covering rule of VOGP, that means for all  $z \in \mathcal{W}_t$ ,  $(\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap (\mathbf{R}_t(z) - C) = \emptyset$ .

We complete the proof by considering two cases:

- **Case 1:**  $x' \in \mathcal{W}_t$ .
- **Case 2:**  $x' \notin \mathcal{W}_t$ . **Case 1:** If  $x' \in \mathcal{W}_t$ , then

$$(\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap (\mathbf{R}_t(x') - C) = \emptyset. \tag{24}$$

By the properties of the Minkowski sum, we have  $\mathbf{R}_t(x) + \epsilon \mathbf{u}^* \subseteq \mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C$ . This together with (24) results in

$$(\mathbf{R}_t(x) + \epsilon \mathbf{u}^*) \cap (\mathbf{R}_t(x') - C) = \emptyset. \tag{25}$$

According to Lemma 3, under the good event  $E$ ,  $\mathbf{f}(x) \in \mathbf{R}_t(x)$  and  $\mathbf{f}(x') \in \mathbf{R}_t(x')$ . Combining this with (25), we conclude that

$$\mathbf{f}(x) + \epsilon \mathbf{u}^* \notin \mathbf{f}(x') - C. \tag{26}$$

Because  $\mathbf{f}(x) + \epsilon \mathbf{u}^* \notin \mathbf{f}(x') - \text{int}(C)$ , by the definition of  $m(x, x')$ ,  $m(x, x') \leq \epsilon$  and we get a contradiction for the case of  $x' \in \mathcal{W}_t$ .

**Case 2:** If  $x' \notin \mathcal{W}_t$ , it must have been discarded at an earlier round  $s_1 < t$ . By the discarding rule,  $\exists z_1 \in \mathcal{P}_{\text{pess}, s_1}$  such that

$$\mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C, \forall \mathbf{y} \in \mathbf{R}_{s_1}(x'). \tag{27}$$

We proceed in Case 2, by considering the following two cases based on the status of  $z_1$ :

- **Case 2.1:**  $z_1 \in \mathcal{W}_t$ .
- **Case 2.2:**  $z_1 \notin \mathcal{W}_t$ . **Case 2.1:** If  $z_1 \in \mathcal{W}_t$ , since  $\mathbf{R}_t(z_1) \subseteq \mathbf{R}_{s_1}(z_1)$ , (27) also implies that

$$\begin{aligned} & \mathbf{R}_t(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C, \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') \\ \iff & \forall \mathbf{y}_{x', s_1} \in \mathbf{R}_{s_1}(x'), \forall \mathbf{y}_{z_1, t} \in \mathbf{R}_t(z_1) : \mathbf{y}_{z_1, t} \in \mathbf{y}_{x', s_1} - \epsilon \mathbf{u}^* + C \\ \iff & \forall \mathbf{y}_{x', s_1} \in \mathbf{R}_{s_1}(x'), \forall \mathbf{y}_{z_1, t} \in \mathbf{R}_t(z_1) : \mathbf{y}_{x', s_1} - \epsilon \mathbf{u}^* \preccurlyeq_C \mathbf{y}_{z_1, t}. \end{aligned} \tag{28}$$

Since  $z_1 \in \mathcal{W}_t$ , the  $\epsilon$ -covering rule at round  $t$  between  $x$  and  $z_1$  pairs should hold. Combined with the fact that  $\mathbf{R}_t(x) + \epsilon \mathbf{u}^* \subseteq \mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C$ , it holds that

$$(\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap (\mathbf{R}_t(z_1) - C) = \emptyset \tag{29}$$

$$\begin{aligned} & \implies (\mathbf{R}_t(x) + \epsilon \mathbf{u}^*) \cap (\mathbf{R}_t(z_1) - C) = \emptyset \\ \iff & \forall \mathbf{y}_{z_1, t} \in \mathbf{R}_t(z_1), \forall \mathbf{y}_{x, t} \in \mathbf{R}_t(x) : \mathbf{y}_{x, t} + \epsilon \mathbf{u}^* \not\preccurlyeq_C \mathbf{y}_{z_1, t}. \end{aligned} \tag{30}$$

Then, by combining (28) and (30), we get  $\mathbf{y}_{x, t} + \epsilon \mathbf{u}^* \not\preccurlyeq_C \mathbf{y}_{x', s_1} - \epsilon \mathbf{u}^*, \forall \mathbf{y}_{x, t} \in \mathbf{R}_t(x)$  and  $\forall \mathbf{y}_{x', s_1} \in \mathbf{R}_{s_1}(x')$ . Therefore, according to Lemma 3, under the good event  $\mathcal{F}_1$  it holds that  $\mathbf{f}(x) + 2\epsilon \mathbf{u}^* \not\preccurlyeq_C \mathbf{f}(x')$ . Since  $2\epsilon \mathbf{u}^* \in \mathbb{B}(2\epsilon) \cap C$ , by the definition of  $m(x, x')$ ,  $m(x, x') \leq 2\epsilon$  which is a contradiction.

**Case 2.2:** Next, we examine the case  $z_1 \notin \mathcal{W}_t$ . Particularly, consider the collection of designs denoted by  $z_1, \dots, z_{n-1}, z_n$  where  $z_i$  has been discarded at some round  $s_{i+1}$  by being removed from  $\mathcal{P}_{\text{pess}, s_{i+1}}$  by  $z_{i+1}$ , as they fulfill the condition  $\mathbf{R}_{s_{i+1}}(z_{i+1}) + C \subsetneq \mathbf{R}_{s_{i+1}}(z_i) + C$ . Notice that this implies  $\mathbf{R}_{s_{i+1}}(z_{i+1}) \subseteq \mathbf{R}_{s_{i+1}}(z_i) + C$ . Assume that  $z_n \in \mathcal{W}_t$ . Notice that it's always possible to identify such a design, because of the way the sequence was defined. Together with (19), we observe that the following set relations hold.

$$\begin{aligned}
 \mathbf{R}_{s_2}(z_1) + C &\supseteq \mathbf{R}_{s_2}(z_2) \supseteq \mathbf{R}_{s_3}(z_2) \\
 \mathbf{R}_{s_3}(z_2) + C &\supseteq \mathbf{R}_{s_3}(z_3) \supseteq \mathbf{R}_{s_4}(z_3) \\
 &\vdots \\
 \mathbf{R}_{s_{n-1}}(z_{n-2}) + C &\supseteq \mathbf{R}_{s_{n-1}}(z_{n-1}) \supseteq \mathbf{R}_{s_n}(z_{n-1}) \\
 \mathbf{R}_{s_n}(z_{n-1}) + C &\supseteq \mathbf{R}_{s_n}(z_n) \supseteq \mathbf{R}_t(z_n) .
 \end{aligned}
 \tag{31}$$

In particular, (31) holds since  $t \geq s_n$ .

By the definition of Minkowski sum and the convexity of  $C$ , for any  $i \in \{2, \dots, n - 1\}$ , we have

$$\begin{aligned}
 \mathbf{R}_{s_i}(z_{i-1}) + C \supseteq \mathbf{R}_{s_{i+1}}(z_i) &\implies \mathbf{R}_{s_i}(z_{i-1}) + C + C \supseteq \mathbf{R}_{s_{i+1}}(z_i) + C \\
 &\implies \mathbf{R}_{s_i}(z_{i-1}) + C \supseteq \mathbf{R}_{s_{i+1}}(z_i) + C .
 \end{aligned}$$

Hence, it holds that

$$\begin{aligned}
 \mathbf{R}_{s_2}(z_1) + C &\supseteq \mathbf{R}_{s_3}(z_2) + C \\
 \mathbf{R}_{s_3}(z_2) + C &\supseteq \mathbf{R}_{s_4}(z_3) + C \\
 &\vdots \\
 \mathbf{R}_{s_{n-1}}(z_{n-2}) + C &\supseteq \mathbf{R}_{s_n}(z_{n-1}) + C \\
 \implies \mathbf{R}_{s_2}(z_1) + C &\supseteq \mathbf{R}_{s_n}(z_{n-1}) + C .
 \end{aligned}
 \tag{32}$$

Using the fact that  $\mathbf{R}_{s_1}(z_1) \supseteq \mathbf{R}_{s_2}(z_1)$  when  $s_1 \leq s_2$ , (32), and (31) in order, we obtain

$$\mathbf{R}_{s_1}(z_1) + C \supseteq \mathbf{R}_{s_2}(z_1) + C \supseteq \mathbf{R}_{s_n}(z_{n-1}) + C \supseteq \mathbf{R}_t(z_n) .
 \tag{33}$$

Next, by combining (27) and (33), and using the properties of Minkowski sum, we have

$$\begin{aligned}
 \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C &\iff \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_{s_1}(z_1) \subseteq \mathbf{y} - \epsilon \mathbf{u}^* + C \\
 \implies \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_{s_1}(z_1) + C \subseteq \mathbf{y} - \epsilon \mathbf{u}^* + C & \\
 \implies \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_t(z_n) \subseteq \mathbf{y} - \epsilon \mathbf{u}^* + C . &
 \end{aligned}
 \tag{34}$$

Alternatively, (34) can be re-written as

$$\forall \mathbf{y}_{x',s_1} \in \mathbf{R}_{s_1}(x'), \forall \mathbf{y}_{z_n,t} \in \mathbf{R}_t(z_n) : \mathbf{y}_{x',s_1} - \epsilon \mathbf{u}^* \preccurlyeq_C \mathbf{y}_{z_n,t} .
 \tag{35}$$

Since  $z_n \in \mathcal{W}_t$ , the  $\epsilon$ -covering rule at round  $t$  between  $x$  and  $z_n$  pairs should hold. Combined with  $\mathbf{R}_t(x) + \epsilon \mathbf{u}^* \subseteq \mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C$ , it holds that

$$(\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap (\mathbf{R}_t(z_n) - C) = \emptyset
 \tag{36}$$

$$\begin{aligned}
 \implies (\mathbf{R}_t(x) + \epsilon \mathbf{u}^*) \cap (\mathbf{R}_t(z_n) - C) &= \emptyset \\
 \implies \forall \mathbf{y}_{z_n,t} \in \mathbf{R}_t(z_n), \forall \mathbf{y}_{x,t} \in \mathbf{R}_t(x) : \mathbf{y}_{x,t} + \epsilon \mathbf{u}^* &\not\prec_C \mathbf{y}_{z_n,t} .
 \end{aligned}
 \tag{37}$$

Then, by combining (35) and (37), we get  $\mathbf{y}_{x,t} + \epsilon \mathbf{u}^* \not\prec_C \mathbf{y}_{x',s_1} - \epsilon \mathbf{u}^*$ ,  $\forall \mathbf{y}_{x,t} \in \mathbf{R}_t(x)$  and  $\forall \mathbf{y}_{x',s_1} \in \mathbf{R}_{s_1}(x')$ . Then, according to Lemma 3, under the good event  $E$  it holds that  $\mathbf{f}(x) + 2\epsilon \mathbf{u}^* \not\prec_C \mathbf{f}(x')$ . Since  $2\epsilon \mathbf{u}^* \in \mathbb{B}(2\epsilon) \cap C$ , by the definition of  $m(x, x')$ ,  $m(x, x') \leq 2\epsilon$  which is a contradiction.  $\square$

The following two lemmas show that, as confidence regions shrink over rounds, a point is reached where domination relations between any two points within a confidence hyperrectangle hold when  $\epsilon \mathbf{u}^*$  help is included. This ensures that the algorithm can make all the necessary decisions, leading to eventual termination. Specifically, we prove that given that all the confidence hyperrectangles have shrunk below a cone dependent threshold, at that round, any design that is not to be added in the predicted Pareto set must have been discarded. Consequently, the algorithm terminates.

**Lemma 7** *Let  $t \in \mathbb{N}$  and suppose that  $\bar{\omega}_t = \max_{x \in \mathcal{W}_t} \omega_t(x) < \frac{\epsilon}{d_C}$ . Then, for every  $x \in \mathcal{W}_t$  and  $\mathbf{y}, \mathbf{y}' \in \mathbf{R}_t(x)$ , it holds*

$$\mathbf{y} \prec_C \mathbf{y}' + \epsilon \mathbf{u}^* .$$

**Proof** Let  $x \in \mathcal{W}_t$  and  $\mathbf{y}, \mathbf{y}' \in \mathbf{R}_t(x)$ . By the definition of  $\omega_t(x)$ , we have

$$\tilde{\mathbf{p}} := \mathbf{y}' - \mathbf{y} \in \text{int} \left( \mathbb{B} \left( \frac{\epsilon}{d_C} \right) \right) .$$

Moreover, we have  $\mathbf{y} + \tilde{\mathbf{p}} = \mathbf{y}' \preceq_C \mathbf{y}'$  trivially. Hence, by Lemma 1(ii), we directly get  $\mathbf{y} \prec_C \mathbf{y}' + \epsilon \mathbf{u}^*$ .  $\square$

**Lemma 8** *Let  $t \in \mathbb{N}$  and suppose that  $\bar{\omega}_t < \frac{\epsilon}{d_C(I)}$ . Then, VOGP terminates at round  $t$ .*

**Proof** Let  $\mathcal{S}_{t,0}(\mathcal{P}_{t,0})$ ,  $\mathcal{S}_{t,1}(\mathcal{P}_{t,1})$ , and  $\mathcal{S}_{t,2}(\mathcal{P}_{t,2})$  denote the set  $\mathcal{S}_t(\mathcal{P}_t)$  at the end of modeling, discarding and Pareto identification phases, respectively. For a node  $x \in \mathcal{S}_{t,0} \setminus \mathcal{P}_{t,2}$ , if  $x \notin \mathcal{S}_{t,1}$ , then it must have been discarded at round  $t$ . Hence, VOGP terminates at round  $t$  if and only if every  $x \in \mathcal{S}_{t,0}$  is either discarded or moved to  $\mathcal{P}_{t,2}$ . In order to prove this, we show that if  $x \in \mathcal{S}_{t,0} \setminus \mathcal{P}_{t,2}$  holds, then  $x \notin \mathcal{S}_{t,1}$ . To prove this by contradiction, let  $x \in \mathcal{S}_{t,0} \setminus \mathcal{P}_{t,2}$  and assume that  $x \in \mathcal{S}_{t,1}$ . Since  $x$  has not been added to  $\mathcal{P}_{t,2}$ , according to the  $\epsilon$ -covering rule in Algorithm 4, there exists some  $z^* \in \mathcal{P}_{t,1} \cup \mathcal{S}_{t,1}$  such that  $(\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap \mathbf{R}_t(z^*) \neq \emptyset$ . Hence, there exist  $\mathbf{y}_{z^*} \in \mathbf{R}_t(z^*)$  and  $\mathbf{y}_x \in \mathbf{R}_t(x)$  such that

$$\mathbf{y}_x + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z^*} . \tag{38}$$

Since we assume that  $x \in \mathcal{S}_{t,1} \subseteq \mathcal{P}_{t,1} \cup \mathcal{S}_{t,1} = \mathcal{W}_t$ , by Lemma 7, we have

$$\forall \tilde{\mathbf{y}}_x \in \mathbf{R}_t(x) : \tilde{\mathbf{y}}_x \prec_C \mathbf{y}_x + \epsilon \mathbf{u}^* . \tag{39}$$

Again, by Lemma 7, since  $z^* \in \mathcal{W}_t$ , we have

$$\forall \tilde{\mathbf{y}}_{z^*} \in \mathbf{R}_t(z^*) : \mathbf{y}_{z^*} \prec_C \tilde{\mathbf{y}}_{z^*} + \epsilon \mathbf{u}^* . \tag{40}$$

Then, using (38) using (40), we have

$$\forall \tilde{\mathbf{y}}_{z^*} \in \mathbf{R}_t(z^*) : \mathbf{y}_x + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z^*} \prec_C \tilde{\mathbf{y}}_{z^*} + \epsilon \mathbf{u}^*, \tag{41}$$

which implies that

$$\forall \tilde{\mathbf{y}}_{z^*} \in \mathbf{R}_t(z^*) : \mathbf{y}_x \prec_C \tilde{\mathbf{y}}_{z^*}$$

so that

$$\mathbf{R}_t(z^*) \subseteq \mathbf{y}_x + \text{int}(C),$$

which then implies that

$$\mathbf{R}_t(z^*) + C \subseteq \mathbf{y}_x + \text{int}(C) + C = \mathbf{y}_x + \text{int}(C).$$

In particular, we have  $\mathbf{R}_t(z^*) + C \subsetneq \mathbf{y}_x + C$ . Together with  $\mathbf{y}_x + C \subseteq \mathbf{R}_t(x) + C$ , we get

$$\mathbf{R}_t(z^*) + C \subsetneq \mathbf{R}_t(x) + C. \tag{42}$$

On the other hand, using (38) and (39), we have

$$\forall \tilde{\mathbf{y}}_x \in \mathbf{R}_t(x) : \tilde{\mathbf{y}}_x \prec_C \mathbf{y}_x + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z^*},$$

which implies that

$$\forall \tilde{\mathbf{y}}_x \in \mathbf{R}_t(x) : \tilde{\mathbf{y}}_x \prec_C \mathbf{y}_{z^*}. \tag{43}$$

Then, by (43) and (40), we obtain

$$\forall \tilde{\mathbf{y}}_x \in \mathbf{R}_t(x) \forall \tilde{\mathbf{y}}_{z^*} \in \mathbf{R}_t(z^*) : \tilde{\mathbf{y}}_x \prec_C \mathbf{y}_{z^*} \prec_C \tilde{\mathbf{y}}_{z^*} + \epsilon \mathbf{u}^*,$$

which implies that

$$\forall \tilde{\mathbf{y}}_x \in \mathbf{R}_t(x) \forall \tilde{\mathbf{y}}_{z^*} \in \mathbf{R}_t(z^*) : \tilde{\mathbf{y}}_x \preceq_C \tilde{\mathbf{y}}_{z^*} + \epsilon \mathbf{u}^*$$

so that

$$\forall \tilde{\mathbf{y}}_x \in \mathbf{R}_t(x) : \mathbf{R}_t(z^*) + \epsilon \mathbf{u}^* \subseteq \tilde{\mathbf{y}}_x + C. \tag{44}$$

Then, by (42), we obtain  $x \notin \mathcal{P}_{\text{pess},t}$ . If  $z^* \in \mathcal{P}_{\text{pess},t}$ , then (44) shows that  $x$  must be discarded. If  $z^* \notin \mathcal{P}_{\text{pess},t}$ , then  $\exists z' \in \mathcal{A}_t$  such that

$$\forall \mathbf{y}_{z'} \in \mathbf{R}_t(z'), \exists \mathbf{y}_{z''} \in \mathbf{R}_t(z^*) : \mathbf{y}_{z''} \preceq_C \mathbf{y}_{z'} \tag{45}$$

$$\iff \forall \mathbf{y}_{z'} \in \mathbf{R}_t(z'), \exists \mathbf{y}_{z''} \in \mathbf{R}_t(z^*) : \mathbf{y}_{z''} + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z'} + \epsilon \mathbf{u}^*. \tag{46}$$

Fix  $\mathbf{y}_{z''}$  as given in (45). Starting from (44) and using (46), we have

$$\mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_{z''} + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z'} + \epsilon \mathbf{u}^*, \forall \mathbf{y}_{z'} \in \mathbf{R}_t(z') \text{ and } \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) \tag{47}$$

$$\implies \mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_{z'} + \epsilon \mathbf{u}^*, \forall \mathbf{y}_{z'} \in \mathbf{R}_t(z') \text{ and } \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) . \tag{48}$$

(48) shows that if  $z' \in \mathcal{P}_{\text{pess},t}$ ,  $x$  should be discarded. If  $z' \notin \mathcal{P}_{\text{pess},t}$ , a similar argument can be made until the condition to be inside  $\mathcal{P}_{\text{pess},t}$  holds since  $\mathcal{A}_t$  is a finite set and  $\mathcal{P}_{\text{pess},t}$  is not empty. Hence, the lemma is proved.  $\square$

The following lemma establishes an upper bound on the sum of maximum diagonal distances of the confidence regions over rounds, linking it to the maximum information gain. The proof follows by relating the maximum diagonal distance of confidence regions to the sum of variances and leveraging known bounds on information gain. Finally, by applying the stopping condition from Lemma 8, we obtain a bound on the number of rounds required for termination, completing the proof of Theorem 1.

**Lemma 9** *Let  $t_s$  represent the round in which the algorithm terminates. We have*

$$\sum_{t=1}^{t_s} \bar{\omega}_t \leq \sqrt{t_s (8\beta_{t_s} \sigma^2 \eta M \gamma_{t_s})} ,$$

where  $\eta = \frac{\sigma^{-2}}{\ln(1+\sigma^{-2})}$  and  $\gamma_{t_s}$  is the maximum information gain in  $t_s$  evaluations.

**Proof** Since the diagonal distance of the hyperrectangle  $\mathbf{Q}_t(x)$  is the largest distance between any two points in the hyperrectangle, we have

$$\sum_{t=1}^{t_s} \bar{\omega}_t^2 = \sum_{t=1}^{t_s} \max_{\mathbf{y}, \mathbf{y}' \in \mathbf{R}_t(x_t)} \|\mathbf{y} - \mathbf{y}'\|_2^2 \tag{49}$$

$$\begin{aligned} &\leq \sum_{t=1}^{t_s} \max_{\mathbf{y}, \mathbf{y}' \in \mathbf{R}_{t-1}(x_t)} \|\mathbf{y} - \mathbf{y}'\|_2^2 \\ &\leq \sum_{t=1}^{t_s} \max_{\mathbf{y}, \mathbf{y}' \in \mathbf{Q}_{t-1}(x_t)} \|\mathbf{y} - \mathbf{y}'\|_2^2 \end{aligned} \tag{50}$$

$$\begin{aligned} &= \sum_{t=1}^{t_s} \sum_{j=1}^M \left( 2\beta_{t-1}^{1/2} \sigma_{t-1}^j(x_t) \right)^2 \\ &\leq 4\beta_{t_s} \sum_{t=1}^{t_s} \sum_{j=1}^M (\sigma_{t-1}^j(x_t))^2 \end{aligned} \tag{51}$$

$$\begin{aligned}
 &= 4\beta_{t_s}\sigma^2 \sum_{t=1}^{t_s} \sum_{j=1}^M \sigma^{-2}(\sigma_{t-1}^j(x_t))^2 \\
 &\leq 4\beta_{t_s}\sigma^2\eta \left( \sum_{t=1}^{t_s} \sum_{j=1}^M \ln(1 + \sigma^{-2}(\sigma_{t-1}^j(x_t))^2) \right)
 \end{aligned} \tag{52}$$

$$\leq 8\beta_{t_s}\sigma^2\eta MI(\mathbf{y}_{[t_s]}; f_{[t_s]}) \tag{53}$$

$$\leq 8\beta_{t_s}\sigma^2\eta M\gamma_{t_s}, \tag{54}$$

where  $\eta := \sigma^{-2} / \ln(1 + \sigma^{-2})$  and  $\sigma$  is the noise standard deviation; (49) is due to the definition of  $\bar{\omega}_t$ ; (50) follows from (18); (51) holds since  $\beta_t$  is nondecreasing in  $t$ ; (52) follows from the fact that  $s \leq \eta \ln(1 + s)$  for all  $0 \leq s \leq \sigma^{-2}$  and that we have

$$\begin{aligned}
 &\sigma^{-2} (\sigma_{t-1}^j(x_t))^2 \\
 &= \sigma^{-2} \left( k^{jj}(x_t, x_t) - \left( \mathbf{k}_{[t-1]}(x_t) (\mathbf{K}_{[t-1]} + \mathbf{\Sigma}_{[t-1]})^{-1} \mathbf{k}_{[t-1]}(x_t)^\top \right)^{jj} \right) \\
 &\leq \sigma^{-2} k^{jj}(x_t, x_t) \\
 &\leq \sigma^{-2},
 \end{aligned} \tag{55}$$

where (55) follows from the fact that  $k^{jj}(x, x') \leq 1$  for all  $x, x' \in \mathcal{X}$ , and (53) follows from Nika, Bozgan, Elahi, Ararat, and Tekin (2021, Proposition 1). Finally, by Cauchy-Schwarz inequality, we have

$$\sum_{t=1}^{t_s} \bar{\omega}_t \leq \sqrt{t_s \sum_{t=1}^{t_s} \bar{\omega}_t^2} \leq \sqrt{t_s (8\beta_{t_s}\sigma^2\eta M\gamma_{t_s})}.$$

□

By definition,  $\bar{\omega}_t = \omega_t(x_t) \leq \omega_{t-1}(x_t) \leq \max_{x \in \mathcal{W}_{t-1}} \omega_{t-1}(x) = \bar{\omega}_{t-1}$ , where the first inequality is due to  $\mathbf{R}_t(x) \subseteq \mathbf{R}_{t-1}(x)$  for  $x \in \mathcal{X}$  and the last inequality is due to the selection rule of VOGP. Hence, we conclude that  $\bar{\omega}_t \leq \bar{\omega}_{t-1}$ . By above and by Lemma 9, we have

$$\bar{\omega}_{t_s} \leq \frac{\sum_{t=1}^{t_s} \bar{\omega}_t}{t_s} \leq \sqrt{\frac{8\beta_{t_s}\sigma^2\eta M\gamma_{t_s}}{t_s}}.$$

By using the stopping condition of Lemma 8, we get

$$T := \min \left\{ t \in \mathbb{N} : \sqrt{\frac{8\beta_t\sigma^2\eta M\gamma_t}{t}} < \frac{\epsilon}{d_C(1)} \right\}.$$

Hence, the proof of Theorem 1 is complete.

## 5.2 Proof of Theorem 2

The following auxiliary lemma will be helpful.

**Lemma 10** (Broxson (2006), Theorem 15) *Let  $\mathbf{A} \in \mathbb{R}^{n \times n}$ ,  $\mathbf{B} \in \mathbb{R}^{m \times m}$  be two real square matrices, where  $m, n \in \mathbb{N}$ . If  $\lambda$  is an eigenvalue of  $\mathbf{A}$  with corresponding eigenvector  $\mathbf{v} \in \mathbb{R}^n$  and  $\mu$  is an eigenvalue of  $\mathbf{B}$  with corresponding eigenvector  $\mathbf{u} \in \mathbb{R}^m$ , then  $\lambda\mu$  is an eigenvalue of  $\mathbf{A} \otimes \mathbf{B} \in \mathbb{R}^{mn \times mn}$ , the Kronecker product of  $\mathbf{A}$  and  $\mathbf{B}$ , with corresponding eigenvector  $\mathbf{v} \otimes \mathbf{u} \in \mathbb{R}^{mn}$ . Moreover, the set of eigenvalues of  $\mathbf{A} \otimes \mathbf{B}$  is  $\{\lambda_i \mu_j : i \in [n], j \in [m]\}$ , where  $\lambda_1, \dots, \lambda_n$  are the eigenvalues of  $\mathbf{A}$  and  $\mu_1, \dots, \mu_m$  are the eigenvalues of  $\mathbf{B}$  (including algebraic multiplicities). In particular, the set of eigenvalues of  $\mathbf{A} \otimes \mathbf{B}$  is the same as the set of eigenvalues of  $\mathbf{B} \otimes \mathbf{A}$ .*

The following lemma associates the information gain of an  $M$ -output GP with the maximum information gain of single output GPs.

**Lemma 11** *Let  $\mathbf{f}$  be a realization from an  $M$ -output GP with a separable covariance function of the form  $(x, x') \mapsto \mathbf{k}(x, x') = [\tilde{k}(x, x')k^*(p, q)]_{p, q \in [M]}$ , where  $\tilde{k} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  is an RBF or Matérn kernel for the design space and  $k^* : [M] \times [M] \rightarrow \mathbb{R}$  is a kernel for the objective space. For each  $p \in [M]$ , let  $\Psi_p$  be the maximum information gain for a single output GP whose kernel is  $(x, x') \mapsto \tilde{k}(x, x')k^*(p, p)$ . Then, for every  $t \in \mathbb{N}$ , we have*

$$I(\mathbf{y}_{[t]}; \mathbf{f}_{[t]}) \leq M \max_{p \in [M]} \Psi_p.$$

**Proof** Our proof is similar to the proof of C.-Y. Li, Rakitsch, and Zimmer (2022, Theorem 2). The difference comes from the structures of the covariance matrices, where the order of Kronecker products are swapped in our case. In other words, Kronecker product of input and output kernels to form the covariance matrix has swapped order.

Recall that  $\mathbf{K}_t = (\mathbf{k}(x_i, x_j))_{i, j \in [t]}$ . Hence, we have

$$\begin{aligned} I(\mathbf{y}_{[t]}; \mathbf{f}_{[t]}) &= H(\mathbf{y}_{[t]}) - H(\mathbf{y}_{[t]} | \mathbf{f}_{[t]}) \\ &= \frac{1}{2} \ln |2\pi e (\mathbf{K}_t + \sigma^2 \mathbf{I}_{Mt})| - \frac{1}{2} \ln |2\pi e \sigma^2 \mathbf{I}_{Mt}| \\ &= \frac{1}{2} \ln \left( \frac{|2\pi e (\mathbf{K}_t + \sigma^2 \mathbf{I}_{Mt})|}{|2\pi e \sigma^2 \mathbf{I}_{Mt}|} \right) \\ &= \frac{1}{2} \ln \left| (\mathbf{K}_t + \sigma^2 \mathbf{I}_{Mt}) \cdot (\sigma^2 \mathbf{I}_{Mt})^{-1} \right| \\ &= \frac{1}{2} \ln \left| \mathbf{I}_{Mt} + \frac{1}{\sigma^2} \mathbf{K}_t \right|. \end{aligned}$$

By the separable form of  $\mathbf{k}$ , we have

$$\mathbf{K}_t = [\tilde{k}(x_i, x_j)]_{i,j \in [t]} \otimes [k^*(p, q)]_{p,q \in [M]}.$$

Hence, by Lemma 10 and using the identity  $|\mathbf{I} + \mathbf{A} \otimes \mathbf{B}| = |\mathbf{I} + \mathbf{B} \otimes \mathbf{A}|$ , we get

$$I(\mathbf{y}_{[t]}; \mathbf{f}_{[t]}) = \frac{1}{2} \ln |\mathbf{I}_{Mt} + \frac{1}{\sigma^2} [k^*(p, q)]_{p,q \in [M]} \otimes [\tilde{k}(x_i, x_j)]_{i,j \in [t]}|.$$

Notice that

$$\begin{aligned} & \mathbf{I}_{Mt} + \frac{1}{\sigma^2} [k^*(q, p)]_{p,q \in [M]} \otimes [\tilde{k}(x_i, x_j)]_{i,j \in [t]} \\ &= \begin{bmatrix} \mathbf{I}_t + k^*(1, 1)[\tilde{k}(x_i, x_j)]_{i,j \in [t]} \sigma^{-2}, & \dots, & k^*(1, M)[\tilde{k}(x_i, x_j)]_{i,j \in [t]} \sigma^{-2} \\ \vdots & & \vdots \\ k^*(M, 1)[\tilde{k}(x_i, x_j)]_{i,j \in [t]} \sigma^{-2}, & \dots, & \mathbf{I}_t + k^*(M, M)[\tilde{k}(x_i, x_j)]_{i,j \in [t]} \sigma^{-2} \end{bmatrix}. \end{aligned}$$

Since the matrix itself and all of its diagonal blocks are positive definite symmetric real matrices, we can apply Fischer’s inequality and obtain

$$I(\mathbf{y}_{[t]}; \mathbf{f}_{[t]}) \leq \frac{1}{2} \sum_{p=1}^M \ln |\mathbf{I}_t + k^*(p, p) \sigma^{-2} [\tilde{k}(x_i, x_j)]_{i,j \in [t]}|.$$

This is actually the sum of mutual informations of single output GPs  $f_i \sim \mathcal{GP}(\mathbf{0}, k^*(l, l)[\tilde{k}(x_i, x_j)]_{i,j \in [t]})$ . Notice that a positive constant multiple of an RBF (resp. Matérn) kernel is still an RBF (resp. Matérn) kernel. Since the mutual information is bounded by maximum information gain, we obtain  $I(\mathbf{y}_{[t]}; \mathbf{f}_{[t]}) \leq M \max_{p \in [M]} \Psi_p$ , which completes the proof.

Now, we can continue with the proof of Theorem 2. By Lemma 8, we have to show that taking

$$t = \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta \cdot \ln^{D+2} \left( \sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \cdot \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta \right)$$

satisfies  $\frac{8\beta_t \sigma^2 \eta M \gamma_t}{t} < \frac{\epsilon^2}{d_C^2(1)}$  where  $\phi$  is the multiplicative constant that comes from  $\mathcal{O}(\cdot)$  notation of the bound on single output GP’s information gain. We will first upper-bound the LHS of this inequality and show that its smaller than or equal to the RHS of the inequality. We are using an RBF kernel  $\tilde{k}$  for the design space. Then, by Lemma 11 and the bounds on maximum information gain established in Vakili et al. (2021), we have

$$I(\mathbf{y}_{[t]}; \mathbf{f}_{[t]}) \leq M \cdot \mathcal{O}(\ln^{D+1}(t)) \implies \gamma_t \leq M \cdot \mathcal{O}(\ln^{D+1}(t)).$$

Notice that in Theorem 1, as  $\epsilon$  goes to 0,  $T$  goes to infinity. Therefore, we can use the bounds on maximum information gain established in Srinivas et al. (2012). We have  $\beta_t = \ln(\frac{M\pi^2|\mathcal{X}|t^2}{3\delta})$ .

$$\begin{aligned}
 \frac{8\beta_t \sigma^2 \eta M \gamma_t}{t} &= \frac{8 \ln\left(\frac{M\pi^2 |\mathcal{X}| t^2}{3\delta}\right) \sigma^2 \eta M \gamma_t}{t} \\
 &\leq \frac{8 \ln\left(\frac{M\pi^2 |\mathcal{X}| t^2}{3\delta}\right) \sigma^2 \eta M^2 \phi \ln^{D+1}(t)}{t} \\
 &= \frac{16 \ln\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \cdot t\pi\right) \sigma^2 \eta M^2 \phi \ln^{D+1}(t)}{t} \\
 &< \frac{16 \ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \cdot t\pi\right) \sigma^2 \eta M^2 \phi}{t}
 \end{aligned}$$

where  $\phi$  is the multiplicative constant that comes from the  $\mathcal{O}(\cdot)$  notation. The last inequality follows from  $M \geq 1, \delta \in (0, 1), |\mathcal{X}| \geq 1$  and  $\pi > \sqrt{3}$ . Now, we can plug in

$$t = \frac{32d_C^2}{\epsilon^2} M^2 \sigma^2 \eta \phi \cdot \ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \cdot \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta\right)$$

to the last expression:

$$\begin{aligned}
 &\frac{16 \ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta \cdot \ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \cdot \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta\right)\right) \phi M^2 \sigma^2 \eta}{\frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta \cdot \ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta\right)} \\
 &= \frac{\ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta \cdot \ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \cdot \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta\right)\right)}{2 \ln^{D+2}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta\right)} \cdot \frac{\epsilon^2}{d_C^2}
 \end{aligned}$$

To show that the last expression is smaller than or equal to  $\frac{\epsilon^2}{d_C^2}$ , we need to show that

$$\ln^{D+2}\left(B \ln^{D+2}(B)\right) \leq 2 \ln^{D+2}(B) = \ln^{D+2}\left(B \left(2^{\frac{1}{D+2}}\right)\right)$$

where  $B = \sqrt{\frac{M|\mathcal{X}|}{3\delta}} \pi \frac{32d_C^2}{\epsilon^2} \phi M^2 \sigma^2 \eta$  for convenience. This holds because  $\ln^{D+2} B \leq B \left(2^{\frac{1}{D+2}} - 1\right)$  for sufficiently large  $B$ . Hence, the proof is complete.  $\square$

### 6 Numerical Results

In this section, we analyze where VOGP stands when compared to the state-of-the-art.

### 6.1 Datasets, Cones and Performance Benchmarks

**BC** (Branin-Currin,  $D = 2, M = 2, |\mathcal{X}| = 500$ ) Simultaneous optimization of two widely used single-output test benchmarks.

**LAC** (Lactose,  $D = 2, M = 2, |\mathcal{X}| = 250$ ) A chemical reaction for the isomerization of lactulose from lactose (Hashemi & Ashtiani, 2010; Müller et al., 2022).

**VS** (Vehicle Safety,  $D = 5, M = 3, |\mathcal{X}| = 500$ ) Optimization of vehicle structures concerning safety, i.e., crashworthiness (Liao et al., 2008).

**SnAr** ( $D = 4, M = 2, |\mathcal{X}| = 2000$ ) Optimization of the nucleophilic aromatic substitution reaction (SnAr) involving 2,4-difluoronitrobenzene and pyrrolidine (Hone et al., 2017).

**BCC** (Branin-Currin Continuous,  $D = 2, M = 2$ ) We use the same test benchmark functions as in BC dataset with continuous domain over  $[0, 1]^2$ .

**ZDT3** ( $D = 2, M = 2$ ) A commonly used test benchmark function with continuous domain over  $[0, 1]^2$  (Zitzler et al., 2000).

In our experiments, we use three cones which we call *acute*, *right*, and *obtuse*. For  $M = 2$ , they correspond to  $C_{60^\circ}, C_{90^\circ}$ , and  $C_{120^\circ}$  in Definition 5 below, respectively. Figure 5 illustrates these ordering cones. However, for VC dataset, which is a tri-objective problem ( $M = 3$ ), we define acute and obtuse cones with  $W$  matrices given as

$$\text{Acute cone: } \mathbf{W} = \sqrt{21} \begin{bmatrix} 1 & -2 & 4 \\ 4 & 1 & -2 \\ -2 & 4 & 1 \end{bmatrix}, \quad \text{Obtuse cone: } \mathbf{W} = \sqrt{3.72} \begin{bmatrix} 1 & 0.4 & 1.6 \\ 1.6 & 1 & 0.4 \\ 0.4 & 1.6 & 1 \end{bmatrix}$$


alongside the positive orthant cone (right cone) with identity matrix.

**Definition 5** Given  $\theta \in (0^\circ, 180^\circ)$ , the cone  $C_{\theta^\circ} \subseteq \mathbb{R}^2$  is defined as the convex cone whose two boundary rays make  $\frac{\pm\theta}{2}$  degrees with the identity line (see Figure 5).

**Remark 9** Cone  $C_{90^\circ}$  in Definition 5 corresponds to multi-objective optimization with  $M = 2$ .

Since our task is simply the classification of the Pareto designs, in the spirit of the widely used F1 score, we use the following definition that integrates our success condition (Definition 1) into a performance measure.

Fig. 5 Acute, right and obtuse ordering cones for  $M = 2$



**Definition 6** (Karagözlü et al., 2024) Given a finite set of designs  $\mathcal{X}$ , set of Pareto designs  $P^*$  and a set of predicted Pareto designs  $P$  w.r.t. the ordering cone  $C$ , and an  $\epsilon$ , we define; a positive class as  $\Pi_\epsilon = \{x \in \mathcal{X} : \Delta_x^* \leq \epsilon\}$ , where  $\Delta_x^*$  is the suboptimality gap, and an  $\epsilon$  lenient F1 classification score ( $\epsilon$ -F1) as

$$\epsilon\text{-F1} = \frac{2|\Pi_\epsilon \cap P|}{2|\Pi_\epsilon \cap P| + |\Pi_{\epsilon \setminus P}| + |P \setminus \Pi_\epsilon|},$$

where  $\Pi_{\epsilon \setminus P}$  is the set of Pareto optimal arms that do not satisfy Definition 1 part (i).

In multi-objective optimization, hypervolume indicator, denoted  $\text{HV}(\hat{\mathcal{P}})$ , is an ubiquitous metric and it corresponds to the volume dominated by a set representing the found Pareto front, i.e.,  $\hat{\mathcal{P}}$ , and a reference point. This metric is especially useful when working with continuous domains. Further, HV discrepancy, defined by

$$d_{\text{HV}} = |\text{HV}(\mathcal{P}^*) - \text{HV}(\hat{\mathcal{P}})|,$$


is a well-known measure of how well the set  $\hat{\mathcal{P}}$  represents a known, true Pareto front, i.e.,  $\mathcal{P}^*$ . Below, we extend the definition of HV to vector optimization. We also establish a theoretical equivalence with the original definition in the case where  $M = N$ , using the Jacobian of cone matrix  $\mathbf{W}$ .

**Definition 7** Given a reference point  $\mathbf{r} \in \mathbb{R}^M$  and  $N \times M$  matrix  $\mathbf{W}$  of ordering cone  $C$ , hypervolume indicator ( $\text{HV}_C$ ) of a finite approximate Pareto set  $\hat{\mathcal{P}}$  is the  $N$ -dimensional Lebesgue measure  $\lambda_N$  of the space dominated by  $\mathcal{P}$  and bounded from below by  $\mathbf{r}$ :  $\text{HV}_C(\hat{\mathcal{P}}, \mathbf{r}) = \lambda_N \left( \bigcup_{i=1}^{|\hat{\mathcal{P}}|} [\mathbf{W}\mathbf{r}, \mathbf{W}\mathbf{y}_i] \right)$ , where  $[\mathbf{W}\mathbf{r}, \mathbf{W}\mathbf{y}_i]$  denotes the hyperrectangle bounded by vertices  $\mathbf{W}\mathbf{r}$  and  $\mathbf{W}\mathbf{y}_i$ .

**Remark 10** The cone hypervolume  $\text{HV}_C$  defined in Definition 7 provides a general definition of hypervolume for any dimensions  $N, M$ . However, when  $N = M$ , by Theorem 2.3.15 of Heil (2019), the  $\text{HV}_C$  simplifies to  $|\det \mathbf{W}| \lambda_M \left( \bigcup_{i=1}^{|\hat{\mathcal{P}}|} (\mathbf{r} + C) \cap (\mathbf{y}_i - C) \right)$  where  $\mathbf{r} \in \mathbb{R}^M$  is the reference point. This represents the hypervolume that is cone dominated by the Pareto front. A two dimensional illustration of this transformation and the corresponding hypervolume scaling is provided in Figure 6.

## 6.2 Experimental Setup and Competing Algorithms

We apply min-max normalization to inputs. We also standardize the outputs to work on a unified scale of  $[0, 1]$ , i.e., for  $\epsilon$  and  $\sigma$ . We fix  $\epsilon = 0.1$ ,  $\delta = 0.05$ ,  $\sigma = 0.1$ . The GP model is initialized using a single randomly selected design point to start the optimization. We use RBF kernel featuring automatic relevance determination for all experiments as it is fairly standard, and in line with our GP assumption in Section 3, we assume known hyperparameters and choose them using maximum likelihood estimation prior to optimization. Similar to the other works with the same confidence hyperrectangle definition (Zuluaga et al., 2016),



**Fig. 6** Illustration of Definition 7 in a simple two dimensional case. The left plot shows the dominated region  $V$  under the cone  $C$ , and the right plot shows its image after the linear transformation induced by  $W$ , where the region becomes axis-aligned and its volume scales by  $|\det W|$

**Table 2**  $\epsilon$ -F1 score comparison of VOGP with NE and PaVeBa on several datasets. S.C. refers to sample complexities of the algorithms

| Dataset | Alg.   | Acute Cone        |                               | Right Cone       |                               | Obtuse Cone      |                               |
|---------|--------|-------------------|-------------------------------|------------------|-------------------------------|------------------|-------------------------------|
|         |        | S.C.              | $\epsilon$ -F1 ( $\uparrow$ ) | S.C.             | $\epsilon$ -F1 ( $\uparrow$ ) | S.C.             | $\epsilon$ -F1 ( $\uparrow$ ) |
| BC      | NE     | 500               | $0.89 \pm .07$                | 500              | $0.92 \pm .07$                | 500              | $0.98 \pm .06$                |
|         | PaVeBa | $517.7 \pm 5.3$   | $0.95 \pm .03$                | $504.5 \pm 3.0$  | $0.96 \pm .04$                | $500.9 \pm 1.1$  | $0.93 \pm .09$                |
|         | VOGP   | $93.5 \pm 32.6$   | $0.93 \pm .04$                | $28.2 \pm 3.8$   | $0.96 \pm .08$                | $18.3 \pm 2.6$   | $0.99 \pm .03$                |
| LAC     | NE     | 250               | $0.95 \pm .02$                | 250              | $0.93 \pm .03$                | 250              | $0.93 \pm .05$                |
|         | PaVeBa | $290.0 \pm 4.2$   | $1.00 \pm .00$                | $266.4 \pm 2.1$  | $0.99 \pm .01$                | $264.1 \pm 3.7$  | $0.98 \pm .02$                |
|         | VOGP   | $69.9 \pm 13.7$   | $1.00 \pm .00$                | $27.4 \pm 3.8$   | $0.99 \pm .01$                | $37.9 \pm 6.5$   | $0.99 \pm .01$                |
| VS      | NE     | 500               | $0.93 \pm .03$                | 500              | $0.95 \pm .04$                | 500              | $0.90 \pm .10$                |
|         | PaVeBa | $515.2 \pm 3.1$   | $0.90 \pm .03$                | $503.3 \pm 1.7$  | $0.94 \pm .03$                | $502.4 \pm 1.8$  | $0.90 \pm .05$                |
|         | VOGP   | $406.2 \pm 158.9$ | $0.93 \pm .02$                | $34.8 \pm 10.2$  | $0.77 \pm .06$                | $23.6 \pm 3.9$   | $0.87 \pm .11$                |
| SnAr    | NE     | 2000              | $0.89 \pm .08$                | 2000             | $0.89 \pm .12$                | 2000             | $0.86 \pm .30$                |
|         | PaVeBa | $2016.9 \pm 5.7$  | $0.91 \pm .03$                | $2005.2 \pm 2.6$ | $0.87 \pm .10$                | $2001.8 \pm 1.7$ | $0.92 \pm .11$                |
|         | VOGP   | $102.5 \pm 15.2$  | $0.97 \pm .02$                | $41.4 \pm 3.8$   | $0.87 \pm .12$                | $36.4 \pm 3.7$   | $1.00 \pm .00$                |

we scale down  $\beta_t$  by an empirical factor, in this case  $2^5 = 32$ . We report our results as average of 10 runs.

We use a series of algorithms to compare against. Previous work introduced Naïve Elimination (NE) (Ararat, 2023) algorithm that is designed for vector optimization using ordering cones. Even though it is an  $(\epsilon, \delta)$ -PAC algorithm, since its theoretical sampling budget is impractically high, we treat NE as a fixed-budget algorithm. It uses a per-design sampling budget  $L$ , and we set  $L = \lceil T/|\mathcal{X}| \rceil$  where  $T$  is the sampling complexity of VOGP for that specific experiment. Another work on vector optimization that does  $(\epsilon, \delta)$ -PAC Pareto set identification is PaVeBa (Karagözlü et al., 2024). PaVeBa, samples all designs that are not discarded at all rounds. We scale down their confidence radius  $r_t$  by 32 in this method to limit high sample complexities. As NE and PaVeBa also assume finite design set, we start by comparing VOGP with these two and report the results in Table 2. Also, we provide the convergence plots for the results in Table 2 in Appendix C. To introduce further baselines, we refer to multi-objective optimization literature, which is a special case of vector optimization where the ordering cone is the positive orthant cone. We devise fixed-budget variants for widely used information theoretic acquisition functions MESMO (Belakaria et al., 2019) and JES (Tu et al., 2022). As these algorithms work in continuous domains, we propose an adaptive discretization (Bubeck et al., 2011) based VOGP as a variant to compare. We run


these acquisition functions over the design set  $\mathcal{X}$  with GPs using the same kernel as VOGP and with the same sample complexity as VOGP. We compute the predicted Pareto set w.r.t. the ordering cone on the posterior mean for MESMO, JES and VOGP. A naive approach of adapting these acquisition functions would be transforming the problem by applying the linear mapping induced by matrix  $\mathbf{W}$  and then using the methods as if doing a multi-objective optimization. However, this approach proves ineffective and a more in depth discussion on this can be found in Section 6.4. We do not compare VOGP with  $\epsilon$ -PAL since in the multi-objective setting, VOGP can be considered as an extension of it and adapting  $\epsilon$ -PAL to the ordering cones is non-trivial.

### 6.3 Ablation Study for Continuous Domain Adaptation

Although theoretical bounds for VOGP are obtained for a finite set of designs  $|\mathcal{X}|$ , as an ablation, we implement a heuristic on top of VOGP to make it suitable for continuous problems. We adaptively discretize (Bubeck et al., 2011; Nika et al., 2021) the continuous design space into cells, forming a tree structure with a hyperparameter that decides the maximum depth the tree can reach. We fix the maximum depth as 5 in this case. We treat leaf nodes as designs while running VOGP. We split a leaf node into its  $2^M$  child branches if it has high enough confidence and is not discarded (pruned) yet. We delay epsilon covering until the tree is maximally expanded and we run the remaining phases of VOGP as usual. Since the returned Pareto set after this process is not highly informative of what the Pareto front looks like in the continuous sense (as it is sparse), we treat VOGP like an exploration algorithm with adaptive maximum variance reduction as its acquisition function. Concretely, after the termination of VOGP we calculate the cone-dependent Pareto set using a uniform discretization over the posterior distribution of the GP. We do the same for MESMO and JES as well. We, then, calculate logarithm HV discrepancy for the found Pareto set for all algorithms. Since our theoretical confidence intervals assume finite domain, we consult to the literature for a confidence interval that is tight and has a theoretical underpinnings for continuous domains. We use the bound in Fiedler, Scherer, and Trimpe (2021, Theorem 1). Even though their reproducing kernel Hilbert space (RKHS) assumption does not coincide with our GP sample assumption, we found that this bound is empirically very sound. We run VOGP by taking the RKHS bound as  $B = 0.1$  and  $\beta_t$  scaled by a factor 32 for BCC and 48 for ZDT3 problem. We present the results of this experiment in Table 3 and Figure 7.

**Table 3** Logarithm hypervolume discrepancy comparison of VOGP with MESMO and JES on BCC and ZDT3 problems in continuous domain. S.C. refers to sample complexities of the algorithms. Lower is better for both metrics

| Dataset | Alg.  | Acute Cone        |                  | Right Cone       |                  | Obtuse Cone     |                  |
|---------|-------|-------------------|------------------|------------------|------------------|-----------------|------------------|
|         |       | S.C.              | $\log(d_{HV})$   | S.C.             | $\log(d_{HV})$   | S.C.            | $\log(d_{HV})$   |
| BCC     | MESMO |                   | $-0.96 \pm 0.50$ |                  | $-1.52 \pm 0.89$ |                 | $-1.37 \pm 2.22$ |
|         | JES   |                   | $-0.75 \pm 0.73$ |                  | $-2.22 \pm 0.61$ |                 | $-4.28 \pm 1.35$ |
|         | VOGP  | $168.6 \pm 75.02$ | $0.05 \pm 0.72$  | $27.5 \pm 8.3$   | $-2.31 \pm 0.72$ | $23.3 \pm 7.29$ | $-4.30 \pm 1.30$ |
| ZDT3    | MESMO |                   | $-2.35 \pm 0.65$ |                  | $-1.59 \pm 0.87$ |                 | $-1.12 \pm 0.52$ |
|         | JES   |                   | $-2.62 \pm 0.71$ |                  | $-1.28 \pm 0.89$ |                 | $-1.12 \pm 0.97$ |
|         | VOGP  | $152.0 \pm 69.0$  | $-2.65 \pm 0.67$ | $64.8 \pm 101.2$ | $-1.75 \pm 0.67$ | $38.9 \pm 15.8$ | $-1.68 \pm 0.96$ |



**Fig. 7** Mean logarithm hypervolume discrepancy vs. rounds on BCC (Left) and ZDT3 (Right) datasets, using right cone. Number of rounds is selected as the median S.C. Shaded regions correspond to half standard deviation

### 6.4 Adapting MO Methods to Vector Optimization via Linear Transformations

The cone ordering can be characterized by the usual componentwise ordering via the linear mapping induced by matrix  $W$ . This can be confirmed by observing that

$$x \preceq_C y \iff W(y - x) \geq 0 \iff Wy \geq Wx \iff Wx \leq Wy, \tag{56}$$

where  $\leq, \geq$  are componentwise orders (i.e. positive orthant order). A naive approach of adapting multi-objective optimization (MO) methods to vector optimization is using them per usual with the transformed objective function  $W \circ f$ . Please note that this linear transformation is  $\mathbb{R}^M \rightarrow \mathbb{R}^N$ , where  $M$  is the number of objectives and  $N$  is the number of halfspaces that define the polyhedral cone. In the paper, the concept of vector optimization is demonstrated mostly for the case  $M = 2$  since it is possible to visualize this case easily. Note that when  $M = 2$ , we necessarily have  $N = 2$  for a solid pointed ordering cone  $C$ . However, when  $M \geq 3$ , in general, the number of halfspaces  $N$  can be (much) larger than the number of objectives  $M$ . Hence, applying componentwise optimization methods (e.g. JES) to the transformed predicted objective values results in running these methods in the  $N$ -dimensional objective space (hence, an  $N$ -objective GP), which increases the computational cost of these methods dramatically, rendering them ineffective.

In this section we illustrate these challenges, particularly when the objective space dimension is modest (for instance,  $M = 3$ ) yet the complexity introduced by a large number of halfspaces defining the polyhedral cone ( $N > M$ ) is significant—a scenario common in practice (e.g. P1 of Hunt et al. (2010)). This is a crucial distinction because it underscores a unique advantage of VOGP: its ability to efficiently navigate the solution space in vector optimization problems, even when the the number of halfspaces defining the polyhedral cone (indicated by  $N$ ) is high.

Consider the ice cream cone (expressed by the inequality  $x_3 \geq \sqrt{x_1^2 + x_2^2}$ ) for  $M = 3$ , and take its rotated version whose symmetry axis lies in the ray in  $\mathbb{R}_+^3$  given by the equation  $x_1 = x_2 = x_3$ . Notice that the rotated ice cream cone is not polyhedral but one can construct a polyhedral approximation  $C$  that is defined as the intersection of  $N \geq 3$  halfspaces. The cone in Figure 2 is a polyhedral approximation of the ice cream cone with  $N = 6$ . We conduct an experiment on VS dataset using different  $W$  matrices that yield polyhedral approximations of the rotated ice cream cone with an increasing number ( $N$ ) of halfspaces. We conduct the experiments on a discrete domain to limit acquisition evaluation of MO

methods and since it conveys the message well enough. We use the original VOGP with  $\beta_t$  set as the same as the continuous domain ablation. From Table 4, it can be seen that the computation times of MO methods (MESMO and JES) increase drastically with the number of halfspaces ( $N$ ), highlighting that the formulation suggested is not really practical. On the contrary, we are able to obtain the results for  $N = 81$  with VOGP in reasonable computation time.

For  $N = 9$  (and also  $N = 27$  and  $N = 81$ ), while running JES algorithm, we observed hundreds of gigabytes of RAM usage and the algorithm raised out-of-memory error in our system with 128GB of RAM and 256GB of swap area.

## 6.5 Ablation Study for Unknown GP Hyperparameters

In some situations, neither the correct hyperparameters are inaccessible to us, nor do we possess reliable initial estimates for them. In such situations, even though the initial hyperparameters are not quite correct, a common approach is to gradually learn and adjust the hyperparameters as more queries get included in observations. So, in this ablation study we start with a random set of GP hyperparameters and after each evaluation phase, we update them by a maximum likelihood estimation (MLE) using the observations made so far.

Normally, after VOGP performs discarding and Pareto identification phases, it does not check the validity of decisions made in these phases in the subsequent rounds. This is because given that the GP model is accurate, these decisions are guaranteed to be good enough with high probability. However, in the unknown parameters setting, the parameters are yet to be meaningful in the earlier rounds. Hence, the decisions made in the earlier rounds can be erroneous. To adapt VOGP to this setting, at the end of each round, we reset the sets  $\mathcal{S}_t$  and  $\mathcal{P}_t$ . With this adaptation, VOGP only terminates when all the designs are either discarded or moved to  $\mathcal{P}_t$  in a single round. This ultimately allows VOGP to defer its decisions to when the kernel hyperparameters had been learned with the cost of doing repetitive work. We do not alter the rest of the algorithm.

The results, as given in Table 5, seem to improve over the known hyperparameters setting. We speculate this is because two things: (i) sample complexity of VOGP increases and (ii) these problems are not actually GP samples. First, sample complexity of VOGP is likely to be increased because it needs a high level of confidence of all designs to terminate in the span of a single round. Second, under the known hyperparameters setting, when MLE is done before the optimization process to find the "correct" hyperparameters, overall smoothness of the whole problem space is learned. Instead, if we could focus on and learn the local smoothness around the Pareto set, it would yield better results. Indeed, we think that this

**Table 4** Comparison using different number of half spaces for ice cream cone approximation experiment. "DNF" means "Did not finish"

| N  | Alg.  | Time (sec.) | S.C.              | $\epsilon$ -F1  |
|----|-------|-------------|-------------------|-----------------|
| 9  | VOGP  | 22.38       | $28.50 \pm 3.8$   | $0.88 \pm 0.09$ |
| 9  | MESMO | 104.26      |                   | $0.83 \pm 0.08$ |
| 9  | JES   | DNF         |                   | –               |
| 27 | VOGP  | 59.68       | $28.30 \pm 3.200$ | $0.86 \pm 0.09$ |
| 27 | MESMO | 299.50      |                   | $0.82 \pm 0.06$ |
| 27 | JES   | DNF         |                   | –               |
| 81 | VOGP  | 173.56      | $28.30 \pm 3.200$ | $0.86 \pm 0.09$ |
| 81 | MESMO | 879.76      |                   | $0.84 \pm 0.08$ |
| 81 | JES   | DNF         |                   | –               |

**Table 5**  $\epsilon$ -F1 scores of VOGP on several datasets under unknown hyperparameters setting. We only report results for acute cone. S.C. refers to sample complexity

| Dataset | S.C.                | $\epsilon$ -F1 |
|---------|---------------------|----------------|
| BC      | 117.10 $\pm$ 20.27  | 0.99 $\pm$ .01 |
| VS      | 555.10 $\pm$ 147.16 | 1.00 $\pm$ .00 |
| SnAr    | 126.60 $\pm$ 19.77  | 0.96 $\pm$ .01 |
| LAC     | 99.70 $\pm$ 18.63   | 1.00 $\pm$ .00 |

is what is happening when we start with unknown hyperparameters setting and learn them iteratively as VOGP focuses on identifying the Pareto set.

## 6.6 Discussion of Numerical Results

Table 2 demonstrates VOGP's high sample efficiency compared to the state-of-the-art methods in vector optimization. Specifically, we observe  $\sim 18.1\times$  and  $\sim 18.3\times$  lower sample complexity compared with NE and PaVeBa, respectively. These results are calculated as the mean of ratios of our method's average samples to those of NE's and PaVeBa's for each dataset/cone configuration in Table 2. This drastic difference in sample complexity is possible since the theory of VOGP, specifically the use of surrogate GPs, allows for sequential decision making to happen after each observation, whereas other methods make decisions less frequently. Despite its lower sample complexity, VOGP still outperforms the other methods in most cases. Table 3 and Figure 7 illustrate VOGP's robust performance when extended to continuous domains using adaptive discretization. VOGP remains competitive with the state-of-the-art information theoretic acquisition methods. Additionally, Table 5 shows that VOGP can be adapted to unknown hyperparameter settings and still perform very effectively. Most notably, Table 4, which uses linear transformations, emphasizes that VOGP or vector optimization in general is non-trivial. These results showcase the effectiveness of VOGP and its adaptability to different scenarios.

## 7 Conclusion

In this work we studied the problem of black-box vector optimization with Gaussian process bandits. We proposed VOGP, a sample efficient vector optimization algorithm, and provided success guarantees and upper bounds on its sample complexity. To the best of our knowledge, this is the first work that provides theoretical guarantees for black-box vector optimization with Gaussian process bandits. Through extensive experiments and ablations on various datasets and ordering cones, we showed our algorithms' effectiveness over the existing works on vector optimization and straightforward adaptations of the MO methods. A limitation of our approach is that the theoretical guarantees are established under the assumption that the underlying surrogate model is a Gaussian process and its hyperparameters are known a priori. Extending the theoretical analysis to unknown hyperparameters or to more general model classes, such as Bayesian neural networks remains an open challenge. Future research could explore such extensions, potentially broadening the applicability of vector optimization in settings where alternative surrogate models offer advantages in scalability or expressiveness. Other interesting future research directions include providing theoretical guarantees for the problem in continuous spaces and developing counterparts of information theoretic acquisition functions for vector optimization.

## Appendix A Table of Notation

Below, we present a table of symbols that are used in the proofs.

| Symbol              | Description  |
|---------------------|--|
| $\mathcal{X}$       | The design space   |
| $M$                 | The dimension of the objective space   |
| $\mathbb{S}^{M-1}$  | The unit sphere in $\mathbb{R}^M$  |
| $C$                 | The polyhedral ordering cone   |
| $\zeta$             | The ordering complexity of the cone $C$  |
| $f$                 | The objective function   |
| $\epsilon$          | Accuracy level given as input to the algorithm   |
| $\mathbf{y}_{[t]}$  | Vector that represents the first $t$ noisy observations where $\mathbf{y}_{[0]} = \emptyset$ .   |
| $\mu_t(x)$          | The posterior mean of design $x$ at round $t$ whose $j^{\text{th}}$ component is $\mu_t^j(x)$  |
| $\sigma_t(x)$       | The posterior variance of design $x$ at round $t$ whose $j^{\text{th}}$ component is $\sigma_t^j(x)$                                     |
| $\beta_t$           | The confidence term at round $t$   |
| $\mathcal{P}_t$     | The predicted Pareto set of designs at round $t$   |
| $\mathcal{S}_t$     | The undecided sets of designs at round $t$   |
| $\hat{\mathcal{P}}$ | The estimated Pareto set of designs returned by VOGP   |
| $P^*$               | The set of true Pareto optimal designs   |
| $P_\theta^*$        | The set of true Pareto optimal designs when $M = 2$ and $C = C_\theta$   |
| $\mathcal{A}_t$     | The union of sets $\mathcal{S}_t$ and $\mathcal{P}_t$ at the beginning of round $t$  |
| $\mathcal{W}_t$     | The union of sets $\mathcal{S}_t$ and $\mathcal{P}_t$ at the end of the discarding phase of round $t$                                    |
| $Q_t(x)$            | The confidence hyperrectangle associated with design $x$ at round $t$  |
| $R_t(x)$            | The cumulative confidence hyperrectangle associated with design $x$ at round $t$   |
| $x_t$               | The design evaluated at round $t$  |
| $\omega_t(x)$       | The diameter of the cumulative confidence hyperrectangle of design $x$ at round $t$  |
| $\bar{\omega}_t$    | The maximum value of $\omega_t(x)$ over all active designs $x$ at round $t$  |
| $m(x, x')$          | $\inf \{s \geq 0 \mid \exists \mathbf{u} \in \mathbb{B}(1) \cap C : \mathbf{f}(x) + s\mathbf{u} \notin \mathbf{f}(x') - \text{int}(C)\}$ |
| $\gamma_t$          | The maximum information that can be gained about $\mathbf{f}$ in $t$ evaluations   |
| $t_s$               | The round in which VOGP terminates   |

## Appendix B Analysis for a Matérn Kernel

The algorithm takes at most

$$O\left(\frac{d_C^2(1)}{\epsilon^2} \left(\frac{2\nu}{2\nu + D} - \frac{4\nu + D}{2\nu + D}\right) M^2 \sigma^2 \eta \cdot \ln \frac{4\nu + D}{2\nu + D} \left(\left(\frac{\pi^2 M |\mathcal{X}|}{3\delta}\right)^{\left(\frac{\nu}{2\nu + D}\right)} \cdot \left(\frac{2\nu}{2\nu + D} - \frac{4\nu + D}{2\nu + D}\right) \frac{d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta\right)\right)$$

samples, when it is run using a Matérn kernel where  $\nu$  is the smoothness parameter.

**Proof** By Lemma 8, we have to show that taking

$$t = \left(\frac{32d_C^2(1)}{\epsilon^2} \left(\frac{2\nu}{2\nu + D} - \frac{4\nu + D}{2\nu + D}\right) M^2 \sigma^2 \eta \phi \cdot \ln \frac{4\nu + D}{2\nu + D} \left(\left(\frac{\pi^2 M |\mathcal{X}|}{3\delta}\right)^{\left(\frac{\nu}{2\nu + D}\right)} \cdot \left(\frac{2\nu}{2\nu + D} - \frac{4\nu + D}{2\nu + D}\right) \frac{32d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta \phi\right)\right)^{\frac{2\nu + D}{2\nu}}$$

satisfies  $\frac{8\beta_t \sigma^2 \eta M \gamma_t}{t} < \frac{\epsilon^2}{d_C^2(1)}$ .

$\phi$  is the multiplicative constant that comes from  $\mathcal{O}(\cdot)$  notation of the bound on single output GP’s information gain. We will first upper-bound the LHS of this inequality and show that its smaller than or equal to the RHS of the inequality. We are using an Matérn kernel kernel  $\tilde{k}$  for the design space. Then, by Lemma 11 and the bounds on maximum information gain established in Vakili et al. (2021), we have

$$I(\mathbf{y}_{[t]}, \mathbf{f}_{[t]}) \leq M \cdot \mathcal{O}\left(T^{\frac{D}{2\nu+D}} \ln^{\frac{2\nu}{2\nu+D}}(T)\right) \implies \gamma_t \leq M \cdot \mathcal{O}\left(T^{\frac{D}{2\nu+D}} \ln^{\frac{2\nu}{2\nu+D}}(T)\right).$$

Notice that in Theorem 1, as  $\epsilon$  goes to 0,  $T$  goes to infinity. Therefore, we can use the bounds on maximum information gain established in Srinivas et al. (2012). We have  $\beta_t = \ln\left(\frac{M\pi^2|\mathcal{X}|t^2}{3\delta}\right)$ .

$$\begin{aligned} \frac{8\beta_t \sigma^2 \eta M \gamma_t}{t} &= \frac{8 \ln\left(\frac{M\pi^2|\mathcal{X}|t^2}{3\delta}\right) \sigma^2 \eta M \gamma_t}{t} \\ &\leq \frac{8 \ln\left(\frac{M\pi^2|\mathcal{X}|t^2}{3\delta}\right) \sigma^2 \eta M^2 \phi t^{\frac{D}{2\nu+D}} \ln^{\frac{2\nu}{2\nu+D}}(t)}{t} \\ &= \frac{16 \ln\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \cdot t\pi\right) \sigma^2 \eta M^2 \phi \ln^{\frac{2\nu}{2\nu+D}}(t)}{t^{\frac{2\nu}{2\nu+D}}} \\ &< \frac{16 \ln^{\frac{4\nu+D}{2\nu+D}}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \cdot t\pi\right) \sigma^2 \eta \phi M^2}{t^{\frac{2\nu}{2\nu+D}}} \\ &= \frac{16\alpha \ln^{\frac{4\nu+D}{2\nu+D}}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \cdot t\pi\right) \sigma^2 \eta \phi M^2}{\alpha t^{\frac{2\nu}{2\nu+D}}} \\ &= \frac{16\left(\alpha^{\frac{2\nu+D}{4\nu+D}}\right)^{\frac{4\nu+D}{2\nu+D}} \ln^{\frac{4\nu+D}{2\nu+D}}\left(\sqrt{\frac{M|\mathcal{X}|}{3\delta}} \cdot t\pi\right) \sigma^2 \eta \phi M^2}{\alpha t^{\frac{2\nu}{2\nu+D}}} \\ &= \frac{16 \ln^{\frac{4\nu+D}{2\nu+D}}\left(\frac{\pi^2 M |\mathcal{X}|}{3\delta} \left(\alpha^{\frac{2\nu+D}{4\nu+D}}\right) \cdot t \left(\alpha^{\frac{2\nu+D}{4\nu+D}}\right)\right) \sigma^2 \eta \phi M^2}{\alpha t^{\frac{2\nu}{2\nu+D}}} \end{aligned}$$

The strict inequality follows from  $M \geq 1, \delta \in (0, 1), |\mathcal{X}| \geq 1$  and  $\pi > \sqrt{3}$ . Now, we can plug in  $\alpha = \frac{2\nu}{2\nu+D} \frac{4\nu+D}{2\nu+D}$  and after cancellations we get:

$$= \frac{16 \left(\frac{2\nu}{2\nu+D} - \frac{4\nu+D}{2\nu+D}\right) \ln^{\frac{4\nu+D}{2\nu+D}} \left(\left(\frac{\pi^2 M |\mathcal{X}|}{3\delta}\right)^{\left(\frac{\nu}{2\nu+D}\right)} \cdot t^{\left(\frac{2\nu}{2\nu+D}\right)}\right) \sigma^2 \eta \phi M^2}{t^{\frac{2\nu}{2\nu+D}}}$$

Now, we can plug in

$$t^{\frac{2\nu}{2\nu+D}} = \frac{32d_C^2(1)}{\epsilon^2} \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} M^2 \sigma^2 \eta \phi \cdot \ln \frac{4\nu+D}{2\nu+D} \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \frac{32d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta \phi \right)$$

to the last expression:

$$\begin{aligned} &= \left\{ 16 \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \ln \frac{4\nu+D}{2\nu+D} \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \frac{32d_C^2(1)}{\epsilon^2} \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} M^2 \sigma^2 \eta \phi \right. \right. \\ &\quad \cdot \left. \ln \frac{4\nu+D}{2\nu+D} \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \frac{32d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta \phi \right) \right\} \left. \right\} / \left\{ \frac{32d_C^2(1)}{\epsilon^2} \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \right. \\ &\quad \left. M^2 \sigma^2 \eta \phi \cdot \ln \frac{4\nu+D}{2\nu+D} \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \frac{32d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta \phi \right) \right\} \\ &= \frac{\ln \frac{4\nu+D}{2\nu+D} \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \frac{32d_C^2(1)}{\epsilon^2} \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} M^2 \sigma^2 \eta \phi \cdot \ln \frac{4\nu+D}{2\nu+D} \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \frac{32d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta \phi \right) \right)}{2 \ln \frac{4\nu+D}{2\nu+D} \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \frac{32d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta \phi \right)} \\ &\quad \cdot \frac{\epsilon^2}{d_C^2(1)} \end{aligned}$$

To show that the last expression is smaller than or equal to  $\frac{\epsilon^2}{d_C^2(1)}$ , we need to show that

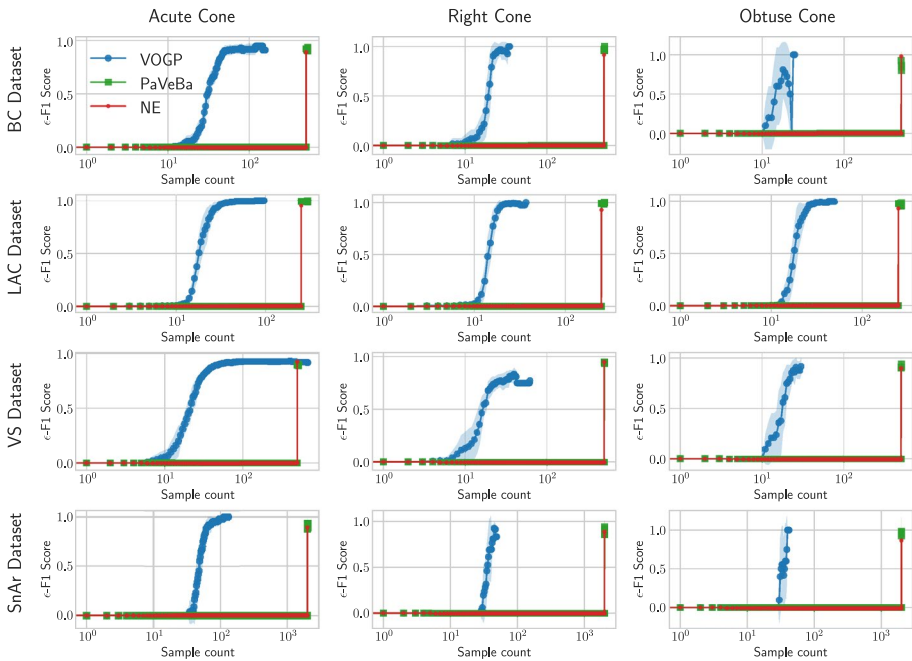
$$\ln \frac{4\nu+D}{2\nu+D} \left( B \ln \frac{4\nu+D}{2\nu+D} (B) \right) \leq 2 \ln \frac{4\nu+D}{2\nu+D} (B) = \ln \frac{4\nu+D}{2\nu+D} \left( B \left( 2^{\frac{2\nu+D}{4\nu+D}} \right) \right)$$

where  $B = \left( \left( \frac{\pi^2 M |\mathcal{X}|}{3\delta} \right)^{\left( \frac{\nu}{2\nu+D} \right)} \cdot \left( \frac{2\nu}{2\nu+D} \right)^{-\frac{4\nu+D}{2\nu+D}} \frac{32d_C^2(1)}{\epsilon^2} M^2 \sigma^2 \eta \phi \right)$  for convenience. This

holds because  $\ln \frac{4\nu+D}{2\nu+D} B \leq B \left( 2^{\frac{2\nu+D}{4\nu+D}} - 1 \right)$  for sufficiently large  $B$ . Hence, the proof is complete.  $\square$

### Appendix C Convergence Plots for Table 2

Figure 8 provides convergence plots of vector optimization methods, presenting  $\epsilon$ -F1 scores over sample counts. The experiment setup is exactly same as Table 2. Each subfigure corresponds to a specific dataset and a specific ordering cone, indicated in the subfigure titles. In these plots, we present average  $\epsilon$ -F1 scores over sample counts of different iterations. Naturally, for each sample count, only the iterations that reach that specific sample count is considered. Therefore, note that some plots illustrate only some iterations around the end. The Naive Elimination and PaVeBa algorithms require a large number of samples before making decisions, which leads to stagnant regions in the plots. In contrast, VOGP makes decisions at each round, resulting in a smoother decision trajectory and superior sample efficiency.



**Fig. 8** Convergence plots of vector optimization methods, presenting  $\epsilon$ -F1 scores over sample counts (logarithmic scale is used for sample counts)

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1007/s10994-026-07024-y>.

**Acknowledgements** This work was supported by the Scientific and Technological Research Council of Türkiye (TÜBİTAK) Grant No 121E159. İ.O.K. and Y.C.Y. were also supported by Türk Telekom as part of 5G and Beyond Joint Graduate Support Programme coordinated by Information and Communication Technologies Authority. C.T. acknowledges support by the Turkish Academy of Sciences Distinguished Young Scientist Award Program (TÜBA-GEBİP-2023) and the Scientific and Technological Research Council of Türkiye (TÜBİTAK) 2024 Incentive Award.

**Author Contributions** İ.O.K. authored most of the manuscript and derived the theoretical contributions. He also took part in development of the VOGP algorithms. İ.O.K. contributed in designing and executing the experiments and figure preparation. Y.C.Y. authored the experimental section, conducted all experiments, and significantly contributed to the manuscript editing and figure preparation. Ç.A. and C.T. provided critical oversight, verified theoretical proofs, and guided the strategic direction of the project. Both also contributed to the manuscript's writing and editing, and conceptualized the idea.

**Funding** Open access funding provided by the Scientific and Technological Research Council of Türkiye (TÜBİTAK).

**Data Availability** All data supporting the findings of this study is available at <https://github.com/Bilkent-CYBORG/VOGP>.

## Declarations

**Competing Interests** The authors declare that there are no competing interests.

**Code Availability** The code is available at <https://github.com/Bilkent-CYBORG/VOGP>.

**Ethics Approval** Not applicable.

**Consent to Participate** Not applicable.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Abdolshah, M., Shilton, A., Rana, S., Gupta, S., & Venkatesh, S. (2019). Multi-objective Bayesian optimization with preferences over objectives. *Advances in Neural Information Processing Systems*, 32, 52.
- Agrawal, A., Verschueren, R., Diamond, S., & Boyd, S. (2018). A rewriting system for convex optimization problems. *Journal of Control and Decision*, 5(1), 42–60.
- Ahmadianshalchi, A., Belakaria, S. & 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* (pp. 182–191).
- Ararat, Ç., & Tekin, C. (2023). Vector optimization with stochastic bandit feedback. In: *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, (Vol. 206, pp. 2165–2190).
- Astudillo, R., & Frazier, P. I. (2019). Multi-attribute Bayesian optimization with interactive preference learning. In: *Proceedings of The 23rd International Conference on Artificial Intelligence and Statistics*, (Vol. 108, pp. 4496–4507).
- Audiffren, J., & Ralaivola, L. (2017). Bandits dueling on partially ordered sets. *Advances in Neural Information Processing Systems*, 30, 25.
- Auer, P., Chiang, C-K., Ortner, R., & Drugan, M. (2016). Pareto front identification from stochastic bandit feedback. In: *Proceedings of The 19th International Conference on Artificial Intelligence and Statistics* (Vol. 51, pp. 939–947).
- Bader, J., & Zitzler, E. (2011). HypE: An algorithm for fast hypervolume-based many-objective optimization. *Evolutionary Computation*, 19(1), 45–76.
- Barracosa, B., Bect, J., Baraffe, H.D., Morin, J., Fournel, J., & Vazquez, E. (2022). *Bayesian multi-objective optimization for stochastic simulators: an extension of the Pareto active learning method*. arXiv preprint [arXiv:2207.03842](https://arxiv.org/abs/2207.03842)
- Batista, L. S., Campelo, F., Guimarães, F. G., & Ramírez, J. A. (2011). Pareto cone  $\epsilon$ -dominance: Improving convergence and diversity in multiobjective evolutionary algorithms. *Evolutionary Multi-Criterion Optimization*, 6576, 76–90.
- Belakaria, S., Deshwal, A., & Doppa, J.R. (2019). Max-value entropy search for multi-objective Bayesian optimization. *Advances in Neural Information Processing Systems* (Vol. 32).
- Bemporad, A., & Piga, D. (2021). Global optimization based on active preference learning with radial basis functions. *Machine Learning*, 110(2), 417–448.
- Berkenkamp, F., Schoellig, A. P., & Krause, A. (2019). No-regret bayesian optimization with unknown hyperparameters. *Journal of Machine Learning Research*, 20(50), 1–24.
- Beume, N., Naujoks, B., & Emmerich, M. (2007). SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3), 1653–1669.
- Branke, J., & Deb, K. (2005). Integrating user preferences into evolutionary multi-objective optimization. *Knowledge Incorporation in Evolutionary Computation* (pp. 461–477). Springer.
- Broxson, B.J. (2006). *The Kronecker Product* (Unpublished master's thesis). University of North Florida.
- Bubeck, S., Munos, R., Stoltz, G., & Szepesvári, C. (2011). X-armed bandits. *Journal of Machine Learning Research*, 12(5), 1655–1695.

- Candelieri, A., Ponti, A., & Archetti, F. (2024). Fair and green hyperparameter optimization via multi-objective and multiple information source Bayesian optimization. *Machine Learning*, 113(5), 2701–2731.
- Chau, S.L., Gonzalez, J., & Sejdinovic, D. (2022). Learning inconsistent preferences with gaussian processes. *International Conference on Artificial Intelligence and Statistics* (pp. 2266–2281).
- Chu, W., & Ghahramani, Z. (2005a). Preference learning with Gaussian processes. In: *Proceedings of the 22nd International Conference on Machine Learning* (pp. 137–144).
- Chu, W., & Ghahramani, Z. (2005b). Preference learning with gaussian processes. In: *Proceedings of the 22nd International Conference on Machine Learning* (pp. 137–144).
- Coello, C.A.C. (2000). Handling preferences in evolutionary multiobjective optimization: a survey. *Proceedings of the 2000 Congress on Evolutionary Computation* (pp. 30–37).
- Dang, Q., Yang, S., Liu, Q., & Ruan, J. (2024). Adaptive and communication-efficient zeroth-order optimization for distributed internet of things. *IEEE Internet of Things Journal*
- Dang, Q., Zhang, G., Wang, L., Yang, S., & Zhan, T. (2024). A generative adversarial networks model based evolutionary algorithm for multimodal multi-objective optimization. *IEEE Transactions on Emerging Topics in Computational Intelligence*,
- Dang, Q., Zhang, G., Wang, L., Yang, S., & Zhan, T. (2023). Hybrid iot device selection with knowledge transfer for federated learning. *IEEE Internet of Things Journal*, 11(7), 12216–12227.
- Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2), 182–197.
- Diamond, S., & Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 1–5.
- Drugan, M.M. (2017). PAC models in stochastic multi-objective multi-armed bandits. In: *Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 409–416).
- Emmerich, M.T.M., Deutz, A.H., & Klinkenberg, J.W. (2011). Hypervolume-based expected improvement: monotonicity properties and exact computation. *IEEE Congress of Evolutionary Computation* (pp. 2147–2154).
- Ferreira, W. M., Meneghini, I. R., Brandao, D. I., & Guimarães, F. G. (2020). Preference cone based multi-objective evolutionary algorithm applied to optimal management of distributed energy resources in microgrids. *Applied Energy*, 274, Article 115326.
- Fiedler, C., Scherer, C.W., & Trimpe, S. (2021). Practical and rigorous uncertainty bounds for gaussian process regression. In: *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 35, pp. 7439–7447).
- Garnett, R. (2023). *Bayesian Optimization*. Cambridge University Press.
- Grant, E., Pan, Y., Richardson, J., Martinelli, J.R., Armstrong, A., Galindo, A., & Adjiman, C.S. (2018). Multi-objective computer-aided solvent design for selectivity and rate in reactions. In: *13th International Symposium on Process Systems Engineering* (Vol. 44, p.2437-2442). Elsevier.
- Hakanen, J., & Knowles, J. (2017). On using decision maker preferences with ParEGO. *Evolutionary Multi-Criterion Optimization*, 10173, 282–297.
- Hamel, A.H., Heyde, F., Löhne, A., Rudloff, B., & Schrage, C. (2015). Set optimization - a rather short introduction. A.H. Hamel, F. Heyde, A. Löhne, B. Rudloff, and C. Schrage (Eds), *Set Optimization and Applications - the state of the art. From set relations to set-valued risk measures*. Springer-Verlag, Berlin.
- Hashemi, S. A., & Ashtiani, F. Z. (2010). The isomerization kinetics of lactose to lactulose in the presence of sodium hydroxide at constant and variable pH. *Food and Bioproducts Processing*, 88(2), 181–187.
- Heil, C. (2019). *Introduction to real analysis* (1st ed.). Springer Cham.
- Hernandez-Lobato, D., Hernandez-Lobato, J., Shah, A., & Adams, R. (2016). Predictive entropy search for multi-objective Bayesian optimization. In: *Proceedings of the 33rd International Conference on Machine Learning* (Vol. 48, pp. 1492–1501).
- Hone, C. A., Holmes, N., Akien, G. R., Bourne, R. A., & Muller, F. L. (2017). Rapid multistep kinetic model generation from transient flow data. *Reaction Chemistry & Engineering*, 2, 103–108.
- Hunt, B. J., Wiecek, M. M., & Hughes, C. S. (2010). Relative importance of criteria in multiobjective programming: a cone-based approach. *European Journal of Operational Research*, 207(2), 936–945.
- Ignatenko, T., Kondrashov, K., Cox, M., & de Vries, B. (2021). *On preference learning based on sequential Bayesian optimization with pairwise comparison*. arXiv preprint [arXiv:2103.13192](https://arxiv.org/abs/2103.13192)
- Jahn, J. (2010). *Vector Optimization* (2nd ed.). Springer.
- Kabanov, Y.M. (1999). *Hedging and liquidation under transaction costs in currency markets* (Vol. 3).
- Karagözlü, E. M., Yıldırım, Y. C., Ararat, Ç., & Tekin, C. (2024). Learning the Pareto set under incomplete preferences: pure exploration in vector bandits. In: *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, (Vol. 238, pp. 3070–3078).
- Katz-Samuels, J., & Scott, C. (2018). Feasible arm identification. In: *Proceedings of the 35th International Conference on Machine Learning*, (Vo. 80, pp. 2535–2543).

- Khan, F.A., Dietrich, J.P., & Wirth, C. (2022). *Efficient utility function learning for multi-objective parameter optimization with prior knowledge*. arXiv preprint [arXiv:2208.10300](https://arxiv.org/abs/2208.10300)
- Knowles, J. (2006). ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multi-objective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1), 50–66.
- Li, H., & Silva, D.L. (2008). Evolutionary multi-objective simulated annealing with adaptive and competitive search direction. In: *Proceedings of the 2008 IEEE Congress on Evolutionary Computation* (3311–3318).
- Li, C.-Y., Rakitsch, B., & Zimmer, C. (2022). Safe active learning for multi-output Gaussian processes. In: *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 151, pp. 4512–4551.
- Liao, X., Li, Q., Yang, X., Zhang, W., & 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. I., & Bakshy, E. (2022). Preference exploration for efficient Bayesian optimization with multiple outcomes. In: *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, 151, pp. 4235–4258.
- Lyu, W., Yang, F., Yan, C., Zhou, D., & Zeng, X. (2018). Multi-objective Bayesian optimization for Analog/RF circuit synthesis. In: *55th ACM/ESDA/IEEE Design Automation Conference* (pp. 1–6).
- Mathern, A., Steinholtz, O. S., Sjöberg, A., Önnheim, M., Ek, K., Rempling, R., & Jirstrand, M. (2021). Multi-objective constrained Bayesian optimization for structural design. *Structural and Multidisciplinary Optimization*, 63(2), 689–701.
- Mukaidaisi, M., Vu, A., Grantham, K., Tchagang, A., & Li, Y. (2022). Multi-objective drug design based on graph-fragment molecular representation and deep evolutionary learning. *Frontiers in Pharmacology*, 13, Article 920747.
- Müller, P., Clayton, A. D., Manson, J., Riley, S., May, O. S., Govan, N., & Bourne, R. A. (2022). Automated multi-objective reaction optimisation: which algorithm should I use? *Reaction Chemistry and Engineering*, 7(4), 987–993.
- Nika, A., Bozgan, K., Elahi, S., Ararat, Ç., & Tekin, C. (2021). *Pareto active learning with Gaussian processes and adaptive discretization*. arXiv preprint [arXiv:2006.14061](https://arxiv.org/abs/2006.14061)
- Obayashi, S., Yamaguchi, Y., & Nakamura, T. (1997). Multiobjective genetic algorithm for multidisciplinary design of transonic wing planform. *Journal of Aircraft*, 34(5), 690–693.
- Phelps, S. P., & Köksalan, M. (2003). An interactive evolutionary metaheuristic for multiobjective combinatorial optimization. *Management Science*, 49(12), 1726–1738.
- Picheny, V. (2013). Multiobjective optimization using Gaussian process emulators via stepwise uncertainty reduction. *Statistics and Computing*, 25, 1265–1280.
- Ponweiser, W., Wagner, T., Biermann, D., & Vincze, M. (2008). Multiobjective optimization on a limited budget of evaluations using model-assisted s-metric selection. *Parallel Problem Solving from Nature*, 5199, 784–794.
- Seada, H., & Deb, K. (2015). U-NSGA-III: A unified evolutionary optimization procedure for single, multiple, and many objectives: proof-of-principle results. *Evolutionary Multi-Criterion Optimization*, 9019, 34–49.
- Shu, L., Jiang, P., Shao, X., & Wang, Y. (2020). A new multi-objective Bayesian optimization formulation with the acquisition function for convergence and diversity. *Journal of Mechanical Design*, 142(9), Article 091703.
- Sinha, A., Korhonen, P., Wallenius, J., & Deb, K. (2010). An interactive evolutionary multi-objective optimization method based on polyhedral cones. *Learning and Intelligent Optimization*, 6073, 318–332.
- Srinivas, N., Krause, A., Kakade, S. M., & Seeger, M. W. (2012). Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5), 3250–3265.
- Suzuki, S., Takeno, S., Tamura, T., Shitara, K., & Karasuyama, M. (2020). Multi-objective Bayesian optimization using pareto-frontier entropy. In: *Proceedings of the 37th International Conference on Machine Learning* (Vol. 119, pp. 9279–9288).
- Svenson, J., & Santner, T. (2016). Multiobjective optimization of expensive-to-evaluate deterministic computer simulator models. *Computational Statistics & Data Analysis*, 94, 250–264.
- Takao, T., & Inagaki, A. (2022). Mono- and bis-cyclopentadienyl complexes of ruthenium and osmium. *Comprehensive Organometallic Chemistry IV*, 7, 294–443.
- Taylor, K., Ha, H., Li, M., Chan, J., & Li, X. (2021). Bayesian preference learning for interactive multi-objective optimisation. In: *Proceedings of the Companion Conference on Genetic and Evolutionary Computation* (pp. 466–475).
- Tu, B., Gandy, A., Kantas, N., & Shafei, B. (2022). Joint entropy search for multi-objective Bayesian optimization. *Advances in Neural Information Processing Systems*, 35, 9922.

- Ungredda, J., & Branke, J. (2023). When to elicit preferences in multi-objective Bayesian optimization. In: *Proceedings of the Companion Conference on Genetic and Evolutionary Computation* (pp. 1997–2003).
- Vakili, S., Khezeli, K., & Picheny, V. (2021). On information gain and regret bounds in Gaussian process bandits. In: *Proceedings of The 24th International Conference on artificial Intelligence and Statistics* (pp. 82–90).
- Yang, K., Li, L., Deutz, A., Back, T., & Emmerich, M. (2016). Preference-based multiobjective optimization using truncated expected hypervolume improvement. In: *12th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery* (pp. 276–281).
- Yang, Z., Wang, H., Yang, K., Back, T., & Emmerich, M. (2016). SMS-EMOA with multiple dynamic reference points. In: *12th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery* (pp. 282–288).
- Yue, Y., Broder, J., Kleinberg, R., & Joachims, T. (2012). The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5), 1538–1556.
- Zhang, Q., & Li, H. (2007). MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6), 712–731.
- Zhou, A., Qu, B.-Y., Li, H., Zhao, S.-Z., Suganthan, P. N., & Zhang, Q. (2011). Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation*, 1(1), 32–49.
- Ziomek, J., Adachi, M., & Osborne, M. A. (2024). Bayesian optimisation with unknown hyperparameters: Regret bounds logarithmically closer to optimal. *Advances in Neural Information Processing Systems*, 37, 86346–86374.
- Zitzler, E., Deb, K., & Thiele, L. (2000). Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation*, 8(2), 173–195.
- Zuluaga, M., Sergent, G., Krause, A., & Püschel, M. (2013). Active learning for multi-objective optimization. In: *Proceedings of the 30th International Conference on Machine Learning* (Vol. 28, pp. 462–470).
- Zuluaga, M., Krause, A., & Püschel, M. (2016). e-PAL: An active learning approach to the multi-objective optimization problem. *Journal of Machine Learning Research*, 17(104), 1–32.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.