

BAYESIAN VECTOR OPTIMIZATION WITH GAUSSIAN PROCESSES

Anonymous authors

Paper under double-blind review

ABSTRACT

Learning problems in which multiple conflicting objectives must be considered simultaneously often arise in various fields, including engineering, drug design, and environmental management. Traditional methods of multi-objective optimization, such as scalarization and identification of the Pareto set under componentwise order, have limitations in incorporating objective preferences and exploring the solution space accordingly. While vector optimization offers improved flexibility and adaptability via specifying partial orders based on ordering cones, current techniques designed for sequential experiments suffer from high sample complexity, which makes them unfit for large-scale learning problems. To address this issue, we propose VOGP, an (ϵ, δ) -PAC adaptive elimination algorithm that performs vector optimization using Gaussian processes. VOGP allows users to convey objective preferences through ordering cones while performing efficient sampling by exploiting the smoothness of the objective function, resulting in a more effective optimization process that requires fewer evaluations. We first establish provable theoretical guarantees for VOGP, and then derive information gain based and kernel specific sample complexity bounds. VOGP demonstrates strong empirical results on both real-world and synthetic datasets, outperforming previous work in sequential vector optimization and its special case multi-objective optimization. This work highlights the potential of VOGP as a powerful preference-driven method for addressing complex sequential vector optimization problems.

1 INTRODUCTION

In diverse fields such as engineering, economics, and computational biology, the problem of identifying the best set of designs across multiple 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. For instance, in aerospace engineering, the design of a wing structure must balance objectives like aerodynamic drag, wing weight, and fuel weight stored in the wing, each evaluation of which may require costly physical or computational experiments (Obayashi et al., 1997). In the realm of drug discovery, a new compound must be optimized for effectiveness, drug-likeness, and synthesizability, with each trial involving time-consuming lab tests or simulations (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 \(i.e., domination relation requires superiority in all dimensions\)](#).

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, 2011).

Let us consider an example. A farmer is planning for the next season and has two key objectives: to maximize crop yield (Y) and to minimize water usage (W). Water usage is something the farmer wants to minimize, so we can convert it into a maximization problem by considering the negative of

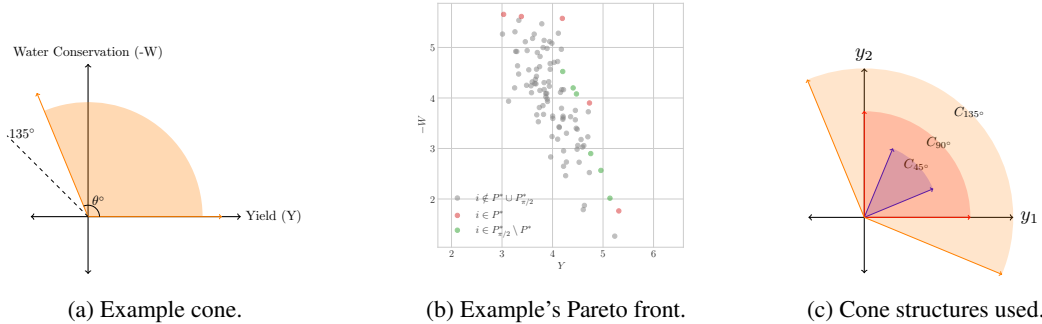


Figure 1: An example cone (1a), its Pareto front (1b) and three different cones used in experiments: C_{45° , C_{90° , and C_{135° (1c).

water usage ($-W$). The two objectives become: maximize Y and maximize $-W$. If the farmer is in a region where water resources are scarce, then they might want to prioritize water conservation over crop yield. However, they still cannot disregard yield entirely, as it directly impacts their income. Then, the use of a wide cone that has the following boundaries is appropriate: a boundary that has 0° angle with Y axis and another boundary that makes 90° to 135° degrees with Y axis. The visualization of this cone can be seen in Figure 1a. This indicates the preference for water conservation over crop yield, but not an extreme one. In Figure 1b, it can be seen that some of the designs (shown in green) in the Pareto front of traditional multiobjective order are not in the Pareto front of vector optimization. That is, they are dominated by designs that have better water conservation and slightly worse crop yield. Apart from conveying trade-off preferences of objectives, as discussed in Ararat & Tekin (2023), employing wider or narrower cones may be used as a control mechanism to tune the size of the returned Pareto optimal set. Choosing a narrower cone, as opposed to the positive orthant cone, \mathbb{R}_+^M , can potentially yield a more extensive set of Pareto optimal designs. This enriched set includes not only the designs that would be deemed optimal under the standard positive orthant cone, but also encompasses additional designs. These additional designs are superior to arbitrary selections due to their already demonstrated potential in the narrower cone.

The use of ordering cones extends the definition of the Pareto set to objective values that cannot be improved by another one in the sense of this preference cone (Boş et al., 2009). While there exists previous work addressing Pareto set identification with respect to a preference cone in noisy environments (Ararat & Tekin, 2023), they tend to require vast amounts of samples from the expensive objective function, making them impractical in real-world scenarios. In this paper, we address the problem of identifying the Pareto set with respect to a preference cone C from a discrete set of designs with a minimum number of noisy black-box function queries. We propose *Vector Optimization with Gaussian processes* (VOGP), an (ϵ, δ) -probably approximately correct (PAC) adaptive elimination algorithm that performs vector optimization by utilizing the modeling power of Gaussian processes. In particular, VOGP leverages confidence regions formed by Gaussian processes to perform cone-dependent, probabilistic elimination and Pareto identification steps, which adaptively explore the Pareto front with Bayesian optimization (Garnett, 2023). Similar to PAL and ϵ -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 componentwise order, it can be seen as a variant of PAL and ϵ -PAL that can handle dependent objectives. We prove strong theoretical guarantees on the convergence of this algorithm and derive an information gain-based sample complexity bound. Our main contributions can be listed as follows:

- We propose VOGP, an (ϵ, δ) -PAC adaptive elimination algorithm that performs vector optimization while utilizing Gaussian processes. This is the first work that considers vector optimization within the framework of Bayesian Optimization.
- We theoretically prove that VOGP returns an (ϵ, δ) -PAC Pareto set with a cone dependent sample complexity bound in terms of the maximum information gain. We also provide kernel specific sample complexity bounds in $\mathcal{O}(\cdot)$ notation.

- 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.

Related work: There is considerable existing work on Pareto set identification that utilizes Gaussian processes (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). 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 maximum value of the function (Belakaria et al., 2019). JESMO tries to sample designs that maximize the mutual information between observations and the joint distribution of the optimal points (Tu et al., 2022). Different from entropy search-based methods, some other methods utilize high-probability confidence regions formed by the Gaussian process posterior. For instance, PALS (Barracosa et al., 2022), PAL and ϵ -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 input 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 ϵ -PAL work in the fixed confidence setting and accommodates a variable sampling budget. On the other hand, entropy search-based methods are often investigated in fixed-budget setting without any theoretical guarantee on the accuracy of Pareto set identification. However, all the algorithms mentioned above try to identify the Pareto set according to the componentwise order. They are not able to capture 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) proposes 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 & Ghahramani, 2005; Hakanen & Knowles, 2017; Taylor et al., 2021; Ungredda & Branke, 2023; Ignatenko et al., 2021), where the user interacts sequentially with the algorithm to learn the user preference (Astudillo & Frazier, 2019; Lin et al., 2022). Astudillo & Frazier (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 Gaussian Process. Ahmadianshalchi et al. (2023) 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.

There are techniques that transform a multi-objective optimization problem into a single-objective problem by assigning weights to the objectives (Ponweiser et al., 2008; Knowles, 2006). 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 & Scott (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 & Tekin (2023) propose a Naïve Elimination algorithm for Pareto set identification with polyhedral ordering cones. They provide sample complexity bounds on (ϵ, δ) -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 Naive Elimination impractical in real-world problems of interest. Table 1 compares our work with the most closely related prior works in terms of key differences.

Table 1: Comparison with Related Works

Works	Cone-based preferences	Sample complexity bounds	(ϵ, δ) -PAC	Utilizes GPs
This work	Yes	Yes	Yes	Yes
Ararat & Tekin (2023)	Yes	Yes	Yes	No
Zuluaga et al. (2016)	No	Yes	Yes	Yes
Belakaria et al. (2019)	No	No	No	Yes
Hernandez-Lobato et al. (2016)	No	No	No	Yes
Tu et al. (2022)	No	No	No	Yes

2 BACKGROUND AND PROBLEM DEFINITION

We consider the problem of sequentially optimizing a vector-valued function $\mathbf{f} : \mathcal{X} \rightarrow \mathbb{R}^M$ over a finite set of designs $\mathcal{X} \subset \mathbb{R}^D$ 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 σ^2 denotes the variance of the noise. To define the optimality of designs, we use the partial order induced by C .

Partial order induced by a cone We assume that the partial order among designs is induced by a known polyhedral ordering cone C whose interior is nonempty. 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 . Rows $\mathbf{w}_1^\top, \dots, \mathbf{w}_N^\top$ are the 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: $\mathbf{y}' \preceq_C \mathbf{y} \Leftrightarrow \mathbf{w}_n^\top(\mathbf{y} - \mathbf{y}') \geq 0 \forall n \in [N] := \{1, \dots, N\}$. The Pareto set with respect to a cone C is the set of designs that are not weakly dominated by another design with respect to C : $P_C^* = \{x \in \mathcal{X} \mid \nexists x' \in \mathcal{X} \setminus \{x\} : \mathbf{f}(x) \preceq_C \mathbf{f}(x')\}$.

(ϵ, δ) -PAC Pareto set We use the cone-dependent, direction-free suboptimality gaps defined in Ararat & Tekin (2023). The gap between designs $x, x' \in \mathcal{X}$ is given as $m(x, x') := \inf\{s \geq 0 \mid \exists \mathbf{u} \in \mathbb{B}(1) \cap C : \mathbf{f}(x) + s\mathbf{u} \notin \mathbf{f}(x') - \text{int}(C)\}$, where $\text{int}(C)$ denotes the interior of C and $-$ operator over 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 (ϵ, δ) -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 1. 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.

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

M -output Gaussian process 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:

$$\begin{aligned} \boldsymbol{\mu}_t(x) &= \mathbf{k}_t(x) (\mathbf{K}_t + \sigma^2 \mathbf{I}_{Mt})^{-1} \mathbf{y}_{[t]}^\top, \\ \mathbf{k}_t(x, x') &= \mathbf{k}(x, x') - \mathbf{k}_t(x) (\mathbf{K}_t + \sigma^2 \mathbf{I}_{Mt})^{-1} \mathbf{k}_t(x')^\top. \end{aligned}$$

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.




Figure 2: Illustration of VOGP algorithm's phases.

3 THE ALGORITHM: VOGP

We propose *Vector Optimization with Gaussian processes* (VOGP), an (ϵ, δ) -PAC adaptive elimination algorithm that performs vector optimization by utilizing the modeling power of GPs. VOGP takes as input the design set \mathcal{X} , cone C , accuracy parameter ϵ , and confidence parameter δ . 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. Designs that are classified as Pareto optimal are moved to the predicted Pareto set $\hat{\mathcal{P}}_t$. 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*. An illustration of these phases is given in Figure 2. Below, we discuss each phase in detail.

Algorithm 1 VOGP Algorithm

- 1: **Input:** Design set \mathcal{X} , accuracy parameter $\epsilon \geq 0$, GP prior μ_0, σ_0 , polyhedral ordering cone C
 - 2: **Initialize:** $\mathcal{P}_1 = \emptyset, \mathcal{S}_1 = \mathcal{X}, \mathbf{R}_0(x) = \mathbb{R}^m$ for each $x \in \mathcal{S}_1, t = 1$
 - 3: **Compute:** Accuracy vector $\mathbf{u}^* \in C$
 - 4: **while** $\mathcal{S}_t \neq \emptyset$ **do**
 - 5: MODELING($\mathcal{S}_t, \mathcal{P}_t$);
 - 6: DISCARDING(\mathcal{S}_t);
 - 7: PARETOIDENTIFICATION($\mathcal{S}_t, \mathcal{P}_t$);
 - 8: EVALUATING($\mathcal{S}_t, \mathcal{W}_t$);
 - 9: $\mathcal{P}_{t+1} = \mathcal{P}_t, \mathcal{S}_{t+1} = \mathcal{S}_t$;
 - 10: $t = t + 1$;
 - 11: **end while**
 - 12: **return** Predicted Pareto set $\hat{\mathcal{P}} = \mathcal{P}_t$
-

Modeling 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 a function of δ . The probability that these hyperrectangles include the true objective values of the designs is at least $1 - \delta$, where δ is the confidence parameter of the algorithm. Using these confidence hyperrectangles allows for the formulation of probabilistic success criteria and the quantification of uncertainty with designs. The hyperrectangles are defined by (1) below, where β_t is a scalar chosen to ensure that the hyperrectangles include the true objective values at least with probability $1 - \delta$. We take $\beta_t = 2 \ln(M\pi^2|\mathcal{X}|t^2/(3\delta))$.

$$\mathcal{Q}_t(x) = \left\{ \mathbf{y} \in \mathbb{R}^M \mid \mu_t^j(x) - \beta_t^{1/2} \sigma_t^j(x) \leq y^j \leq \mu_t^j(x) + \beta_t^{1/2} \sigma_t^j(x), \forall j \in \{1, 2, \dots, M\} \right\}. \quad (1)$$

At each round, the calculated hyperrectangles are intersected cumulatively to obtain the cumulative confidence intervals for the designs. Once the algorithm has modeled the designs using confidence hyperrectangles, it proceeds to the discarding phase.

Algorithm 2 Modeling Subroutine

```

1:  $\mathcal{A}_t = \mathcal{S}_t \cup \mathcal{P}_t$ ;
2: for  $x \in \mathcal{A}_t$  do
3:   Obtain GP posterior for  $x$ :  $\boldsymbol{\mu}_t(x)$  and  $\boldsymbol{\sigma}_t(x)$ ;
4:   Construct the confidence hyperrectangle  $\mathbf{Q}_t(x)$ ;
5:    $\mathbf{R}_t(x) = \mathbf{R}_{t-1}(x) \cap \mathbf{Q}_t(x)$ ;
6: end for

```

Discarding VOGP discards undecided designs that are dominated with respect to C with high probability. In order to speed up the discarding phase, the designs that are to be discarded with low probability are identified and are used to form the pessimistic Pareto set $\mathcal{P}_{\text{pess},t}^{(C)}(\mathcal{S}_t)$ as defined below.

Definition 3 (Pessimistic Pareto Set). *Let $t \geq 1$ and let $D \subseteq \mathcal{A}_t$ be a 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 other node $x' \in D \setminus \{x\}$ such that $\mathbf{R}_t(x') \subseteq \mathbf{R}_t(x) + C$.*

VOGP computes pessimistic Pareto set of \mathcal{S}_t by checking for each design in \mathcal{S}_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 $\mathbf{R}_t(x')$ can be expressed as the sum of a vector in C and a vector in $\mathbf{R}_t(x)$.

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", 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.

The accuracy vector is of the form $\epsilon \mathbf{u}^*$, where the accuracy parameter ϵ gives its norm and \mathbf{u}^* is a unit direction vector, called the accuracy direction. The latter is selected strategically to maximize the robustness of the ordering relation induced by C . More precisely, the accuracy direction dominates the origin, even with a small perturbation in an arbitrary direction $\mathbf{u} \in \mathbb{S}^{M-1}$, where \mathbb{S}^{M-1} denotes the unit sphere in \mathbb{R}^M (see Lemma 1 below). For this purpose, \mathbf{u}^* is defined as follows:

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

Lemma 1. *For every $\mathbf{y}, \mathbf{z} \in \mathbb{R}^M$ and $\tilde{\mathbf{p}} \in \mathbb{B}(\frac{\epsilon}{d(1)})$, having $\mathbf{y} + \tilde{\mathbf{p}} \preceq_C \mathbf{z}$ implies $\mathbf{y} \preceq_C \mathbf{z} + \epsilon \mathbf{u}^*$.*

Algorithm 3 Discarding Subroutine

```

1: Compute:  $\mathcal{P}_{\text{pess},t} = \mathcal{P}_{\text{pess},t}^{(C)}(\mathcal{S}_t)$ 
2: for  $x \in \mathcal{S}_t \setminus \mathcal{P}_{\text{pess},t}$  do
3:   if  $\exists x' \in \mathcal{P}_{\text{pess},t} \forall \mathbf{y} \in \mathbf{R}_t(x) : \mathbf{R}_t(x') + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C$  then
4:      $\mathcal{S}_t = \mathcal{S}_t \setminus \{x\}$ ;
5:   end if
6: end for

```

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.

Algorithm 4 Pareto Identification Subroutine

```

1:  $\mathcal{W}_t = \mathcal{S}_t \cup \mathcal{P}_t$ ;
2: for  $x \in \mathcal{S}_t$  do
3:   if  $\nexists x' \in \mathcal{W}_t: (\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap \mathbf{R}_t(x') \neq \emptyset$  then
4:      $\mathcal{S}_t = \mathcal{S}_t \setminus \{x\}, \mathcal{P}_t = \mathcal{P}_t \cup \{x\}$ ;
5:   end if
6: end for

```

Evaluating VOGP selects the design whose confidence hyperrectangle has the widest diagonal. The diagonal of the hyperrectangle $\mathbf{R}_t(x)$ is given by $\omega_t(x) = \max_{\mathbf{y}, \mathbf{y}' \in \mathbf{R}_t(x)} \|\mathbf{y} - \mathbf{y}'\|_2^2$. 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.

Algorithm 5 Evaluating Subroutine

```

1: if  $\mathcal{S}_t \neq \emptyset$  then
2:   Select design  $x_t \in \arg \max_{x \in \mathcal{W}_t} \omega_t(x)$  (break ties randomly);
3:   Observe  $\mathbf{y}_t = \mathbf{f}(x_t) + \boldsymbol{\nu}_t$ ;
4: end if

```

4 THEORETICAL ANALYSIS

In this section, we provide theoretical guarantees for VOGP and sketch their derivation. Existing adaptive elimination methods on multi-objective optimization compare designs in componentwise order. Our proof techniques diverge from the this line of prior work by elaborately using properties of the ordering cone.

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

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

Here, we sketch the idea of the proof. The full proof can be found in supplemental document. First, we define an event E under which the confidence hyperrectangles of designs include their true objective values. We prove that $\mathbb{P}(E) \geq 1 - \delta$. We introduce a new cone-dependent accuracy vector $\epsilon \mathbf{u}^*$ as a means of performing comparisons under uncertainty. Next, we derive some useful properties of $d(1)$ and $A(1)$ (see Definition 4). Then, by leveraging cone domination properties together with the subroutines of the algorithm, we establish that the set $\hat{\mathcal{P}}$ returned by VOGP is an (ϵ, δ) -PAC Pareto set (see Definition 1). We determine the stopping criterion for VOGP in a unique way, by utilizing the uncertainty in designs based on the properties of $\epsilon \mathbf{u}^*$ (see Lemma 1). We use the stated stopping criterion to derive an upper bound on the sample complexity of VOGP.

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 or a Matérn kernel. Let $\Lambda := 4\alpha_2 d^2(1) (2 + \ln(\alpha_1))$, $\alpha_1 := M\pi^2 |\mathcal{X}| / (3\delta)$, $\alpha_2 := 16\sigma^2 \eta M^2$. Then, the sample complexity of VOGP is given by*

$$T = \mathcal{O} \left(e^{\left(\tau \sqrt{2(\ln(\tau) - \frac{\ln(\omega)}{\tau} - 1)} \right)} \cdot \frac{\tau^\tau}{\omega} \right). \quad (3)$$

Here, $\tau = D + 2$ and $\omega = \epsilon^2 / \Lambda$ when \tilde{k} is a squared exponential kernel; $\tau = (D + 2)(2\nu + D(D + 1)) / 2\nu$ and $\omega = (\epsilon^2 / \Lambda)^{\frac{2\nu + D}{2\nu}}$ when \tilde{k} is a Matérn kernel with smoothness parameter ν .

To prove Theorem 2, we use some bounds on the mutual information of the observations and the true values for an M -output GP to bound the γ_t term in (2). Then, we find an upper bound on the sample complexity which is expressed in terms of the Lambert W function. Finally, we use some tight bounds for W in Chatzigeorgiou (2013a) to bound the sample complexity of VOGP.

Remark 2. *Theorem 1 provides a bound on sample complexity that is applicable in a very general context where the kernel \mathbf{k} can be any bounded positive definite kernel. However, since it has the form of an argmin over an inequality, it may be hard to interpret. Theorem 2 provides a more explicit and practical formulation of sample complexity that shows how it scales in relation to VOGP’s parameters.*

5 EXPERIMENTS

We evaluate the performance of VOGP on two synthetic and two real-world datasets. We investigate its adherence to the proven theoretical guarantees ([success rates SR1, SR2](#)), [function evaluation count \(SC\)](#), and Pareto statistics (namely, Pareto accuracy (PA), recall (PR), and precision (PP) values in classifying the designs as Pareto or not.) We will investigate the performance of VOGP under C_{45° , C_{90° , and C_{135° in comparison to the state of the art method for vector optimization: Naïve Elimination (NE). There are many established multi-objective optimization (MOO) methods mentioned in the related works section. However, adapting these methods to general cone structures is non-trivial. A comparison with these methods is still possible since VOGP can be used for MOO purposes by choosing $C = C_{90^\circ}$. Therefore, we provide an extensive comparison of VOGP with other state-of-the-art MOO methods in the context of C_{90° in the supplemental document (Section A.1). These methods are MESMO, JESMO, PESMO and EHVI (Daulton et al., 2020). The reported results are obtained by repeating all experiments 10 times and averaging the scores. The results of the experiments can be seen in Table 2.

- **SNW dataset** This dataset was introduced by Zuluaga et al. (2012). The objective here is to optimize the area and throughput for the synthesis of a field-programmable gate array (FPGA) platform ($M = 2, D = 3$). The dataset comprises 206 distinct hardware design implementations of a sorting network.
- **BC (Branin Currin) dataset** For this dataset, we utilize two commonly used benchmark functions, namely Branin and Currin, each serving as an output dimension. The input space for this dataset is two-dimensional ($M = 2, D = 2$).
- **OKA dataset** This dataset involves the evaluation of two functions as defined in Okabe et al. (2004) ($M = 2, D = 3$).
- **SnAr (chemical reaction) dataset** This benchmark examines the nucleophilic aromatic substitution reaction (SnAr) involving 2,4-difluoronitrobenzene and pyrrolidine in ethanol, which results in a primary product and two additional side-products (Hone et al., 2017). The objectives are the space-time yield and the environmental impact. The dataset comprises 950 designs ($M = 2, D = 4$).
- **MAR (marine) dataset** This dataset focuses on optimizing bulk carrier architectures while obeying specific restrictions required for vessels navigating the Panama Canal (Parsons & Scott, 2004). The objective includes maximizing annual cargo, while minimizing transportation cost and ship weight. The mentioned constraints are also transformed into objectives. The dataset comprises of 1000 designs ($M = 4, D = 6$).

Definition 5. *Given an angle θ° , 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.*

Definition 6. *Given a finite set of designs \mathcal{X} , set of Pareto designs P_θ^* and a set of predicted Pareto designs P_θ with respect to the cone C_{θ° , Pareto accuracy (PA), Pareto recall (PR), and Pareto precision (PP), rate of satisfying success conditions (i) (SR1) and (ii) (SR2) are defined as*

$$\text{PA} := \frac{|P_\theta^* \cap P_\theta| + |(\mathcal{X} \setminus P_\theta^*) \cap (\mathcal{X} \setminus P_\theta)|}{|\mathcal{X}|} \times 100, \quad \text{PR} := \frac{|P_\theta^* \cap P_\theta|}{|P_\theta^*|} \times 100, \quad \text{PP} := \frac{|P_\theta^* \cap P_\theta|}{|P_\theta|} \times 100$$

$$\text{SR1} := \frac{|\{x \in P_\theta^* : \mathbf{f}(x^*) \in \bigcup_{x \in P_\theta} (\mathbf{f}(x) + \mathbb{B}(\epsilon) \cap C - C)\}|}{|P_\theta^*|}, \quad \text{SR2} := \frac{|\{x \in P_\theta : \Delta_x^* \leq 2\epsilon\}|}{|P_\theta|}.$$

Experimental Setup Prior to the experiments, we normalize the datasets for ϵ parameter to be applicable. Since we normalize the datasets, we apply a noise standard deviation of 0.1. We fix $\epsilon = 0.1$ and $\delta = 0.05$. Similar to the other works with the same confidence hyperrectangle definition, we scale down β_t by a factor of 20 (Zuluaga et al. (2016)). Initially, for each dataset, we learn the kernel hyperparameters by training on the entire dataset. This is done to ensure that the algorithm

capture the smoothness of the dataset correctly. Data used in this phase is not used to update the GP posterior. All algorithms start the learning phase with an empty set of observations. In this experiment, for OKA dataset, Matérn kernel was used. For SNW, BC, SnAr datasets, RBF kernel was used. NE is a fixed budget algorithm and it takes at least one sample from each design. Prior work derived a theoretically sufficient sampling budget for NE under which an (ϵ, δ) -PAC Pareto set is guaranteed to be returned; however, this budget is too large to be useful in practice. To make results comparable, NE was given a sample-per-design budget that makes total number of evaluations as close as possible $(\lceil T/|\mathcal{X}| \rceil)$.

Discussion As seen in Table 2, despite the disparity in sampling budgets, VOGP performs very similarly but moderately below NE in terms of Pareto accuracy (PA), Pareto precision (PP), and Pareto recall (PR). VOGP performs marginally below NE in terms of satisfying the (ϵ, δ) -PAC conditions. These can be attributed to the disparity between the sample counts between NE and VOGP. Since NE takes at least one sample from each design, in cases where VOGP samples less designs than $|\mathcal{X}|$, NE gets to sample more designs than VOGP. As can be confirmed from Table 2, this is a lot more than VOGP’s total sample count for most cases; specifically, for all cases except when SNW dataset is used with C_{45° . As an example, for SnAr dataset with C_{135° , VOGP had a total budget of around 33 whereas NE had 950. As can be seen in Table 2, despite NE’s advantage of sample counts, VOGP was capable of achieving superior (ϵ, δ) -PAC success rates for multiple settings and comparable Pareto statistics while requiring significantly fewer samples in almost all of the cases. This is an expected result because VOGP utilizes an M -output GP to obtain information from close observations while also taking correlations of output dimensions into account whereas NE uses empirical means only.

Table 2: Performance comparison between VOGP (Our method) and Naive Elimination under C_{45° , C_{90° , and C_{135° . SR1 and SR2: The success rate of satisfying (ϵ, δ) -PAC conditions (i) and (ii) respectively. PA: Pareto accuracy rate. PR: Pareto recall rate. PP: Pareto precision rate. SC: The number of evaluations.

C	D	VOGP					SC	Naive Elimination					
		SR1 \uparrow	SR2 \uparrow	PA \uparrow	PR \uparrow	PP \uparrow		SR1 \uparrow	SR2 \uparrow	PA \uparrow	PR \uparrow	PP \uparrow	SC
C_{45°	SNW	95.96 \pm 2.78	98.65 \pm 1.46	85.53 \pm 1.35	76.54 \pm 4.93	69.49 \pm 3.25	777.8 \pm 220.65	96.92 \pm 3.87	100.00 \pm 0.00	88.40 \pm 2.42	72.31 \pm 6.50	79.95 \pm 5.21	3 \times 250
C_{45°	BC	95.17 \pm 7.43	99.37 \pm 1.27	95.36 \pm 1.59	86.90 \pm 8.13	76.56 \pm 6.35	446.40 \pm 131.14	91.03 \pm 3.52	99.63 \pm 1.11	92.60 \pm 1.26	67.93 \pm 8.45	68.25 \pm 5.16	2 \times 250
C_{45°	OKA	88.00 \pm 7.38	92.84 \pm 5.41	95.28 \pm 0.99	76.40 \pm 4.54	76.87 \pm 6.89	650.3 \pm 153.36	98.00 \pm 2.00	100.00 \pm 0.00	96.56 \pm 1.11	82.40 \pm 7.42	83.77 \pm 7.61	3 \times 250
C_{45°	SnAr	78.75 \pm 14.03	98.00 \pm 6.00	98.46 \pm 0.42	55.62 \pm 16.64	53.63 \pm 12.25	517.0 \pm 148.79	88.12 \pm 5.90	97.86 \pm 3.27	98.56 \pm 0.22	55.00 \pm 6.12	58.52 \pm 8.07	1 \times 950
C_{90°	SNW	94.62 \pm 4.28	97.5 \pm 2.22	89.85 \pm 2.14	63.85 \pm 10.35	59.61 \pm 8.75	112.70 \pm 28.49	92.69 \pm 7.78	99.23 \pm 1.54	90.00 \pm 1.19	56.54 \pm 7.10	61.23 \pm 4.96	1 \times 250
C_{90°	BC	98.33 \pm 5.00	100.0 \pm 0.00	99.20 \pm 0.59	80.00 \pm 20.82	88.33 \pm 14.53	33.60 \pm 2.73	98.33 \pm 5.00	100.00 \pm 0.00	98.32 \pm 0.61	68.33 \pm 17.40	67.03 \pm 15.97	1 \times 250
C_{90°	OKA	87.14 \pm 4.29	99.09 \pm 2.73	96.24 \pm 0.57	57.14 \pm 6.39	38.91 \pm 5.84	110.60 \pm 16.38	90.00 \pm 11.16	98.33 \pm 5.00	97.96 \pm 0.66	60.00 \pm 16.66	68.93 \pm 16.40	1 \times 250
C_{90°	SnAr	67.5 \pm 40.39	95.00 \pm 10.00	99.46 \pm 0.14	40.00 \pm 25.50	34.42 \pm 21.34	39.80 \pm 20.26	95.00 \pm 15.00	98.57 \pm 4.29	99.43 \pm 0.18	55.00 \pm 10.00	40.19 \pm 11.51	1 \times 950
R_1^+	MAR	100.00 \pm 0.00	99.67 \pm 0.81	97.70 \pm 0.31	79.32 \pm 2.96	66.84 \pm 4.29	716.43 \pm 173.89	98.42 \pm 1.29	99.00 \pm 1.22	96.88 \pm 0.51	71.58 \pm 5.98	57.53 \pm 6.21	1 \times 1000
C_{135°	SNW	92.00 \pm 9.80	94.39 \pm 6.41	96.12 \pm 1.76	72.00 \pm 12.49	60.07 \pm 15.58	72.80 \pm 11.91	89.00 \pm 11.36	100.00 \pm 0.00	96.41 \pm 1.13	63.00 \pm 13.45	64.19 \pm 13.28	1 \times 250
C_{135°	BC	100.00 \pm 0.00	100.00 \pm 0.00	99.88 \pm 0.18	85.00 \pm 22.91	100.00 \pm 0.00	17.50 \pm 2.77	100.00 \pm 0.00	100.00 \pm 0.00	99.76 \pm 0.20	85.00 \pm 22.91	90.00 \pm 15.28	1 \times 250
C_{135°	OKA	65.00 \pm 22.91	85.00 \pm 22.91	99.60 \pm 0.31	65.00 \pm 22.91	85.00 \pm 22.91	44.60 \pm 5.22	95.00 \pm 15.00	100.00 \pm 0.00	99.96 \pm 0.12	95.00 \pm 15.00	100.00 \pm 0.00	1 \times 250
C_{135°	SnAr	73.33 \pm 13.33	100.0 \pm 0.0	99.64 \pm 0.10	26.67 \pm 20.00	38.33 \pm 30.78	32.90 \pm 9.41	83.33 \pm 30.73	100.00 \pm 0.00	99.75 \pm 0.15	43.33 \pm 26.03	65.00 \pm 39.76	1 \times 950

6 CONCLUSION, LIMITATIONS AND FUTURE RESEARCH

We consider black-box vector optimization with noisy evaluations where the objective function is expensive to evaluate. We propose a sample-efficient adaptive elimination algorithm, VOGP. We prove that VOGP returns an (ϵ, δ) -PAC Pareto set and derive information gain-based and kernel-specific sample complexity bounds.

In our experiments, we observe that VOGP outperforms NE in terms of sampling budget needed to return an (ϵ, δ) -PAC Pareto set. Furthermore, VOGP requires significantly fewer function evaluations to obtain Pareto statistics similar to those of NE. In the experiments given in the supplemental document, we observe that VOGP is the most consistently high-performing method among other MOO methods across the datasets and metrics presented. In real-world datasets (SnAr, SNW), the lead of VOGP over MOO methods (JESMO, MESMO, PESMO, EHVI) is amplified. In conclusion, VOGP surpasses the state-of-the-art methods of vector optimization and the specific case of multi-objective optimization in terms of the needed function evaluation count and returned Pareto set quality.

Our work opens up many research directions in which vector optimization with Gaussian Processes can be improved. One of those directions is to extend VOGP to continuous design setting, which would require a more nuanced theoretical analysis and incorporating techniques like adaptive discretization. Another direction worth exploring is extending VOGP to the case of missing rewards where observations have incomplete objective values. This could broaden its applicability to real-world problems where data might be partial or missing.

REFERENCES

- Majid Abdolshah, Alistair Shilton, Santu Rana, Sunil Gupta, and Svetha Venkatesh. Multi-objective Bayesian optimisation with preferences over objectives. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Alaleh Ahmadianshalchi, Syrine Belakaria, and Janardhan Rao Doppa. Preference-aware constrained multi-objective Bayesian optimization. In *Proceedings of the Women in Machine Learning Workshop at Advances in Neural Information Processing Systems Conference*, 2023.
- Çağın Ararat and Cem Tekin. Vector optimization with stochastic bandit feedback. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 2165–2190, 25–27 Apr 2023.
- Raul Astudillo and Peter I. Frazier. Multi-attribute Bayesian optimization with interactive preference learning. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- Johannes Bader and Eckart Zitzler. HypE: An algorithm for fast hypervolume-based many-objective optimization. *Evolutionary Computation*, 19(1):45–76, 2011.
- Bruno Barracosa, Julien Bect, Héloïse Dutrieux Baraffe, Juliette Morin, Josselin Fournel, and Emmanuel Vazquez. Bayesian multi-objective optimization for stochastic simulators: an extension of the pareto active learning method. *arXiv preprint arXiv:2207.03842*, 2022.
- Lucas S. Batista, Felipe Campelo, Frederico G. Guimarães, and Jaime A. Ramírez. Pareto cone ϵ -dominance: Improving convergence and diversity in multiobjective evolutionary algorithms. In *Evolutionary Multi-Criterion Optimization*, pp. 76–90, 2011.
- Syrine Belakaria, Aryan Deshwal, and Janardhan Rao Doppa. Max-value entropy search for multi-objective Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Nicola Beume, Boris Naujoks, and Michael Emmerich. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3):1653–1669, 2007.
- Radu Ioan Bot, Sorin-Mihai Grad, and Gert Wanka. Duality in vector optimization. In *Vector Optimization*, 2009.
- Jürgen Branke and Kalyanmoy Deb. *Integrating User Preferences into Evolutionary Multi-Objective Optimization*, pp. 461–477. 2005.
- Bobbi Jo Broxson. The kronecker product, 2006.
- Ioannis Chatzigeorgiou. Bounds on the Lambert function and their application to the outage analysis of user cooperation. *IEEE Communications Letters*, 17(8):1505–1508, 2013a.
- Ioannis Chatzigeorgiou. Bounds on the lambert function and their application to the outage analysis of user cooperation. *IEEE Communications Letters*, 17(8):1505–1508, aug 2013b.
- Wei Chu and Zoubin Ghahramani. Preference learning with Gaussian processes. In *Proceedings of the 22nd International Conference on Machine Learning*, ICML '05, pp. 137–144, 2005.
- C.A.C. Coello. Handling preferences in evolutionary multiobjective optimization: a survey. In *Proceedings of the 2000 Congress on Evolutionary Computation. CEC00 (Cat. No.00TH8512)*, volume 1, pp. 30–37 vol.1, 2000.
- Samuel Daulton, Maximilian Balandat, and Eytan Bakshy. Differentiable expected hypervolume improvement for parallel multi-objective bayesian optimization. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 9851–9864. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/6fec24eac8f18ed793f5eaad3dd7977c-Paper.pdf.

- Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.
- Michael T. M. Emmerich, André H. Deutz, and Jan Willem Klinkenberg. Hypervolume-based expected improvement: Monotonicity properties and exact computation. In *2011 IEEE Congress on Evolutionary Computation (CEC)*, pp. 2147–2154, 2011. doi: 10.1109/CEC.2011.5949880.
- Willian M. Ferreira, Ivan R. Meneghini, Danilo I. Brandao, and Frederico G. Guimarães. Preference cone based multi-objective evolutionary algorithm applied to optimal management of distributed energy resources in microgrids. *Applied Energy*, 274:115326, 2020.
- Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2023.
- Jussi Hakanen and Joshua Knowles. On using decision maker preferences with parego. In *Evolutionary Multi-Criterion Optimization*, volume 10173 of *Lecture Notes in Computer Science*, pp. 282–297, 2017. doi: 10.1007/978-3-319-54157-0_20. 9th International Conference on Evolutionary Multi-Criterion Optimization ; Conference date: 19-03-2017 Through 22-03-2017.
- Daniel Hernandez-Lobato, Jose Hernandez-Lobato, Amar Shah, and Ryan Adams. Predictive entropy search for multi-objective Bayesian optimization. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1492–1501, 20–22 Jun 2016.
- Christopher A. Hone, Nicholas Holmes, Geoffrey R. Akien, Richard A. Bourne, and Frans L. Muller. Rapid multistep kinetic model generation from transient flow data. *React. Chem. Eng.*, 2:103–108, 2017.
- Tanya Ignatenko, Kirill Kondrashov, Marco Cox, and Bert de Vries. On preference learning based on sequential bayesian optimization with pairwise comparison. *arXiv preprint arXiv:2103.13192*, 2021.
- Johannes Jahn. *Vector Optimization*. 2nd edition, 2011.
- Julian Katz-Samuels and Clay Scott. Feasible arm identification. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2535–2543, 10–15 Jul 2018.
- Farha A Khan, Jörg P Dietrich, and Christian Wirth. Efficient utility function learning for multi-objective parameter optimization with prior knowledge. *arXiv preprint arXiv:2208.10300*, 2022.
- Joshua Knowles. ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation*, 10(1): 50–66, 2006.
- Cen-You Li, Barbara Rakitsch, and Christoph Zimmer. Safe active learning for multi-output Gaussian processes. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 4512–4551, 28–30 Mar 2022.
- Hui Li and Dario Landa Silva. Evolutionary multi-objective simulated annealing with adaptive and competitive search direction. *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*, pp. 3311–3318, 2008.
- Zhiyuan Jerry Lin, Raul Astudillo, Peter I. Frazier, and Eytan Bakshy. Preference exploration for efficient Bayesian optimization with multiple outcomes. *ArXiv*, abs/2203.11382, 2022.
- Wenlong Lyu, Fan Yang, Changhao Yan, Dian Zhou, and Xuan Zeng. Multi-objective Bayesian optimization for Analog/RF circuit synthesis. In *Proc. 55th Annual Design Automation Conference*, pp. 1–6, 2018.
- Alexandre Mathern, Olof Skogby Steinholtz, Anders Sjöberg, Magnus Önnheim, Kristine Ek, Rasmus Rempling, Emil Gustavsson, and Mats Jirstrand. Multi-objective constrained Bayesian optimization for structural design. *Structural and Multidisciplinary Optimization*, 63(2):689–701, Feb 2021.

- Muhetaer Mukaidaisi, Andrew Vu, Karl Grantham, Alain Tchagang, and Yifeng Li. Multi-objective drug design based on graph-fragment molecular representation and deep evolutionary learning, May 2022.
- Andi Nika, Kerem Bozgan, Sepehr Elahi, Çağın Ararat, and Cem Tekin. Pareto active learning with Gaussian processes and adaptive discretization, 2021.
- Shigeru Obayashi, Yoshihiro Yamaguchi, and Takashi Nakamura. Multiobjective genetic algorithm for multidisciplinary design of transonic wing planform. *Journal of Aircraft*, 34(5):690–693, 1997.
- Tatsuya Okabe, Yaochu Jin, Markus Olhofer, and Bernhard Sendhoff. On test functions for evolutionary multi-objective optimization. In *Parallel Problem Solving from Nature - PPSN VIII*, pp. 792–802, 2004.
- Michael G. Parsons and Randall L. Scott. Formulation of Multicriterion Design Optimization Problems for Solution With Scalar Numerical Optimization Methods. *Journal of Ship Research*, 48(01):61–76, 03 2004. doi: 10.5957/jsr.2004.48.1.61.
- Selcen (Pamuk) Phelps and Murat Köksalan. An interactive evolutionary metaheuristic for multiobjective combinatorial optimization. *Management Science*, 49(12):1726–1738, 2003.
- Victor Picheny. Multiobjective optimization using gaussian process emulators via stepwise uncertainty reduction. *Statistics and Computing*, 25:1265–1280, 2013.
- Wolfgang Ponweiser, Tobias Wagner, Dirk Biermann, and Markus Vincze. Multiobjective optimization on a limited budget of evaluations using model-assisted s-metric selection. In *Parallel Problem Solving from Nature – PPSN X*, pp. 784–794, 2008.
- Haitham Seada and Kalyanmoy Deb. U-NSGA-III: A unified evolutionary optimization procedure for single, multiple, and many objectives: Proof-of-principle results. In *Evolutionary Multi-Criterion Optimization*, pp. 34–49, 2015.
- Leshi Shu, Ping Jiang, Xinyu Shao, and Yan Wang. A New Multi-Objective Bayesian Optimization Formulation With the Acquisition Function for Convergence and Diversity. *Journal of Mechanical Design*, 142(9), 03 2020.
- Ankur Sinha, Pekka Korhonen, Jyrki Wallenius, and Kalyanmoy Deb. An interactive evolutionary multi-objective optimization method based on polyhedral cones. In *Learning and Intelligent Optimization*, pp. 318–332, 2010.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias W Seeger. Information-theoretic regret bounds for Gaussian process optimization in the bandit setting. *IEEE Transactions on Information Theory*, 58(5):3250–3265, 2012.
- Shinya Suzuki, Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Multi-objective Bayesian optimization using pareto-frontier entropy. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 9279–9288, 13–18 Jul 2020.
- Joshua Svenson and Thomas Santner. Multiobjective optimization of expensive-to-evaluate deterministic computer simulator models. *Computational Statistics & Data Analysis*, 94:250–264, 2016.
- Kendall Taylor, Huong Ha, Minyi Li, Jeffrey Chan, and Xiaodong Li. Bayesian preference learning for interactive multi-objective optimisation. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO ’21*, pp. 466–475, 2021. doi: 10.1145/3449639.3459299.
- Ben Tu, Axel Gandy, Nikolas Kantas, and Behrang Shafei. Joint entropy search for multi-objective Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 35, pp. 9922–9938, 2022.
- Juan Ungredda and Juergen Branke. When to elicit preferences in multi-objective bayesian optimization. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation, GECCO ’23 Companion*, pp. 1997–2003. Association for Computing Machinery, 2023. doi: 10.1145/3583133.3596342.

- Sattar Vakili, Nacime Bouziani, Sepehr Jalali, Alberto Bernacchia, and Da shan Shiu. Optimal order simple regret for gaussian process bandits, 2021a.
- Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in gaussian process bandits. In *Proceedings of International Conference on Artificial Intelligence and Statistics*, pp. 82–90, 2021b.
- Kaifeng Yang, Longmei Li, André Deutz, Thomas Back, and Michael Emmerich. Preference-based multiobjective optimization using truncated expected hypervolume improvement. In *2016 12th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD)*, pp. 276–281, 2016a. doi: 10.1109/FSKD.2016.7603186.
- Zhiwei Yang, Hao Wang, Kaifeng Yang, Thomas Back, and Michael Emmerich. SMS-EMOA with multiple dynamic reference points. In *2016 12th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD)*, pp. 282–288, 2016b.
- Qingfu Zhang and Hui Li. MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007.
- Aimin Zhou, Bo-Yang Qu, Hui Li, Shi-Zheng Zhao, Ponnuthurai Nagarathnam Suganthan, and Qingfu Zhang. Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation*, 1(1):32–49, 2011.
- Marcela Zuluaga, Peter Milder, and Markus Püschel. Computer generation of streaming sorting networks. In *DAC Design Automation Conference 2012*, pp. 1241–1249, 2012.
- Marcela Zuluaga, Guillaume Sergent, Andreas Krause, and Markus Püschel. Active learning for multi-objective optimization. In *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pp. 462–470, 17–19 Jun 2013.
- Marcela Zuluaga, Andreas Krause, and Markus Püschel. e-PAL: An active learning approach to the multi-objective optimization problem. *Journal of Machine Learning Research*, 17(104):1–32, 2016.

SUPPLEMENTAL DOCUMENT

Overview In Section A, we present additional experiments by providing the comparison of VOGP with other MOO methods. In Section B, we provide the comparison of VOGP and random search. In Section C, we present an ablation study where the absence of known kernel hyperparameters assumption is investigated. In Section D, we provide an extension of our related works that encapsulates related works on evolutionary algorithms. In Section E, we provide the derivations of Theorems 1, 2.

A COMPARISON OF VOGP WITH OTHER MOO METHODS

A.1 COMPARISON OF VOGP WITH OTHER MOO METHODS UNDER $C = C_{90^\circ}$

Here, we undertake a comparative study of the VOGP algorithm and other MOO methods within the framework of a 90° ordering cone. These methods are MESMO, JESMO, PESMO and EHVI. Our objective is to compare the performance of algorithms in terms of their sample complexity and the quality of the resulting Pareto front.

Table 3: Comparison of VOGP with other MOO methods under C_{90° . SR1 and SR2: The success rate of satisfying (ϵ, δ) -PAC conditions (i) and (ii) respectively. HV: Hypervolume. PA: Pareto accuracy rate. PR: Pareto recall rate. PP: Pareto precision rate. SC: The number of evaluations.

Method	Dataset	SR1 \uparrow	SR2 \uparrow	HV \uparrow	PA \uparrow	PR \uparrow	PP \uparrow	SC \downarrow
VOGP (Ours)	SNW	94.62 \pm 4.28	97.5 \pm 2.22	45.89 \pm 0.09	89.85 \pm 2.14	63.85 \pm 10.35	59.61 \pm 8.75	112.70 \pm 28.49
MESMO	SNW	89.23 \pm 6.62	92.7 \pm 3.27	45.84 \pm 0.12	88.40 \pm 1.40	56.54 \pm 7.89	53.77 \pm 5.33	112
EHVI	SNW	26.54 \pm 16.80	46.62 \pm 25.92	32.09 \pm 5.10	72.57 \pm 12.83	16.92 \pm 13.01	14.25 \pm 13.07	112
JESMO	SNW	0.00 \pm 0.00	1.43 \pm 4.29	27.48 \pm 1.27	85.87 \pm 0.91	0.00 \pm 0.00	0.00 \pm 0.00	112
PESMO	SNW	0.00 \pm 0.00	0.00 \pm 0.00	26.34 \pm 2.15	86.12 \pm 0.73	0.00 \pm 0.00	0.00 \pm 0.00	112
VOGP (Ours)	BC	98.33 \pm 5.00	100.0 \pm 0.00	42.66 \pm 0.06	99.20 \pm 0.59	80.00 \pm 20.82	88.33 \pm 14.53	33.6 \pm 2.73
MESMO	BC	78.33 \pm 29.86	94.24 \pm 8.94	42.59 \pm 0.14	97.36 \pm 1.78	65.00 \pm 28.33	50.12 \pm 24.41	33
EHVI	BC	36.67 \pm 22.11	90.67 \pm 19.6	42.14 \pm 0.29	97.08 \pm 0.74	31.67 \pm 8.98	40.02 \pm 16.70	33
JESMO	BC	46.67 \pm 26.67	100.0 \pm 0.00	42.41 \pm 0.09	96.92 \pm 0.44	36.67 \pm 6.67	36.57 \pm 7.79	33
PESMO	BC	60.00 \pm 32.66	100.0 \pm 0.00	42.45 \pm 0.12	96.92 \pm 0.54	41.67 \pm 11.18	37.89 \pm 7.62	33
VOGP (Ours)	OKA	87.14 \pm 4.29	99.09 \pm 2.73	51.68 \pm 0.02	96.24 \pm 0.57	57.14 \pm 6.39	38.91 \pm 5.84	110.60 \pm 16.38
MESMO	OKA	77.14 \pm 7.0	98.89 \pm 3.33	50.39 \pm 1.96	96.96 \pm 0.86	48.57 \pm 9.48	48.17 \pm 13.87	110
EHVI	OKA	88.57 \pm 5.71	100.0 \pm 0.00	51.72 \pm 0.03	97.76 \pm 0.82	65.71 \pm 9.48	60.85 \pm 14.92	110
JESMO	OKA	90.00 \pm 6.55	100.0 \pm 0.00	51.70 \pm 0.03	97.28 \pm 0.5	62.86 \pm 7.00	51.66 \pm 7.27	110
PESMO	OKA	87.14 \pm 7.69	100.0 \pm 0.00	51.62 \pm 0.20	96.96 \pm 0.65	61.43 \pm 9.15	47.50 \pm 9.88	110
VOGP (Ours)	SnAr	67.5 \pm 40.39	95.00 \pm 10.00	48.65 \pm 1.17	99.46 \pm 0.14	40.00 \pm 25.50	34.42 \pm 21.34	39.80 \pm 20.26
MESMO	SnAr	15.0 \pm 25.5	56.0 \pm 35.77	45.69 \pm 1.76	99.39 \pm 0.15	5.00 \pm 10.00	13.33 \pm 30.55	39
EHVI	SnAr	0.00 \pm 0.00	0.00 \pm 0.00	14.29 \pm 0.00	99.47 \pm 0.00	0.00 \pm 0.0	0.0 \pm 0.00	39
JESMO	SnAr	0.00 \pm 0.00	0.00 \pm 0.00	14.29 \pm 0.00	99.47 \pm 0.00	0.00 \pm 0.0	0.0 \pm 0.00	39
PESMO	SnAr	0.00 \pm 0.00	0.00 \pm 0.00	14.29 \pm 0.00	99.47 \pm 0.00	0.00 \pm 0.0	0.0 \pm 0.00	39

Experimental Setup VOGP algorithm was run with the same setup as in the main paper’s experiments section. As in the case of VOGP, for MOO methods, initially we learn the kernel hyperparameters by training on the entire dataset. This is done to ensure that the algorithms capture the smoothness of the dataset correctly. Data used in this phase is not used to update the GP posterior. All algorithms start the learning phase with an empty set of observations. MOO methods operates in fixed budget setting. To have an accurate comparison between VOGP and MOO algorithms, we set the budget of MOO algorithms based on the average number of evaluations of VOGP in each dataset. All methods used Matérn kernel for OKA dataset and RBF kernel for other datasets. After the optimization loop of MOO methods is finished, their observations are used to train a GP, and the posterior mean values for designs in the design set are used to calculate the predicted Pareto set for these MOO methods.

Discussion As seen in Table 3, VOGP consistently ranks among the top-performing methods across datasets and metrics. For BC, SNW and SnAr datasets, VOGP surpasses other methods in every metric. JESMO performs well in the OKA dataset, especially with SR. EHVI performs well in OKA dataset and has the highest Pareto statistics. MESMO generally has decent performances but does not lead as strongly as VOGP in most cases. In conclusion, VOGP is the most consistently high-performing method across the datasets and metrics presented. In real-world datasets (SnAr, SNW), the lead of VOGP over MOO methods is amplified and PESMO and JESMO fail to return good designs.

The superiority of VOGP over other MOO methods can be attributed to multiple factors. First, VOGP guarantees an (ϵ, δ) -PAC Pareto set to be returned, whereas the compared MOO methods do not have this guarantee. Second, VOGP algorithm does maximum variance reduction-based sampling. In Vakili et al. (2021a), it was shown that tighter than usual confidence bounds hold for sampling methods where noise history is independent of sampling history. For instance, a purely non-adaptive variance reduction-based algorithm satisfies the conditions for the mentioned tight bounds to hold since it samples without looking at the observations. While VOGP performs variance reduction in a similar spirit, it also performs adaptive elimination, which may prevent a design from getting sampled if the design is discarded based on the noise history. Therefore, these tight bounds are not certified for VOGP. However, it is important to observe that at the beginning of the algorithm, the sampling strategy of VOGP is very close to that of a purely non-adaptive variance reduction-based sampling algorithm. Since we empirically scale VOGP’s confidence region in the experiments, VOGP might be getting the benefit of (almost) valid tighter confidence regions, thereby improving its performance.

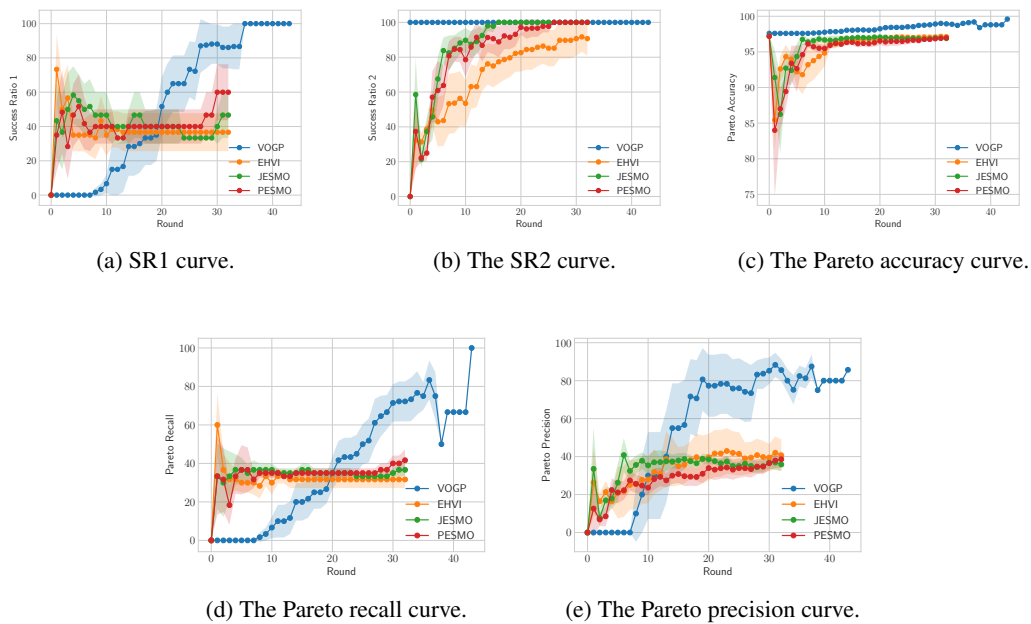


Figure 3: The scores of VOGP, JESMO, EHVI, and PESMO on BC dataset with C_{90° cone. Since VOGP’s runs are of different round counts, averages were taken over possible runs.

A.2 COMPARISON OF VOGP WITH OTHER MOO METHODS UNDER $C = C_{135^\circ}$

In this section, we compare VOGP run with $C = C_{135^\circ}$ to other MOO methods under with the metrics defined over C_{135° .

Table 4: Comparison of VOGP with other MOO methods under C_{135° . SR1 and SR2: The success rate of satisfying (ϵ, δ) -PAC conditions (i) and (ii) respectively. PA: Pareto accuracy rate. PR: Pareto recall rate. PP: Pareto precision rate. SC: The number of evaluations.

Method	Dataset	SR1 \uparrow	SR2 \uparrow	PA \uparrow	PR \uparrow	PP \uparrow	SC \downarrow
VOGP (Ours)	SNW	92.00 \pm 9.80	94.39 \pm 6.41	96.12 \pm 1.76	72.0 \pm 12.49	60.07 \pm 15.58	72.8 \pm 11.91
MESMO	SNW	80.00 \pm 14.14	80.31 \pm 5.91	94.22 \pm 1.86	63.0 \pm 20.52	43.79 \pm 11.95	72
EHVI	SNW	0.00 \pm 0.00	50.00 \pm 50.00	94.66 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	72
JESMO	SNW	0.00 \pm 0.00	0.00 \pm 0.00	94.66 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	72
PESMO	SNW	0.00 \pm 0.00	0.00 \pm 0.00	94.66 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	72
VOGP (Ours)	BC	100.00 \pm 0.00	100.00 \pm 0.00	99.88 \pm 0.18	85.0 \pm 22.91	100.00 \pm 0.00	17.5 \pm 2.77
MESMO	BC	60.00 \pm 48.99	70.00 \pm 45.83	98.0 \pm 3.01	45.0 \pm 41.53	60.00 \pm 48.99	17
EHVI	BC	100.00 \pm 0.00	100.00 \pm 0.00	99.92 \pm 0.16	90.00 \pm 20.00	100.00 \pm 0.00	17
JESMO	BC	100.00 \pm 0.00	100.00 \pm 0.00	99.96 \pm 0.12	95.0 \pm 15.0	100.00 \pm 0.00	17
PESMO	BC	100.00 \pm 0.00	100.00 \pm 0.00	99.84 \pm 0.20	80.00 \pm 24.49	100.00 \pm 0.00	17
VOGP (Ours)	OKA	65.00 \pm 22.91	85.00 \pm 22.91	99.60 \pm 0.31	65.00 \pm 22.91	85.00 \pm 22.91	44.60 \pm 5.22
MESMO	OKA	50.00 \pm 0.00	85.0 \pm 22.91	99.48 \pm 0.18	50.00 \pm 0.00	85.0 \pm 22.91	44
EHVI	OKA	60.00 \pm 20.00	100.00 \pm 0.00	99.68 \pm 0.16	60.00 \pm 20.00	100.00 \pm 0.00	44
JESMO	OKA	80.00 \pm 24.49	96.67 \pm 10.00	99.80 \pm 0.20	80.00 \pm 24.49	96.67 \pm 10.00	44
PESMO	OKA	70.00 \pm 24.49	96.67 \pm 10.00	99.72 \pm 0.18	70.00 \pm 24.49	96.67 \pm 10.00	44
VOGP (Ours)	SnAr	73.33 \pm 13.33	100.00 \pm 0.00	99.64 \pm 0.1	26.67 \pm 20.00	38.33 \pm 30.78	32.9 \pm 9.41
MESMO	SnAr	30.00 \pm 31.45	60.00 \pm 48.99	99.61 \pm 0.09	6.67 \pm 13.33	20.00 \pm 40.00	32
EHVI	SnAr	0.00 \pm 0.00	0.00 \pm 0.00	99.58 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	32
JESMO	SnAr	0.00 \pm 0.00	0.00 \pm 0.00	99.58 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	32
PESMO	SnAr	0.00 \pm 0.00	0.00 \pm 0.00	99.58 \pm 0.00	0.00 \pm 0.00	0.00 \pm 0.00	32

Experimental Setup In this experiment (by all methods), for OKA dataset, Matérn kernel was used. For SNW, BC, SnAr datasets, RBF kernel was used. Gaussian process kernel parameters are known by all methods. MOO methods are run with the budget of VOGP for $C = C_{135^\circ}$. After MOO methods finish the Bayesian optimization loop, a GP is set with the observations of the MOO methods. Then, the posterior mean values of the GP are used (by identifying the non-dominated posterior means) to calculate the estimated Pareto set for C_{135° . PA, PR, PP are then calculated with respect to C_{135° .

Discussion VOGP consistently ranks among the top-performing methods across datasets and metrics. While VOGP consistently demonstrates high success rates (SR1 and SR2), other methods show a broader range of performance. This result is not unexpected since VOGP performs comparisons with the ordering cone whereas for the other methods, there is a mismatch between their optimized acquisition function and the cone-dependent Pareto front.

B COMPARISON OF VOGP WITH RANDOM SEARCH

Here, we compare the VOGP algorithm with random search as a baseline. The random search algorithm samples as many samples as VOGP had sampled at random. Then, the observed data is used to train a Gaussian process model, whose posterior mean values at every design are used to determine the cone-dependent Pareto set.

Table 5: Performance comparison between VOGP (Our method) and random search under C_{45° , C_{90° , and C_{135° . SR1 and SR2: The success rate of satisfying (ϵ, δ) -PAC conditions (i) and (ii) respectively. PA: Pareto accuracy rate. PR: Pareto recall rate. PP: Pareto precision rate. SC: The number of evaluations.

C	D	VOGP					SC	Random search					SC
		SR1 \uparrow	SR2 \uparrow	PA \uparrow	PR \uparrow	PP \uparrow		SR1 \uparrow	SR2 \uparrow	PA \uparrow	PR \uparrow	PP \uparrow	
C_{45°	SNW	95.96 \pm 2.78	98.65 \pm 1.46	85.53 \pm 1.35	76.54 \pm 4.93	69.49 \pm 3.25	777.8 \pm 220.65	95.38 \pm 0.94	99.24 \pm 0.94	85.63 \pm 1.52	73.27 \pm 4.07	70.87 \pm 3.01	777
C_{45°	BC	95.17 \pm 7.43	99.37 \pm 1.27	95.36 \pm 1.59	86.90 \pm 8.13	76.56 \pm 6.35	446.40 \pm 131.14	99.66 \pm 1.03	100.00 \pm 0.00	97.24 \pm 0.9	91.03 \pm 4.14	86.17 \pm 4.71	446
C_{45°	OKA	88.00 \pm 7.38	92.84 \pm 5.41	95.28 \pm 0.99	76.40 \pm 4.54	76.87 \pm 6.89	650.3 \pm 153.36	97.6 \pm 4.80	94.82 \pm 5.31	95.6 \pm 0.88	79.60 \pm 6.56	77.63 \pm 5.98	650
C_{45°	SnAr	78.75 \pm 14.03	98.00 \pm 6.00	98.46 \pm 0.42	55.62 \pm 16.64	53.63 \pm 12.25	517.0 \pm 148.79	87.50 \pm 7.40	96.89 \pm 5.63	98.53 \pm 0.33	53.12 \pm 11.61	57.09 \pm 12.11	517
C_{90°	SNW	94.62 \pm 4.28	97.5 \pm 2.22	89.85 \pm 2.14	63.85 \pm 10.35	59.61 \pm 8.75	112.70 \pm 28.49	85.38 \pm 7.65	93.94 \pm 3.29	87.72 \pm 3.32	50.00 \pm 13.76	51.75 \pm 13.39	112
C_{90°	BC	98.33 \pm 5.00	100.0 \pm 0.00	99.20 \pm 0.59	80.00 \pm 20.82	88.33 \pm 14.53	33.60 \pm 2.73	98.33 \pm 5.00	88.09 \pm 11.26	98.00 \pm 0.91	88.33 \pm 13.02	59.57 \pm 17.02	33
C_{90°	OKA	87.14 \pm 4.29	99.09 \pm 2.73	96.24 \pm 0.57	57.14 \pm 6.39	38.91 \pm 5.84	110.60 \pm 16.38	72.86 \pm 10.00	87.37 \pm 7.85	95.52 \pm 0.59	37.14 \pm 9.48	27.96 \pm 6.62	110
C_{90°	SnAr	67.5 \pm 40.39	95.00 \pm 10.00	99.46 \pm 0.14	40.00 \pm 25.50	34.42 \pm 21.34	39.80 \pm 20.26	62.50 \pm 16.77	84.58 \pm 19.37	99.45 \pm 0.28	20.00 \pm 18.71	38.33 \pm 37.31	39
C_{90°	MAR	100.00 \pm 0.00	99.67 \pm 0.81	97.70 \pm 0.31	79.32 \pm 2.96	66.84 \pm 4.29	716.43 \pm 173.89	98.16 \pm 2.37	100.00 \pm 0.00	97.96 \pm 0.53	73.42 \pm 6.17	74.33 \pm 9.78	716
C_{135°	SNW	92.00 \pm 9.80	94.39 \pm 6.41	96.12 \pm 1.76	72.00 \pm 12.49	60.07 \pm 15.58	72.80 \pm 11.91	76.00 \pm 16.85	91.79 \pm 9.27	95.10 \pm 2.12	57.00 \pm 21.00	50.50 \pm 17.86	72
C_{135°	BC	100.00 \pm 0.00	100.00 \pm 0.00	99.88 \pm 0.18	85.00 \pm 22.91	100.0 \pm 0	17.50 \pm 2.77	75.00 \pm 40.31	84.68 \pm 25.51	98.84 \pm 1.49	70.00 \pm 40.00	55.77 \pm 40.57	17
C_{135°	OKA	65.00 \pm 22.91	85.00 \pm 22.91	99.60 \pm 0.31	65.00 \pm 22.91	85.00 \pm 22.91	44.60 \pm 5.22	30.00 \pm 24.49	33.33 \pm 37.08	98.76 \pm 0.49	30.00 \pm 24.49	33.33 \pm 37.08	44
C_{135°	SnAr	73.33 \pm 13.33	100.00 \pm 0.00	99.64 \pm 0.10	26.67 \pm 20.00	38.33 \pm 30.78	32.90 \pm 9.41	63.33 \pm 23.33	66.67 \pm 37.27	99.54 \pm 0.14	13.33 \pm 16.33	20.00 \pm 31.45	32

Experimental Setup The setup of this experiment is the same as the experiment setup in Section 5.

Discussion As seen in Table 5, VOGP consistently outperforms random search method across datasets and metrics for cones C_{90° and C_{135° . The superiority of VOGP stems from its efficient sampling as a result of variance reduction sampling. By learning as much as possible with few samples, GP used in VOGP learns more about the underlying objective function than random search. For C_{45° , random search method performs marginally better than VOGP. This could be because VOGP samples more designs than are necessary to learn the objective function globally, possibly because of looseness in confidence intervals of designs. This makes the advantages of variance reduction of VOGP diminish, making random search learn the function well enough to become competitive with VOGP.

C ABLATION STUDY: UNKNOWN KERNEL HYPERPARAMETERS

In numerous practical situations, the correct hyperparameters are often inaccessible to us, nor do we possess reliable initial estimates for them. Therefore, a common approach in practical settings is to gradually learn and adjust the hyperparameters as more queries get included in observations. In this section, we investigate the performance of VOGP when the kernel hyperparameters are unknown. Then, as more observations are made, at each round, the kernel hyperparameters are retrained with the observed data so far. We also share the results of an experiment where VOGP’s parameters are fixed after being trained with a small sample of 30 datapoints. In the latter setting, parameters are not updated as more queries get included in observations.

Table 6: Performance analysis of VOGP (Our method) with known and unknown kernel hyperparameters under C_{45° , C_{90° , and C_{135° . SR1 and SR2: The success rate of satisfying (ϵ, δ) -PAC conditions (i) and (ii) respectively. PA: Pareto accuracy rate. PR: Pareto recall rate. PP: Pareto precision rate. SC: The number of evaluations.

C	D	VOGP (Known Parameters)					SC	VOGP (Unknown Parameters)					SC
		SR1 \uparrow	SR2 \uparrow	PA \uparrow	PR \uparrow	PP \uparrow		SR1 \uparrow	SR2 \uparrow	PA \uparrow	PR \uparrow	PP \uparrow	
C_{45°	SNW	95.96 \pm 2.78	98.65 \pm 1.46	85.53 \pm 1.35	76.54 \pm 4.93	69.49 \pm 3.25	777.80 \pm 220.65	98.46 \pm 2.24	86.93 \pm 17.95	72.33 \pm 23.71	83.46 \pm 10.29	57.74 \pm 17.08	283.2 \pm 206.97
C_{45°	BC	95.17 \pm 7.43	99.37 \pm 1.27	95.36 \pm 1.59	86.90 \pm 8.13	76.56 \pm 6.35	446.40 \pm 131.14	70.34 \pm 35.21	81.1 \pm 31.33	90.84 \pm 3.62	57.24 \pm 31.35	52.46 \pm 25.14	180.8 \pm 102.00
C_{45°	OKA	88.00 \pm 7.38	92.84 \pm 5.41	95.28 \pm 0.99	76.40 \pm 4.54	76.87 \pm 6.89	650.3 \pm 153.36	64.50 \pm 27.82	61.42 \pm 29.96	80.9 \pm 26.83	59.0 \pm 27.69	42.83 \pm 22.63	254.25 \pm 179.59
C_{45°	SnAr	78.75 \pm 14.03	98.00 \pm 6.00	98.46 \pm 0.42	55.62 \pm 16.64	53.63 \pm 12.25	517.0 \pm 148.79	73.75 \pm 29.82	66.54 \pm 34.44	78.96 \pm 38.64	55.62 \pm 32.05	36.38 \pm 22.1	93.9 \pm 64.71
C_{90°	SNW	94.62 \pm 4.28	97.5 \pm 2.22	89.85 \pm 2.14	63.85 \pm 10.35	59.61 \pm 8.75	112.70 \pm 28.49	73.85 \pm 37.0	72.51 \pm 36.46	86.99 \pm 2.26	53.85 \pm 27.41	39.53 \pm 20.75	104.4 \pm 86.2
C_{90°	BC	98.33 \pm 5.00	100.0 \pm 0.00	99.20 \pm 0.59	80.00 \pm 20.82	88.33 \pm 14.53	33.60 \pm 2.73	30.00 \pm 38.59	73.41 \pm 31.02	87.24 \pm 28.3	28.33 \pm 38.04	15.57 \pm 21.56	15.0 \pm 12.75
C_{90°	OKA	87.14 \pm 4.29	99.09 \pm 2.73	96.24 \pm 0.57	57.14 \pm 6.39	38.91 \pm 5.84	110.60 \pm 16.38	69.64 \pm 29.83	50.17 \pm 39.29	73.4 \pm 38.21	46.43 \pm 34.07	24.17 \pm 24.81	29.88 \pm 21.07
C_{90°	SnAr	67.5 \pm 40.39	95.00 \pm 10.00	99.46 \pm 0.14	40.00 \pm 25.50	34.42 \pm 21.34	39.80 \pm 20.26	17.5 \pm 35.44	33.57 \pm 39.75	99.24 \pm 0.21	10.00 \pm 20.00	6.86 \pm 13.95	19.7 \pm 11.47
C_{135°	SNW	92.00 \pm 9.80	94.39 \pm 6.41	96.12 \pm 1.76	72.00 \pm 12.49	60.07 \pm 15.58	72.80 \pm 11.91	64.00 \pm 34.99	70.28 \pm 37.72	94.76 \pm 0.84	56.00 \pm 31.37	38.17 \pm 19.84	67.6 \pm 42.26
C_{135°	BC	100.00 \pm 0.00	100.00 \pm 0.00	99.88 \pm 0.18	85.00 \pm 22.91	100.00 \pm 0	17.50 \pm 2.77	60.00 \pm 48.99	51.91 \pm 44.73	88.8 \pm 29.37	60.0 \pm 48.99	35.08 \pm 36.78	13.6 \pm 8.24
C_{135°	OKA	65.00 \pm 22.91	85.00 \pm 22.91	99.60 \pm 0.31	65.00 \pm 22.91	85.00 \pm 22.91	44.60 \pm 5.22	18.75 \pm 34.8	6.35 \pm 16.5	86.20 \pm 32.28	18.75 \pm 34.8	6.35 \pm 16.5	13.75 \pm 12.4
C_{135°	SnAr	73.33 \pm 13.33	100.00 \pm 0.00	99.64 \pm 0.10	26.67 \pm 20.00	38.33 \pm 30.78	32.90 \pm 9.41	10.00 \pm 21.34	25.00 \pm 40.31	99.58 \pm 0.12	10.00 \pm 21.34	15.00 \pm 32.02	12.00 \pm 5.78

Experimental Setup In this experiment (by all methods), for OKA dataset, Matérn kernel was used. For SNW, BC, SnAr datasets, RBF kernel was used. The instance of VOGP with the unknown kernel hyperparameters, as more observations are made at each round, are retrained with the observed data so far.

Table 7: Performance analysis of VOGP (Our method) with known and partially known hyperparameters under C_{45° , C_{90° , and C_{135° . SR1 and SR2: The success rate of satisfying (ϵ, δ) -PAC conditions (i) and (ii) respectively. PA: Pareto accuracy rate. PR: Pareto recall rate. PP: Pareto precision rate. SC: The number of evaluations.

C	D	VOGP (Known Parameters)					SC ↓	VOGP (Partially Known Parameters)					
		SR1 ↑	SR2 ↑	PA ↑	PR ↑	PP ↑		SR1 ↑	SR2 ↑	PA ↑	PR ↑	PP ↑	SC ↓
C_{45°	SNW	95.96 ± 2.78	98.65 ± 1.46	85.53 ± 1.35	76.54 ± 4.93	69.49 ± 3.25	777.8 ± 220.65	96.35 ± 2.50	97.83 ± 1.73	82.67 ± 2.35	68.08 ± 5.52	65.28 ± 5.18	450.7 ± 117.68
C_{45°	BC	95.17 ± 7.43	99.37 ± 1.27	95.36 ± 1.59	86.90 ± 8.13	76.56 ± 6.35	446.40 ± 131.14	92.07 ± 5.78	96.65 ± 3.50	92.48 ± 2.54	78.28 ± 7.24	65.7 ± 10.44	285.3 ± 93.17
C_{45°	OKA	88.00 ± 7.38	92.84 ± 5.41	95.28 ± 0.99	76.40 ± 4.54	76.87 ± 6.89	650.3 ± 153.36	82.0 ± 7.43	89.7 ± 2.53	93.16 ± 1.31	74.0 ± 7.21	63.78 ± 6.0	365.2 ± 154.56
C_{45°	SnAr	78.75 ± 14.03	98.00 ± 6.00	98.46 ± 0.42	55.62 ± 16.64	53.63 ± 12.25	517.0 ± 148.79	90.0 ± 9.76	96.81 ± 4.31	98.72 ± 0.31	60.0 ± 8.93	62.82 ± 9.74	350.8 ± 91.41
C_{90°	SNW	94.62 ± 4.28	97.5 ± 2.22	89.85 ± 2.14	63.85 ± 10.35	59.61 ± 8.75	112.70 ± 28.49	98.08 ± 3.10	93.23 ± 3.60	89.17 ± 1.55	64.62 ± 11.12	56.36 ± 5.29	69.5 ± 13.29
C_{90°	BC	98.33 ± 5.00	100.0 ± 0.00	99.20 ± 0.59	80.00 ± 20.82	88.33 ± 14.53	33.60 ± 2.73	93.33 ± 20.0	98.89 ± 3.33	98.96 ± 0.93	80.0 ± 19.44	79.89 ± 21.36	28.5 ± 8.02
C_{90°	OKA	87.14 ± 4.29	99.09 ± 2.73	96.24 ± 0.57	57.14 ± 6.39	38.91 ± 5.84	110.60 ± 16.38	75.71 ± 6.55	90.44 ± 10.64	96.28 ± 0.67	40.0 ± 12.45	35.66 ± 9.93	52.5 ± 15.31
C_{90°	SnAr	67.5 ± 40.39	95.00 ± 10.00	99.46 ± 0.14	40.00 ± 25.50	34.42 ± 21.34	39.80 ± 20.26	82.5 ± 11.46	82.92 ± 22.55	99.51 ± 0.16	37.5 ± 16.77	51.0 ± 26.78	27.0 ± 11.08
C_{135°	SNW	92.00 ± 9.80	94.39 ± 6.41	96.12 ± 1.76	72.00 ± 12.49	60.07 ± 15.58	72.80 ± 11.91	95.00 ± 8.06	90.06 ± 10.99	96.65 ± 1.56	85.0 ± 8.06	64.18 ± 15.87	42.0 ± 7.21
C_{135°	BC	100.00 ± 0.00	100.00 ± 0.00	99.88 ± 0.18	85.00 ± 22.91	100.00 ± 0.00	17.50 ± 2.77	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	13.3 ± 1.95
C_{135°	OKA	65.00 ± 22.91	85.00 ± 22.91	99.60 ± 0.31	65.00 ± 22.91	85.00 ± 22.91	44.60 ± 5.22	40.0 ± 20.0	55.0 ± 35.0	99.08 ± 0.57	40.0 ± 20.0	55.0 ± 35.0	25.9 ± 6.86
C_{135°	SnAr	73.33 ± 13.33	100.00 ± 0.00	99.64 ± 0.10	26.67 ± 20.00	38.33 ± 30.78	32.90 ± 9.41	66.67 ± 29.81	75.0 ± 33.54	99.68 ± 0.19	36.67 ± 17.95	60.0 ± 30.0	16.7 ± 4.08

Experimental Setup In this experiment, for OKA dataset, Matérn kernel was used. For SNW, BC, SnAr datasets, RBF kernel was used. The instance of VOGP with the partially known kernel hyperparameters has its parameters with randomly selected 30 datapoints initially. These data are not used to update the GP posterior.

Discussion As seen in Table 6, VOGP’s performance is affected negatively when the kernel parameters are initially unknown. Though the VOGP gets to learn the kernel parameters as more queries are added to the observations, VOGP performs Pareto classification operations in early stages without having the right kernel parameters. Since VOGP does not check discarded designs again and does not reconsider added Pareto designs, these erroneous decisions in the cannot be corrected later on. As seen in Table 7, a small set of 30 initial samples to learn the kernel parameters are enough to make VOGP perform marginally worse than the version with known parameters. Additionally, VOGP samples the objective function less times when it has incomplete information of the kernel hyperparameters. A possible reason for this would be higher length scales, which could lead to observations reducing more uncertainty about other designs, leading to faster terminations.

D RELATED WORKS ON EVOLUTIONARY ALGORITHMS

Numerous studies employ multi-objective evolutionary algorithms in order to estimate the Pareto front, thereby accomplishing this task through the iterative development of a population of evaluated designs; see Seada & Deb (2015); Deb et al. (2002); Knowles (2006); Zhang & Li (2007). Some of these methods use hypervolume calculation to guide their method (Yang et al., 2016b; Beume et al., 2007; Bader & Zitzler, 2011). A significant body of research exists on the integration of user preferences into multi-objective evolutionary algorithms (Coello, 2000; Zhou et al., 2011; Branke & Deb, 2005; Phelps & Köksalan, 2003; Li & Silva, 2008; Sinha et al., 2010). 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.

E PROOFS

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
ζ	The ordering complexity of the cone C
\mathbf{f}	The objective function
ϵ	Accuracy level given as input to the algorithm
$\mathbf{y}_{[t]}$	Vector that represents the first t noisy observations where $\mathbf{y}_{[0]} = \emptyset$.
$\boldsymbol{\mu}_t(x)$	The posterior mean of design x at round t whose j^{th} component is $\mu_t^j(x)$
$\boldsymbol{\sigma}_t(x)$	The posterior variance of design x at round t whose j^{th} component is $\sigma_t^j(x)$
β_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_θ^*	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
$\mathbf{Q}_t(x)$	The confidence hyperrectangle associated with design x at round t
$\mathbf{R}_t(x)$	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
$m(x, x')$	$\inf \{s \geq 0 \mid \exists \mathbf{u} \in \mathbb{B}(1) \cap C : \mathbf{f}(x) + s\mathbf{u} \notin \mathbf{f}(x') - \text{int}(C)\}$
γ_t	The maximum information that can be gained about \mathbf{f} in t evaluations
t_s	The round in which VOGP terminates

E.1 DERIVATION OF THEOREM 1

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

Lemma 2. *Let us define the event*

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

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 Ω is the underlying sample space. Note that

$$\begin{aligned} 1 - \mathbb{P}(E) &= \mathbb{E} \left[\mathbb{I} \left(\left\{ \exists j \in [M] \exists t \in \mathbb{N} \exists x \in \mathcal{X} : |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \right) \right] \\ &\leq \mathbb{E} \left[\sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} \mathbb{I} \left(\left\{ |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \right) \right] \\ &= \sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} \mathbb{E} \left[\mathbb{E} \left[\mathbb{I} \left(\left\{ |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \right) \mid \mathbf{y}_{[t-1]} \right] \right] \end{aligned} \quad (4)$$

$$\begin{aligned} &= \sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} \mathbb{E} \left[\mathbb{P} \left(\left\{ |f^j(x) - \mu_t^j(x)| > \beta_t^{1/2} \sigma_t^j(x) \right\} \mid \mathbf{y}_{[t-1]} \right) \right] \\ &\leq \sum_{j=1}^M \sum_{t=1}^{\infty} \sum_{x \in \mathcal{X}} 2e^{-\beta_t/2} \\ &= 2M|\mathcal{X}| \sum_{t=1}^{\infty} e^{-\beta_t/2} \\ &= 2M|\mathcal{X}| \sum_{t=1}^{\infty} \left(\frac{M\pi^2 |\mathcal{X}| t^2}{3\delta} \right)^{-1} \\ &= \frac{6\delta}{\pi^2} \sum_{t=1}^{\infty} \frac{1}{t^2} = \delta, \end{aligned} \quad (5)$$

where (4) uses the tower rule and linearity of expectation and (5) 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

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

Lemma 3. *Let $s > 0$. Then, $A(s) = sA(1) := \{s\mathbf{y} \mid \mathbf{y} \in A(1)\}$ and $d(s) = sd(1)$. Moreover, there exists a unique $\mathbf{z}^s \in A(s)$ such that $d(s) = \|\mathbf{z}^s\|_2$.*

Proof. Let $s > 0$. Then, note that

$$\begin{aligned} \mathbf{y} \in A(s) &\Leftrightarrow \mathbf{y} \in s\mathbf{u} + C, \forall \mathbf{u} \in \mathbb{S}^{M-1} \\ &\Leftrightarrow \mathbf{y} - s\mathbf{u} \in C, \forall \mathbf{u} \in \mathbb{S}^{M-1} \\ &\Leftrightarrow \mathbf{w}_n^\top (\mathbf{y} - s\mathbf{u}) \geq 0, \forall n \in [N], \forall \mathbf{u} \in \mathbb{S}^{M-1} \\ &\Leftrightarrow \mathbf{w}_n^\top \mathbf{y} \geq s\mathbf{w}_n^\top \mathbf{u}, \forall n \in [N], \forall \mathbf{u} \in \mathbb{S}^{M-1} \\ &\Leftrightarrow \mathbf{w}_n^\top \mathbf{y} \geq s \sup_{\mathbf{u} \in \mathbb{S}^{M-1}} (\mathbf{w}_n^\top \mathbf{u}), \forall n \in [N]. \end{aligned} \quad (6)$$

Let $n \in [N]$. By the definition of \mathbf{w}_n , we have

$$\sup_{\mathbf{u} \in \mathbb{S}^{M-1}} \mathbf{w}_n^\top \mathbf{u} = \sup_{\mathbf{u} \in \mathbb{B}(1)} \mathbf{w}_n^\top \mathbf{u} = \|\mathbf{w}_n\|_2 = 1. \quad (7)$$

Combining (6) and (7), we get

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

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

The existence and uniqueness of \mathbf{z}^s is a direct consequence of the strict convexity and continuity of the ℓ_2 -norm $\|\cdot\|_2$ together with the closedness and convexity of $A(s)$. \square

Remark 3. The last part of Lemma 3 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: For every $\mathbf{y}, \mathbf{z} \in \mathbb{R}^M$ and $\tilde{\mathbf{p}} \in \mathbb{B}(\frac{\epsilon}{d(1)})$, having $\mathbf{y} + \tilde{\mathbf{p}} \preceq_C \mathbf{z}$ implies $\mathbf{y} \preceq_C \mathbf{z} + \epsilon \mathbf{u}^*$.

Proof of Lemma 1. Since $\tilde{\mathbf{p}} \in \mathbb{B}(\frac{\epsilon}{d(1)})$, we have

$$(\tilde{\mathbf{p}} + C) \cap \mathbb{S}^{M-1}\left(\frac{\epsilon}{d(1)}\right) \neq \emptyset \iff \exists \tilde{\gamma} \in \mathbb{S}^{M-1}\left(\frac{\epsilon}{d(1)}\right) : \tilde{\mathbf{p}} \preceq_C \tilde{\gamma}. \quad (8)$$

By Definition 4, we have

$$\begin{aligned} \mathbf{u}^* &= \frac{\operatorname{argmin}_{\mathbf{z} \in A(1)} (\|\mathbf{z}\|_2)}{d(1)} \implies d(1)\mathbf{u}^* \in A(1) \\ &\implies \mathbf{u}^* \in \frac{A(1)}{d(1)} \end{aligned} \quad (9)$$

$$\implies \mathbf{u}^* \in A\left(\frac{1}{d(1)}\right) \quad (10)$$

$$\implies \epsilon \mathbf{u}^* \in A\left(\frac{\epsilon}{d(1)}\right), \quad (11)$$

where (9), (10), and (11) follow from Lemma 3. The division in (10) denotes the scaling of the elements of $A(1)$. The proof is completed by observing that, by Definition 4,

$$\forall \gamma \in \mathbb{S}^{M-1}\left(\frac{\epsilon}{d(1)}\right), \forall \mathbf{k} \in A\left(\frac{\epsilon}{d(1)}\right) : \gamma \preceq_C \mathbf{k} \quad (12)$$

$$\implies \tilde{\gamma} \preceq_C \epsilon \mathbf{u}^* \quad (13)$$

$$\implies \tilde{\mathbf{p}} \preceq_C \epsilon \mathbf{u}^*, \quad (14)$$

where (13) follows from $\tilde{\gamma} \in \mathbb{S}^{M-1}(\frac{\epsilon}{d(1)})$ and (11), and (14) follows from combining (13) with (8). \square

Lemma 4. Under event E , the set $\hat{\mathcal{P}}$ returned by VOGP satisfies condition (i) in Definition 1.

Proof. We 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 $\epsilon \mathbf{u}^* \in C$. If $x \notin \hat{\mathcal{P}}$, then x must have been discarded at some round s_1 . By the discarding rule, that means there exists $z_1 \in \mathcal{P}_{\text{pess}, s_1}$ such that

$$\mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C. \quad (15)$$

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

At each round, the initial confidence hyperrectangle of a design x is calculated as

$$\mathbf{Q}_t(x) = \left\{ \mathbf{y} \in \mathbb{R}^M \mid \boldsymbol{\mu}_t(x) - \beta_t^{1/2} \boldsymbol{\sigma}_t(x) \preceq_{\mathbb{R}_+^M} \mathbf{y} \preceq_{\mathbb{R}_+^M} \boldsymbol{\mu}_t(x) + \beta_t^{1/2} \boldsymbol{\sigma}_t(x) \right\} \quad (16)$$

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

$$\mathbf{R}_t(x) = \mathbf{R}_{t-1}(x) \cap \mathbf{Q}_t(x). \quad (17)$$

It can be checked that, due to Lemma 2, we have $\mathbf{f}(x) \in \mathbf{R}_{s_1}(x)$ under event E . Note that, by (15), there exists $z_1 \in \mathcal{P}_{\text{pess},s_1}$ such that

$$\mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{f}(x) + C. \quad (18)$$

Since $\mathbf{f}(z_1) \in \mathbf{R}_{s_1}(x)$, (18) implies that $\mathbf{f}(z_1) + \epsilon \mathbf{u}^* \in \mathbf{f}(x) + C$, which is equivalent to $\mathbf{f}(x) \preceq_C \mathbf{f}(z_1) + \epsilon \mathbf{u}^*$. 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 $S_t \setminus \mathcal{P}_{\text{pess},t}$ at round t , $z_1 \notin \mathcal{P}_{\text{pess},s_2}$. Then, using the definition of the pessimistic set (see Definition 3), there exists $z_2 \in \mathcal{A}_{s_2}$ such that

$$\mathbf{R}_{s_2}(z_2) \subseteq \mathbf{R}_{s_2}(z_1) + C. \quad (19)$$

To proceed, we use (18) and the fact that $C + C = C$, to obtain

$$\begin{aligned} \mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{f}(x) + C &\implies \mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* + C \subseteq \mathbf{f}(x) + C + C \\ &\implies \mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* + C \subseteq \mathbf{f}(x) + C. \end{aligned} \quad (20)$$

In addition, (19) implies the following

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

Combining the above display with (20) yields

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

According to Lemma 2, under event E , $\mathbf{f}(z_2) \in \mathbf{R}_{s_2}(z_2)$. Hence, it holds that

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

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 $S_t \setminus \mathcal{P}_{\text{pess},t}$ at t , $z_2 \notin \mathcal{P}_{\text{pess},s_3}$. Then, using the definition of the pessimistic set, there exists $z_3 \in \mathcal{A}_{s_3}$ such that

$$\mathbf{R}_{s_3}(z_3) \subseteq \mathbf{R}_{s_3}(z_2) + C \implies \mathbf{R}_{s_3}(z_3) \subseteq \mathbf{R}_{s_2}(z_2) + C, \quad (21)$$

where (21) follows from the fact that $\mathbf{R}_{s_3}(z_2) \subseteq \mathbf{R}_{s_2}(z_2)$. Continuing from (19), we have

$$\mathbf{R}_{s_2}(z_2) \subseteq \mathbf{R}_{s_2}(z_1) + C \implies \mathbf{R}_{s_2}(z_2) + C \subseteq \mathbf{R}_{s_2}(z_1) + C + C \quad (22)$$

$$\implies \mathbf{R}_{s_2}(z_2) + C \subseteq \mathbf{R}_{s_2}(z_1) + C, \quad (23)$$

where (22) follows from the definition of Minkowski sum and (23) follows from the convexity property of C .

Combining (21) and (23), we have

$$\mathbf{R}_{s_3}(z_3) \subseteq \mathbf{R}_{s_2}(z_1) + C. \quad (24)$$

Next, using (24), $s_2 \geq s_1$, and (20), we get

$$\mathbf{R}_{s_3}(z_3) + \epsilon \mathbf{u}^* \subseteq \mathbf{R}_{s_2}(z_1) + \epsilon \mathbf{u}^* + C \subseteq \mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* + C \subseteq \mathbf{f}(x) + C.$$

According to Lemma 2, $\mathbf{f}(z_3) \in \mathbf{R}_{s_3}(z_3)$ under event E . Hence, 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 $z_3 \notin \hat{\mathcal{P}}$, a similar argument can be made until $z_n \in \hat{\mathcal{P}}$. In the worst case, there comes a point where $z_n \in A_{t_s}$, in which case it is either discarded or added to the $\hat{\mathcal{P}}$. If it is discarded, then it is removed from $\mathcal{P}_{\text{pess},t_s}$ by some design which is then moved to $\hat{\mathcal{P}}$. If it is not discarded, then $z_n \in \hat{\mathcal{P}}$. \square

Lemma 5. *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 ϵ -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

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

By the properties of the Minkowski sum, we have $\mathbf{R}_t(x) + \epsilon\mathbf{u}^* \subset \mathbf{R}_t(x) + \epsilon\mathbf{u}^* + C$. This together with (25) results in

$$(\mathbf{R}_t(x) + \epsilon\mathbf{u}^*) \cap (\mathbf{R}_t(x') - C) = \emptyset. \quad (26)$$

According to Lemma 2, 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 (26), we conclude that

$$\mathbf{f}(x) + \epsilon\mathbf{u}^* \notin \mathbf{f}(x') - C. \quad (27)$$

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

$$\mathbf{R}_{s_1}(z_1) + \epsilon\mathbf{u}^* \subseteq \mathbf{y} + C, \forall \mathbf{y} \in \mathbf{R}_{s_1}(x'). \quad (28)$$

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 2.1:** $z_1 \notin \mathcal{W}_t$.

Case 2.1: If $z_1 \in \mathcal{W}_t$, since $\mathbf{R}_t(z_1) \subseteq \mathbf{R}_{s_1}(z_1)$, (28) also implies that

$$\begin{aligned} & \mathbf{R}_t(z_1) + \epsilon\mathbf{u}^* \subseteq \mathbf{y} + C, \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') \\ \iff & \forall \mathbf{y}_{x', s_1} \in \mathbf{R}_{s_1}(x'), \forall \mathbf{y}_{z_1, t} \in \mathbf{R}_t(z_1) : \mathbf{y}_{z_1, t} \in \mathbf{y}_{x', s_1} - \epsilon\mathbf{u}^* + C \\ \iff & \forall \mathbf{y}_{x', s_1} \in \mathbf{R}_{s_1}(x'), \forall \mathbf{y}_{z_1, t} \in \mathbf{R}_t(z_1) : \mathbf{y}_{x', s_1} - \epsilon\mathbf{u}^* \preceq_C \mathbf{y}_{z_1, t}. \end{aligned} \quad (29)$$

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

$$(\mathbf{R}_t(x) + \epsilon\mathbf{u}^* + C) \cap (\mathbf{R}_t(z_1) - C) = \emptyset \quad (30)$$

$$\begin{aligned} \implies & (\mathbf{R}_t(x) + \epsilon\mathbf{u}^*) \cap (\mathbf{R}_t(z_1) - C) = \emptyset \\ \iff & \forall \mathbf{y}_{z_1, t} \in \mathbf{R}_t(z_1), \forall \mathbf{y}_{x, t} \in \mathbf{R}_t(x) : \mathbf{y}_{x, t} + \epsilon\mathbf{u}^* \not\preceq_C \mathbf{y}_{z_1, t}. \end{aligned} \quad (31)$$

Then, by combining (29) and (31), we get $\mathbf{y}_{x, t} + \epsilon\mathbf{u}^* \not\preceq_C \mathbf{y}_{x', s_1} - \epsilon\mathbf{u}^*$, $\forall \mathbf{y}_{x, t} \in \mathbf{R}_t(x)$ and $\forall \mathbf{y}_{x', s_1} \in \mathbf{R}_{s_1}(x')$. Therefore, according to Lemma 2, under the good event \mathcal{F}_1 it holds that $\mathbf{f}(x) + 2\epsilon\mathbf{u}^* \not\preceq_C \mathbf{f}(x')$. 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 $\mathbf{R}_{s_{i+1}}(z_{i+1}) \subseteq \mathbf{R}_{s_{i+1}}(z_i) + C$. 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. Due to the definition of the pessimistic Pareto set $\mathcal{P}_{\text{pess}}$ given in Definition 3 and (17), we observe that the following set operations hold.

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

In particular, (32) 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

$$\begin{aligned} \mathbf{R}_{s_2}(z_1) + C &\supseteq \mathbf{R}_{s_3}(z_2) + C \\ \mathbf{R}_{s_3}(z_2) + C &\supseteq \mathbf{R}_{s_4}(z_3) + C \\ &\vdots \\ \mathbf{R}_{s_{n-1}}(z_{n-2}) + C &\supseteq \mathbf{R}_{s_n}(z_{n-1}) + C \\ \implies \mathbf{R}_{s_2}(z_1) + C &\supseteq \mathbf{R}_{s_n}(z_{n-1}) + C. \end{aligned} \quad (33)$$

Using the fact that $\mathbf{R}_{s_1}(z_1) \supseteq \mathbf{R}_{s_2}(z_1)$ when $s_1 \leq s_2$, (33), and (32) in order, we obtain

$$\mathbf{R}_{s_1}(z_1) + C \supseteq \mathbf{R}_{s_2}(z_1) + C \supseteq \mathbf{R}_{s_n}(z_{n-1}) + C \supseteq \mathbf{R}_t(z_n). \quad (34)$$

Next, by combining (28) and (34), and using the properties of Minkowski sum, we have

$$\begin{aligned} \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_{s_1}(z_1) + \epsilon \mathbf{u}^* \subseteq \mathbf{y} + C &\iff \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_{s_1}(z_1) \subseteq \mathbf{y} - \epsilon \mathbf{u}^* + C \\ &\implies \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_{s_1}(z_1) + C \subseteq \mathbf{y} - \epsilon \mathbf{u}^* + C \\ &\implies \forall \mathbf{y} \in \mathbf{R}_{s_1}(x') : \mathbf{R}_t(z_n) \subseteq \mathbf{y} - \epsilon \mathbf{u}^* + C. \end{aligned} \quad (35)$$

Alternatively, (35) can be re-written as

$$\forall \mathbf{y}_{x',s_1} \in \mathbf{R}_{s_1}(x'), \forall \mathbf{y}_{z_n,t} \in \mathbf{R}_t(z_n) : \mathbf{y}_{x',s_1} - \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z_n,t}. \quad (36)$$

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

$$(\mathbf{R}_t(x) + \epsilon \mathbf{u}^* + C) \cap (\mathbf{R}_t(z_n) - C) = \emptyset \quad (37)$$

$$\begin{aligned} &\implies (\mathbf{R}_t(x) + \epsilon \mathbf{u}^*) \cap (\mathbf{R}_t(z_n) - C) = \emptyset \\ &\implies \forall \mathbf{y}_{z_n,t} \in \mathbf{R}_t(z_n), \forall \mathbf{y}_{x,t} \in \mathbf{R}_t(x) : \mathbf{y}_{x,t} + \epsilon \mathbf{u}^* \not\preceq_C \mathbf{y}_{z_n,t}. \end{aligned} \quad (38)$$

Then, by combining (36) and (38), we get $\mathbf{y}_{x,t} + \epsilon \mathbf{u}^* \not\preceq_C \mathbf{y}_{x',s_1} - \epsilon \mathbf{u}^*$, $\forall \mathbf{y}_{x,t} \in \mathbf{R}_t(x)$ and $\forall \mathbf{y}_{x',s_1} \in \mathbf{R}_{s_1}(x')$. Then, according to Lemma 2, under the good event E it holds that $\mathbf{f}(x) + 2\epsilon \mathbf{u}^* \not\preceq_C \mathbf{f}(x')$. 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

Lemma 6. *If $\bar{\omega}_t = \max_{x \in \mathcal{W}_t} \omega_t(x) \leq \frac{\epsilon}{2d(1)}$, then it holds that $\forall x \in \mathcal{W}_t, \forall \mathbf{y}_x, \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) : \mathbf{y}_x \preceq_C \mathbf{y}_{\bar{x}} + \epsilon \mathbf{u}^*$.*

Proof. We have

$$(\mathbf{y}_x - \mathbf{y}_{\bar{x}}) \in \mathbb{B}\left(\frac{2\epsilon}{2d(1)}\right) = \mathbb{B}\left(\frac{\epsilon}{d(1)}\right).$$

Therefore, it holds that

$$(\mathbf{y}_x - \mathbf{y}_{\bar{x}} + C) \cap \mathbb{S}^{M-1} \left(\frac{\epsilon}{d(1)} \right) \neq \emptyset \iff \exists \tilde{\gamma} \in \mathbb{S}^{M-1} \left(\frac{\epsilon}{d(1)} \right) : \mathbf{y}_x - \mathbf{y}_{\bar{x}} \preceq_C \tilde{\gamma}. \quad (39)$$

By Definition 4, we have

$$\begin{aligned} \mathbf{u}^* &= \frac{\operatorname{argmin}_{z \in A(1)} (\|z\|_2)}{d(1)} \implies d(1)\mathbf{u}^* \in A(1) \\ &\implies \mathbf{u}^* \in \frac{A(1)}{d(1)} \end{aligned} \quad (40)$$

$$\implies \mathbf{u}^* \in A \left(\frac{1}{d(1)} \right) \quad (41)$$

$$\implies \epsilon \mathbf{u}^* \in A \left(\frac{\epsilon}{d(1)} \right), \quad (42)$$

where (40), (41), and (42) follow from Lemma 3. The division in (41) denotes the scaling of the elements of $A(1)$. The proof is completed by observing that, by Definition 4,

$$\forall \gamma \in \mathbb{S}^{M-1} \left(\frac{\epsilon}{d(1)} \right), \forall \mathbf{k} \in A \left(\frac{\epsilon}{d(1)} \right) : \gamma \preceq_C \mathbf{k} \quad (43)$$

$$\implies \tilde{\gamma} \preceq_C \epsilon \mathbf{u}^* \quad (44)$$

$$\implies \mathbf{y}_x - \mathbf{y}_{\bar{x}} \preceq_C \epsilon \mathbf{u}^*, \quad (45)$$

where (44) follows from $\tilde{\gamma} \in \mathbb{S}^{M-1} \left(\frac{\epsilon}{d(1)} \right)$ and (42), and (45) follows from combining (44) with (39). \square

Lemma 7. *VOGP terminates at round t if $\bar{\omega}_t < \frac{\epsilon}{2d(1)}$.*

Proof. Let $\mathcal{S}_{t,0}(P_{t,0})$, $\mathcal{S}_{t,1}(P_{t,1})$ and $\mathcal{S}_{t,2}(P_{t,2})$ denote sets $\mathcal{S}_t(\mathcal{P})$ at the end of modeling, discarding and covering phases, respectively. For an $x \in \mathcal{S}_{t,0} \setminus P_{t,2}$, if $x \notin \mathcal{S}_{t,1}$, that means it must have been discarded at round t . So, if the claim holds, any $x \in \mathcal{S}_{t,0}$ is either discarded or moved to $P_{t,2}$. In order to prove the lemma, we will show that if $x \in \mathcal{S}_{t,0} \setminus P_{t,2}$ holds, then x cannot belong to $\mathcal{S}_{t,1}$. To prove this by contradiction, assume that $x \in \mathcal{S}_{t,1}$ (otherwise it is discarded). Since $x \in \mathcal{S}_{t,0} \setminus P_{t,2}$, x has not been added to $P_{t,2}$. According to the ϵ -covering rule, there exists some $z^* \in P_{t,1} \cup \mathcal{S}_{t,1}$ for which

$$\exists \mathbf{y}_{z^*} \in \mathbf{R}_t(z^*), \exists \mathbf{y}_x \in \mathbf{R}_t(x) : \mathbf{y}_x + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z^*} \quad (46)$$

Fix \mathbf{y}_{z^*} and \mathbf{y}_x as given in (46). Since we assume that $x \in P_{t,1} \subseteq P_{t,1} \cup \mathcal{S}_{t,1} = \mathcal{W}_t$, by Lemma 6, we have

$$\forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) : \mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_x + \epsilon \mathbf{u}^*. \quad (47)$$

Again, by Lemma 6, since $z^* \in \mathcal{W}_t$, we have

$$\forall \tilde{\mathbf{y}}_z \in \mathbf{R}_t(z^*) : \mathbf{y}_{z^*} \preceq_C \tilde{\mathbf{y}}_z + \epsilon \mathbf{u}^*. \quad (48)$$

Then, starting from (46), and using (48), we have

$$\mathbf{y}_x + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z^*} \preceq_C \tilde{\mathbf{y}}_z + \epsilon \mathbf{u}^*, \forall \tilde{\mathbf{y}}_z \in \mathbf{R}_t(z^*) \quad (49)$$

$$\iff \mathbf{y}_x \preceq_C \tilde{\mathbf{y}}_z, \forall \tilde{\mathbf{y}}_z \in \mathbf{R}_t(z^*). \quad (50)$$

Now, starting from (46) and using (47), we have

$$\begin{aligned} \mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_x + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z^*}, \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) \\ \implies \mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_{z^*}, \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x). \end{aligned} \quad (51)$$

Starting from (51) and using (48), we have

$$\begin{aligned} \mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_{z^*} \preceq_C \tilde{\mathbf{y}}_z + \epsilon \mathbf{u}^*, \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) \text{ and } \forall \tilde{\mathbf{y}}_z \in \mathbf{R}_t(z^*) \\ \implies \mathbf{y}_{\bar{x}} \preceq_C \tilde{\mathbf{y}}_z + \epsilon \mathbf{u}^*, \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) \text{ and } \forall \tilde{\mathbf{y}}_z \in \mathbf{R}_t(z^*). \end{aligned} \quad (52)$$

Notice that (50) shows that $x \notin \mathcal{P}_{\text{pess},t}$. If $z^* \in \mathcal{P}_{\text{pess},t}$, then (52) shows that x must be discarded. If $z^* \notin \mathcal{P}_{\text{pess},t}$, then $\exists z' \in \mathcal{A}_t$ such that

$$\forall \mathbf{y}_{z'} \in \mathbf{R}_t(z'), \exists \mathbf{y}_{z''} \in \mathbf{R}_t(z^*) : \mathbf{y}_{z''} \preceq_C \mathbf{y}_{z'} \quad (53)$$

$$\iff \forall \mathbf{y}_{z'} \in \mathbf{R}_t(z'), \exists \mathbf{y}_{z''} \in \mathbf{R}_t(z^*) : \mathbf{y}_{z''} + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z'} + \epsilon \mathbf{u}^*. \quad (54)$$

Fix $\mathbf{y}_{z''}$ as given in (53). Starting from (52) and using (54), we have

$$\mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_{z''} + \epsilon \mathbf{u}^* \preceq_C \mathbf{y}_{z'} + \epsilon \mathbf{u}^*, \forall \mathbf{y}_{z'} \in \mathbf{R}_t(z') \text{ and } \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x) \quad (55)$$

$$\implies \mathbf{y}_{\bar{x}} \preceq_C \mathbf{y}_{z'} + \epsilon \mathbf{u}^*, \forall \mathbf{y}_{z'} \in \mathbf{R}_t(z') \text{ and } \forall \mathbf{y}_{\bar{x}} \in \mathbf{R}_t(x). \quad (56)$$

(56) 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{S}_t is a finite set. Hence, the lemma is proved. \square

Lemma 8. *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 γ_{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

$$\sum_{t=1}^{t_s} \bar{\omega}_t^2 = \sum_{t=1}^{t_s} \max_{\mathbf{y}, \mathbf{y}' \in \mathbf{R}_t(x_t)} \|\mathbf{y} - \mathbf{y}'\|_2^2 \quad (57)$$

$$\leq \sum_{t=1}^{t_s} \max_{\mathbf{y}, \mathbf{y}' \in \mathbf{R}_{t-1}(x_t)} \|\mathbf{y} - \mathbf{y}'\|_2^2$$

$$\leq \sum_{t=1}^{t_s} \max_{\mathbf{y}, \mathbf{y}' \in \mathbf{Q}_{t-1}(x_t)} \|\mathbf{y} - \mathbf{y}'\|_2^2$$

$$= \sum_{t=1}^{t_s} \sum_{j=1}^M \left(2\beta_{t-1}^{1/2} \sigma_{t-1}^j(x_t) \right)^2 \quad (58)$$

$$\leq 4\beta_{t_s} \sum_{t=1}^{t_s} \sum_{j=1}^M (\sigma_{t-1}^j(x_t))^2 \quad (59)$$

$$= 4\beta_{t_s} \sigma^2 \sum_{t=1}^{t_s} \sum_{j=1}^M \sigma^{-2} (\sigma_{t-1}^j(x_t))^2$$

$$\leq 4\beta_{t_s} \sigma^2 \eta \left(\sum_{t=1}^{t_s} \sum_{j=1}^M \ln(1 + \sigma^{-2} (\sigma_{t-1}^j(x_t))^2) \right) \quad (60)$$

$$\leq 8\beta_{t_s} \sigma^2 \eta MI(\mathbf{y}_{[t_s]}; f_{[t_s]}) \quad (61)$$

$$\leq 8\beta_{t_s} \sigma^2 \eta M \gamma_{t_s}, \quad (62)$$

where $\eta := \sigma^{-2} / \ln(1 + \sigma^{-2})$ and σ is the noise standard deviation; (57) is due to the definition of $\bar{\omega}_t$; (58) follows from (16); (59) holds since β_t is nondecreasing in t ; (60) follows from the fact that

$s \leq \eta \ln(1 + s)$ for all $0 \leq s \leq \sigma^{-2}$ and that we have

$$\begin{aligned}
& \sigma^{-2} \left(\sigma_{t-1}^j(x_t) \right)^2 \\
&= \sigma^{-2} \left(k^{jj}(x_t, x_t) - \left(\mathbf{k}_{[t-1]}(x_t) \left(\mathbf{K}_{[t-1]} + \boldsymbol{\Sigma}_{[t-1]} \right)^{-1} \mathbf{k}_{[t-1]}(x_t)^\top \right)^{jj} \right) \\
&\leq \sigma^{-2} k^{jj}(x_t, x_t) \\
&\leq \sigma^{-2},
\end{aligned} \tag{63}$$

where (63) follows from the fact that $k^{jj}(x, x') \leq 1$ for all $x, x' \in \mathcal{X}$, and (61) follows from (Nika et al., 2021, Proposition 1). Finally, by Cauchy-Schwartz inequality, we have

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

which completes the proof. \square

E.1.1 PROOF OF THEOREM 1

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 8, we have

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

E.2 DERIVATION OF THEOREM 2

E.2.1 AUXILIARY LEMMATA

Lemma 9. (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 λ is an eigenvalue of \mathbf{A} with corresponding eigenvector $\mathbf{v} \in \mathbb{R}^n$ and μ 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 μ_1, \dots, μ_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 $\mathbf{B} \otimes \mathbf{A}$.

Lemma 10. (Chatzigeorgiou, 2013b, Theorem 1) The Lambert function W_{-1} satisfies

$$-1 - \sqrt{2u} - u < W_{-1}(-e^{-u-1}) < -1 - \sqrt{2u} - \frac{2}{3}u$$

for every $u > 0$.

Lemma 11. Let $\tau \geq 0$ and $t > e^\tau$. Setting $\omega := \frac{\ln^\tau(t)}{t}$, we have

$$t \leq \exp \left(\tau \sqrt{2 \left(\ln(\tau) - \frac{\ln(\omega)}{\tau} - 1 \right)} \right) \frac{\tau^\tau}{\omega}. \tag{64}$$

Proof. For convenience, let us write $a := \ln(t)$. Hence, $\omega = e^{-a} a^\tau$ and

$$a e^{-\frac{a}{\tau}} = \omega^{\frac{1}{\tau}}.$$

Let $b := \frac{a}{\tau}$. Then,

$$b \tau e^{-b} = \omega^{\frac{1}{\tau}},$$

which implies that

$$-b e^{-b} = -\frac{\omega^{\frac{1}{\tau}}}{\tau}.$$

Thus,

$$-b = W_{-1} \left(-\frac{\omega^{\frac{1}{\tau}}}{\tau} \right).$$

Here, we use the Lambert W function with index -1 thanks to the assumption $t > e^\tau$. Then, we obtain

$$\frac{\ln(t)}{\tau} = -W_{-1} \left(-\frac{\omega^{\frac{1}{\tau}}}{\tau} \right)$$

so that

$$t = e^{\left(-\tau \cdot W_{-1} \left(-\frac{\sqrt[\tau]{\omega}}{\tau} \right)\right)}.$$

To be able to use Theorem 10, we set u via

$$\frac{\sqrt[\tau]{\omega}}{\tau} = e^{-u-1} \Leftrightarrow \frac{\ln(\omega)}{\tau} - \ln(\tau) = -u - 1 \Leftrightarrow u = \ln(\tau) - \frac{\ln(\omega)}{\tau} - 1.$$

First, we prove that $u > 0$. To that end, we show that $t \mapsto w(t) := \frac{\ln^\tau(t)}{t}$ is decreasing for $t > e^\tau$. We have

$$\frac{\partial}{\partial t} \left(\frac{\ln^\tau(t)}{t} \right) = \frac{(\tau - \ln(t)) \ln^{\tau-1}(t)}{t^2}.$$

Since $t > e^\tau$, we have $\ln^{\tau-1}(t) > 0$ and $t^2 > 0$. Also notice that $\tau - \ln(t) < 0$ for $t > e^\tau$. Hence, $t \mapsto w(t)$ is decreasing for $t > e^\tau$. Now, if we take $t = e^\tau$, then we have

$$\begin{aligned} u &= \ln(\tau) - \frac{\ln(w(e^\tau))}{\tau} - 1 \\ &= \ln(\tau) - \frac{\ln\left(\frac{\ln^\tau(e^\tau)}{e^\tau}\right)}{\tau} - 1 \\ &= \ln(\tau) - \frac{\ln(\ln^\tau(e^\tau)) - \ln(e^\tau)}{\tau} - 1 \\ &= \ln(\tau) - \frac{\ln(\tau^\tau) - \tau}{\tau} - 1 \\ &= \ln(\tau) - \ln(\tau) + 1 - 1 \\ &= 0 \end{aligned}$$

Since $w(t) = \frac{\ln^\tau(t)}{t}$ is decreasing for $t > e^\tau$, and $u = 0$ when $t = e^\tau$, we can conclude that $u > 0$ whenever $t > e^\tau$.

Therefore, we may apply the bound in Theorem 10 and obtain

$$\begin{aligned} t &\leq \exp\left((1 + \sqrt{2u} + u)\tau\right) \\ &= \exp\left(\tau + \tau\sqrt{2\left(\ln(\tau) - \frac{\ln(\omega)}{\tau} - 1\right)} + \tau\left(\ln(\tau) - \frac{\ln(\omega)}{\tau} - 1\right)\right) \\ &= \exp\left(\tau\sqrt{2\left(\ln(\tau) - \frac{\ln(\omega)}{\tau} - 1\right)} + \tau\ln(\tau) - \ln(\omega)\right) \\ &= \exp\left(\tau\sqrt{2\left(\ln(\tau) - \frac{\ln(\omega)}{\tau} - 1\right)}\right) \cdot \frac{\tau^\tau}{\omega}, \end{aligned}$$

which completes the proof. \square

Lemma 12. *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 Ψ_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, 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 (Li et al., 2022, Theorem 2). The difference comes from the structures of the covariance matrices, where the order of Kronecker products are swapped in our case. (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

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

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 9 and using the identity $|\mathbf{I} + \mathbf{A} \otimes \mathbf{B}| = |\mathbf{I} + \mathbf{B} \otimes \mathbf{A}|$, we get

$$I(\mathbf{y}_{[t]}; \mathbf{f}_{[t]}) = \frac{1}{2} \ln \left| \mathbf{I}_{Mt} + \frac{1}{\sigma^2} [k^*(p, q)]_{p,q \in [M]} \otimes [\tilde{k}(x_i, x_j)]_{i,j \in [t]} \right|. \quad (65)$$

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

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

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. \square

E.2.2 ANALYSIS FOR AN RBF KERNEL

Now, let us assume that we are using an RBF kernel \tilde{k} for the design space. Then, by Lemma 12 and the bounds on maximum information gain established in Srinivas et al. (2012), we have

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

Notice that in Theorem 1, as ϵ 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 $\beta_t = 2 \ln(M\pi^2 |\mathcal{X}| t^2 / (3\delta))$. Let $\alpha_1 = M\pi^2 |\mathcal{X}| / (3\delta)$, $\alpha_2 = 16\sigma^2 \eta M^2$ and $\alpha_3 = \alpha_2 \phi$ where ϕ is the multiplicative constant that comes from the \mathcal{O} notation. We continue by finding an

upper bound on the left of the inequality given in Theorem 1. Note that, for $t \geq 3$, we have $\ln(t) \geq 1$ so that

$$\begin{aligned}
\sqrt{\frac{8\beta_t\sigma^2\eta M\gamma_t}{t}} &= \sqrt{\frac{\ln(\alpha_1 t^2)\gamma_t\alpha_2}{tM}} \\
&\leq \sqrt{\frac{\ln(\alpha_1 t^2)\alpha_3 \ln^{D+1}(t)}{t}} \\
&= \sqrt{\frac{[\ln(\alpha_1) + \ln(t^2)] \ln^{D+1}(t)}{t}} \cdot \alpha_3 \\
&\leq \sqrt{\frac{[\ln(\alpha_1)\ln(t) + \ln(t^2)] \ln^{D+1}(t)}{t}} \cdot \alpha_3 \\
&= \sqrt{\frac{\ln^{D+2}(t)}{t}} \cdot \alpha_3(2 + \ln(\alpha_1)).
\end{aligned} \tag{66}$$

Let us define

$$\begin{aligned}
t^* &:= \min \left\{ t \in \mathbb{N} \mid t \geq 3, \sqrt{\frac{\ln^{D+2}(t)}{t}} \alpha_3(2 + \ln(\alpha_1)) \leq \frac{\epsilon}{2d(1)} \right\} \\
&= \min \left\{ t \in \mathbb{N} \mid t \geq 3, \frac{\ln^{D+2}(t)}{t} \leq \frac{\epsilon^2}{4d^2(1)(2 + \ln(\alpha_1))\alpha_3} \right\}.
\end{aligned}$$

Notice that $t^* \geq T$. Moreover, the definition of t^* also guarantees that

$$\begin{aligned}
t^* - 1 &\leq \min \left\{ t \in \mathbb{R} \mid t \geq 3, \frac{\ln^{D+2}(t)}{t} \leq \frac{\epsilon^2}{4d^2(1)(2 + \ln(\alpha_1))\alpha_3} \right\} \\
&\leq \min \left\{ t \in \mathbb{R} \mid t \geq e^{D+2}, \frac{\ln^{D+2}(t)}{t} \leq \frac{\epsilon^2}{4d^2(1)(2 + \ln(\alpha_1))\alpha_3} \right\},
\end{aligned}$$

which implies that

$$t^* - 1 \leq \tilde{t}, \tag{67}$$

where $t = \tilde{t}$ is the unique solution of the equation

$$\frac{\ln^{D+2}(t)}{t} = \frac{\epsilon^2}{4d^2(1)(2 + \ln(\alpha_1))\alpha_3}$$

in the region $t \in [e^{D+2}, +\infty)$, thanks to the fact that $t \mapsto \frac{\ln^{D+2}(t)}{t}$ is a decreasing function in this region. Since $T - 1 \leq t^* - 1 \leq \tilde{t}$ and we can safely ignore the -1 term, we can use Lemma 11 where $\tau = D + 2$ and $\omega = \frac{\epsilon^2}{4d^2(1)(2 + \ln(\alpha_1))\alpha_3}$. Hence, the analysis for RBF kernel is complete.

E.2.3 ANALYSIS FOR A MATÉRN KERNEL

Next, we suppose that \tilde{k} is a Matérn kernel. The analysis is very similar to the analysis for an RBF kernel. We use the bounds found by Vakili et al. (2021b):

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

By observing only the exponent of the denominator t changes and $D + 1$ is replaced by $\frac{2\nu}{2\nu+d}$ in comparison to (66), we have (67) with $t = \tilde{t}$ solving the equation

$$\frac{\ln^{4\nu+D}(t)}{t^{(2\nu)/(2\nu+D)}} = \frac{\epsilon^2}{4d^2(1)(2 + \ln(\alpha_1))\alpha_3} \Leftrightarrow \frac{\ln^{\frac{(4\nu+D)(2\nu+D)}{2\nu}}(t)}{t} = \frac{\epsilon^{\frac{2\nu+D}{\nu}}}{(4\alpha_3 d^2(1)(2 + \ln(\alpha_1)))^{\frac{2\nu+D}{2\nu}}} \tag{68}$$

The sample complexity is then found by applying Lemma 11 with $\tau = \frac{(4\nu+D)(2\nu+D)}{2\nu}$ and ω being the last term in (68). Hence, the analysis for a Matérn kernel is complete.