
Finding Interior Optimum of Black-box Constrained Objective with Bayesian Optimization

Fengxue Zhang¹

Yuxin Chen¹

¹Computer Science Dept., University of Chicago, Chicago, Illinois, USA

Abstract

Optimizing objectives under constraints, where both the objectives and constraints are black box functions, is a common challenge in real-world applications such as medical therapy design, industrial process optimization, and hyperparameter optimization. Bayesian Optimization (BO) is a popular approach for tackling these complex scenarios. However, constrained Bayesian Optimization (CBO) often relies on heuristics, approximations, or relaxation of objectives, which can lead to weaker theoretical guarantees compared to canonical BO. In this paper, we address this gap by focusing on identifying the interior optimum of the constrained objective, deliberately excluding boundary candidates susceptible to noise perturbations. Our approach leverages the insight that jointly optimizing the objective and learning the constraints can help pinpoint high-confidence *regions of interest* (ROI) likely to contain the interior optimum. We introduce an efficient CBO framework, which intersects these ROIs within a discretized search space to determine a general ROI. Within this ROI, we optimize the acquisition functions, balancing constraints learning and objective optimization. We showcase the efficiency and robustness of our proposed framework by deriving high-probability regret bounds and validating its performance through extensive empirical evaluations.

1 INTRODUCTION

Bayesian optimization (BO) has been extensively studied over the past few decades as a powerful framework for addressing expensive black-box optimization tasks in machine learning, engineering, and science. In many real-world applications, these optimization tasks often involve black-

box constraints that are costly to evaluate [Digabel and Wild, 2015]. Examples include choosing from a plethora of untested medical therapies under safety constraints [Sui et al., 2015]; determining optimal pumping rates in hydrology to minimize operational costs under constraints on plume boundaries [Gramacy et al., 2016]; or tuning hyperparameters of a neural network under memory constraints [Gelbart et al., 2014]. It is common to model constraints analogously to the objectives via Gaussian processes (GP) and then utilize an acquisition function to trade off the learning and optimization to decide subsequent query points.

Recently, significant advancements have been made in several directions to address constrained BO (CBO). For instance, extended Expected Improvement approaches [Bernardo et al., 2011, Gelbart et al., 2014, Gardner et al., 2014, Zhang et al., 2021, Bachoc et al., 2020] learn the constraints passively and calibrate the acquisition with feasibility. The augmented lagrangian (AL) methods [Gramacy et al., 2016, Picheny et al., 2016, Ariafar et al., 2019] convert constrained optimization into unconstrained optimization with additional hyperparameters. The entropy-based methods [Takeno et al., 2022] optimize the lower bound of the mutual information concerning the underlying optimum within the feasible region.

In general, the existing methods extend the unconstrained BO methods to learn the constraints and optimize the objective simultaneously with heuristics, approximation, or relaxation of the objective. Nevertheless, these treatments often hinder the rigorous theoretical analysis, particularly regarding regret, which is well-established for unconstrained BO tasks Srinivas et al. [2009], Chowdhury and Gopalan [2017]. The key challenge lies in the fact that an optimum outside the feasible region or insufficient exploration within the feasible region can lead to unbounded regret. Therefore, it is crucial for CBO algorithms to rigorously guarantee the *accurate and efficient* identification of the entire feasible region. This ensures that the algorithm remains within the feasible region and selects near-optimal candidates with high probability after sufficient iterations. Unfortunately, this guarantee is

Approach	Feasible Region Identification	Regret Guarantees	Constraint Learning and Trade-off
Extension of Unconstrained Methods	Not explicitly guaranteed	No established convergence rate with respect to constraints	No adaptive trade-off between objective and constraints
Violation-Tolerant Objectives	May select infeasible points despite decreasing violations	Provides upper bounds on regret, but only when infeasible points yield meaningful rewards	Penalizes constraint violations separately from the objective
Active Learning of Constraints	Typically limited to one constraint at a time	Guarantees available only for level-set estimation	Splits constraint learning and objective optimization into separate phases
COBAR (Proposed)	Guaranteed identification of the feasible region	Non-asymptotic regret convergence rate	Integrated, adaptive trade-off between objective optimization and constraint learning

Table 1: Comparison of existing CBO approaches and COBAR (Proposed). The first three rows serve as a high-level summary of the respective approaches discussed in the preceding paragraphs of Section 2, where detailed discussion and citations are provided.

absent in existing methods, which motivates our approach to incorporate active learning of unknown constraints in designing a CBO algorithm with a regret guarantee.

In this paper, we propose a novel framework that integrates active learning for level-set estimation (AL-LSE) [Gotovos et al., 2013, Nguyen et al., 2021] with Bayesian optimization for constrained problems. Our approach leverages the theoretical strengths of both paradigms, enabling a rigorous performance analysis of the proposed CBO method in finding the interior optimum. A brief illustration of the framework design is shown in figure 1. The remainder of this paper is organized as follows. We first provide a detailed overview of recent advancements in various facets of CBO. Next, we formally state the CBO problem and introduce the definition of probabilistic regret as a performance metric that facilitates rigorous analysis. Following the problem statement, we present our novel CBO framework, offer the corresponding performance analysis, and provide empirical evidence for its efficacy. Finally, we summarize the key contributions of our work and discuss potential directions for future research.

2 RELATED WORK

Extension of unconstrained methods While the majority of BO research focuses on unconstrained problems [Bernardo et al., 2011, Frazier, 2018, Gramacy, 2020, Binois and Wycoff, 2022, Garnett, 2023], several works also address black-box constraints. The pioneering work by Schonlau et al. [1998] first extended Expected Improvement (EI) to the constrained setting, and subsequent developments [Gelbart et al., 2014, Gardner et al., 2014, Feliot et al., 2017, Letham et al., 2019, Wang et al., 2024] refined this

approach by defining the acquisition function at a given point as the product of the expected improvement and the probability that the point is feasible. In addition, the posterior sampling method (Thompson sampling) was extended to scalable CBO (SCBO) by Eriksson and Poloczek [2021], generalizing the unconstrained TuRBO approach [Eriksson et al., 2019] by incorporating additional samples from the constraint posteriors to weight the objective samples. Methods based on information criteria [Hernández-Lobato et al., 2014, Wang and Jegelka, 2017] have also been extended to the constrained setting [Hernández-Lobato et al., 2015, Perrone et al., 2019, Takeno et al., 2022], although these approaches rely heavily on sampling-based approximations. Another line of work transforms the CBO task into an unconstrained problem via the augmented Lagrangian framework [Gramacy et al., 2016, Picheny et al., 2016, Ariaifar et al., 2019], allowing vanilla BO to be applied as a subroutine, particularly in decoupled settings. In general, these methods do not guarantee the identification of the feasible region during optimization, and consequently, they lack a convergence rate for regret that accounts for feasibility.

Violation-tolerant objectives In addition to the aforementioned empirical approaches, recent works [Zhou and Ji, 2022, Lu and Paulson, 2022, Guo et al., 2023, Xu et al., 2023] have considered a relaxed CBO objective to facilitate theoretical analysis of convergence rates. These works assume that queries outside the feasible region still incur a reward and incorporate constraint violations either as a weighted penalty within the regret or analyze them separately from the objective’s regret. Although they provide upper bounds on both the constraint violations and the regret, these methods do not adequately address the issue of potentially infinite regret arising from evaluations outside the

feasible region. Since diminishing constraint violations do not guarantee the eventual selection of a feasible point, the analysis in these works diverges from our objective without nontrivial modifications.

Active learning of constraints The concept of data selection in active learning dates back to MacKay’s work [MacKay, 1992], and stepwise uncertainty reduction (SUR) has been used to estimate failure probabilities in industrial settings [Bect et al., 2012]. The principled approach of active learning for level-set estimation (AL-LSE) was introduced by Gotovos et al. [2013] to perform classification over the sample space, offering theoretical guarantees on convergence rates. Since both AL-LSE and BO employ Gaussian processes to model underlying functions, Bogunovic et al. [2016] unified these problems using truncated variance reduction and by selecting a kernel that ensures submodularity of the variance reduction. However, a direct application of level-set estimation methods is limited by their focus on a single unknown function and the lack of a straightforward extension to balance the learning of multiple unknown functions. While Malkomes et al. [2021], Komiyama et al. [2022] proposed novel acquisition functions that prioritize diversity in the active search, these approaches do not provide a mechanism for adaptively trading off between constraint learning and objective optimization. A similar challenge is encountered in Antonio [2021], where the algorithm decouples the learning of the feasible region from the optimization of the objective by addressing them in two separate phases. These decoupled, two-phase approaches are fundamentally different from our goal of an integrated algorithm that adaptively balances learning and optimization in every step.

3 PROBLEM STATEMENT

In this section, we introduce a few useful notations and formalize the problem. Consider a compact search space $\mathbf{X} \subseteq \mathbb{R}^d$. We aim to find a maximizer $\mathbf{x}^* \in \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x})$ of a black-box function $f : \mathbf{X} \rightarrow \mathbb{R}$, subject to M black-box constraints $\mathcal{C}_m(\mathbf{x})$ ($m \in \mathbf{M} = \{1, 2, 3, \dots, M\}$) such that each constraint is satisfied by staying above its corresponding threshold h_m . For simplicity and without loss of generality, we let all $h_m = 0$. Thus, formally, our goal can be formulated as finding the *interior optimum*:

$$\max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}) \text{ s.t. } \mathcal{C}_m(\mathbf{x}) > 0, \forall m \in \mathbf{M}$$

We maintain a Gaussian process (\mathcal{GP}) as the surrogate model for each black-box function, pick a point $\mathbf{x}_t \in \mathbf{X}$ at iteration t by maximizing the acquisition function $\alpha : \mathbf{X} \rightarrow \mathbb{R}$, and observe the function values perturbed by additive noise: $y_{f,t} = f(\mathbf{x}_t) + \epsilon$ and $y_{\mathcal{C}_m,t} = \mathcal{C}_m(\mathbf{x}_t) + \epsilon$, with $\epsilon \sim \mathcal{N}(0, \sigma^2)$ being i.i.d. Gaussian noise. Each $\mathcal{GP}(\mu(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ is fully specified by its

prior mean μ and kernel k . With the historical observations $\mathbf{S}_{t-1} = \{(\mathbf{x}_i, y_{f,i}, \{y_{\mathcal{C}_m,i}\}_{m \in \mathbf{M}})\}_{i=1,2,\dots,t-1}$, the posterior also takes the form of a \mathcal{GP} , with mean

$$\mu_t(\mathbf{x}) = k_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma^2 I)^{-1} \mathbf{y}_t \quad (1)$$

and covariance

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

where $k_t(\mathbf{x}) \triangleq [k(\mathbf{x}_1, \mathbf{x}), \dots, k(\mathbf{x}_t, \mathbf{x})]^\top$ and $\mathbf{K}_t \triangleq [k(\mathbf{x}, \mathbf{x}')]_{\mathbf{x}, \mathbf{x}' \in \mathbf{S}_{t-1}}$ is a positive definite kernel matrix [Rasmussen and Williams, 2006].

The definition of reward plays an important role in analyzing online learning algorithms. Throughout the rest of the paper, we define the reward of CBO as the following and defer the detailed discussion of alternative reward choices to Appendix F.

$$r(\mathbf{x}_t) = \begin{cases} y_{f,t} & \text{if } \mathbb{I}(y_{\mathcal{C}_m(\mathbf{x}_t)} \geq 0) \quad \forall m \in \mathbf{M} \\ -\inf & \text{o.w.} \end{cases} \quad (3)$$

We want to locate the global maximizer efficiently

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}, \forall m \in \mathbf{M}, \mathcal{C}_m(\mathbf{x}) > 0} f(\mathbf{x})$$

More specifically, we seek to establish an upper bound on the performance in terms of expected regret at a certain time t , with respect to the distribution over f at time t given historical observation \mathbf{S}_{t-1} ,

$$\mathbf{R}_t(\mathbf{x}) \triangleq \mathbb{E}_f [r(\mathbf{x}^*) \mid \mathbf{S}_{t-1}] - \mathbb{E}_f [r(\mathbf{x}) \mid \mathbf{S}_{t-1}]$$

Formally, given a certain confidence level δ and constant ϵ_f , we want to guarantee that after using up certain budget T dependent on δ and ϵ_f , we could achieve a high probability upper bound of the regret on the identified area $\tilde{\mathbf{X}}$ which is the subset of \mathbf{X} :

$$P \left(\max_{\mathbf{x} \in \tilde{\mathbf{X}}} \mathbf{R}_T(\mathbf{x}) \geq \epsilon_f \right) \leq \delta.$$

Remark 1. As discussed in the literature [Antonio, 2021, Donskoi, 1993, Rudenko, 1994, Sergeyev et al., 2007, Sacher et al., 2018, Bachoc et al., 2020], the reward of CBO is not defined outside the feasible region. This reward definition in equation 3, along with both the aleatoric and epistemic uncertainties of the underlying black-box functions, necessitates excluding boundary candidates from the formulation to ensure the soundness of the CBO objective. For example, consider a scenario where $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}, \forall m \in \mathbf{M}, \mathcal{C}_m(\mathbf{x}) \geq 0} f(\mathbf{x})$ and $\exists m \in \mathbf{M}, \mathcal{C}_m(\mathbf{x}^*) = 0$. In this case, the observation $y_{\mathcal{C}_m,t} = \mathcal{C}_m(\mathbf{x}_t) + \epsilon$ at $\mathbf{x}_t = \mathbf{x}^*$ will be purely noise (ϵ), with $\mathbb{P}[y_{\mathcal{C}_m,t} < 0] = 0.5$. According to equation 3, this results

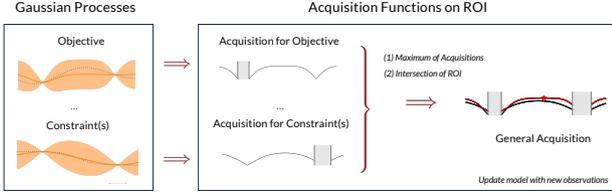


Figure 1: The pipeline of our proposed algorithm, COBAR. In the left box, we maintain Gaussian process surrogates for the unknown objective and each constraint. The dotted curve shows the actual function, the red curve is the predicted mean, and the shaded area is the confidence interval. In the right box, we derive acquisition functions from each Gaussian process. The general acquisition function combines these, but only over specific Regions of Interest (ROIs). The grey area in the acquisition plot represents a region excluded by our ROI identification, a process detailed in Section 4. For a step-by-step visual breakdown of this filtering, please see figure 2.

in $\mathbb{P}[r(\mathbf{x}_t) = -\inf] \geq 0.5$, making it impossible to achieve a high-probability regret bound, as previously discussed. Therefore, we aim to find the interior optimum of the black-box constrained objective.

4 THE COBAR ALGORITHM

We start by introducing necessary notions from active learning for level-set estimation, followed by a detailed description of our proposed algorithm.

4.1 ACTIVE LEARNING FOR LEVEL-SET ESTIMATION

We follow the common practice and assume the objective and each unknown constraint is sampled from a corresponding independent Gaussian process (\mathcal{GP}) [Hernández-Lobato et al., 2015, Gelbart et al., 2014, Gotovos et al., 2013] to treat the epistemic uncertainty.

$$\begin{aligned} f &\sim \mathcal{GP}_f \\ \mathcal{C}_m &\sim \mathcal{GP}_{\mathcal{C}_m} \quad \forall m \in \mathbf{M} \end{aligned}$$

We could derive pointwise confidence interval estimation with the \mathcal{GP} for each black-box function. We define the upper confidence bound $\text{UCB}_t(\mathbf{x}) \triangleq \mu_{t-1}(\mathbf{x}) + \beta_t^{1/2} \sigma_{t-1}(\mathbf{x})$ and lower confidence bound $\text{LCB}_t(\mathbf{x}) \triangleq \mu_{t-1}(\mathbf{x}) - \beta_t^{1/2} \sigma_{t-1}(\mathbf{x})$, where $\sigma_{t-1}(\mathbf{x}) = k_{t-1}(\mathbf{x}, \mathbf{x})^{1/2}$ and β_t acts as a scaling factor corresponding to certain confidence.

For each unknown constraint \mathcal{C}_m , we follow the notations from Gotovos et al. [2013] and define the superlevel-set to be

the areas that meet the constraint \mathcal{C}_m with high confidence

$$S_{\mathcal{C}_m,t} \triangleq \{\mathbf{x} \in \mathbf{X} \mid \text{LCB}_{\mathcal{C}_m,t}(\mathbf{x}) > 0\}$$

We define the sublevel-set to be the areas that do not meet the constraint \mathcal{C}_m with high confidence

$$L_{\mathcal{C}_m,t} \triangleq \{\mathbf{x} \in \mathbf{X} \mid \text{UCB}_{\mathcal{C}_m,t}(\mathbf{x}) < 0\}$$

and the undecided set is defined as

$$U_{\mathcal{C}_m,t} \triangleq \{\mathbf{x} \in \mathbf{X} \mid \text{UCB}_{\mathcal{C}_m,t}(\mathbf{x}) \geq 0, \text{LCB}_{\mathcal{C}_m,t}(\mathbf{x}) \leq 0\}$$

where the points remain to be classified.

4.2 REGION OF INTEREST IDENTIFICATION FOR EFFICIENT CBO

In the CBO setting, we only care about the superlevel-set $S_{\mathcal{C}_m,t}$ and undecided-set $U_{\mathcal{C}_m,t}$, where the global optimum is likely to lie in. Hence, we define the region of interest for each constraint function \mathcal{C}_m as

$$\hat{\mathbf{X}}_{\mathcal{C}_m,t} \triangleq S_{\mathcal{C}_m,t} \cup U_{\mathcal{C}_m,t} = \{\mathbf{x} \in \mathbf{X} \mid \text{UCB}_{\mathcal{C}_m,t}(\mathbf{x}) \geq 0\}$$

Similarly, for the objective function, though there is no pre-specified threshold, we could use the maximum of $\text{LCB}_f(\mathbf{x})$ on the intersection of superlevel-set $S_{\mathcal{C},t} \triangleq \bigcap_m^{\mathbf{M}} S_{\mathcal{C}_m,t}$

$$\text{LCB}_{f,t,\max} \triangleq \begin{cases} \max_{\mathbf{x} \in S_{\mathcal{C},t}} \text{LCB}_{f,t}(\mathbf{x}), & \text{if } S_{\mathcal{C},t} \neq \emptyset \\ -\infty, & \text{otherwise} \end{cases}$$

as the high confidence threshold for the $\text{UCB}_{f,t}(\mathbf{x})$ to identify a region of interest for the optimization of the objective. Given that $\text{UCB}_{f,t}(\mathbf{x}^*) \geq f^* \geq f(\mathbf{x}) \geq \text{LCB}_{f,t}(\mathbf{x})$ with the probability specified by the choice of β_t , we define the ROI for the objective optimization as

$$\hat{\mathbf{X}}_{f,t} \triangleq \{\mathbf{x} \in \mathbf{X} \mid \text{UCB}_{f,t}(\mathbf{x}) \geq \text{LCB}_{f,t,\max}\}$$

By taking the intersection of the ROI of each constraint, we could identify the ROI for identifying the feasible region

$$\hat{\mathbf{X}}_{\mathcal{C},t} \triangleq \bigcap_m^{\mathbf{M}} \hat{\mathbf{X}}_{\mathcal{C}_m,t}$$

The combined ROI for CBO is determined by intersecting the ROIs of constraints and the objective:

$$\hat{\mathbf{X}}_t \triangleq \hat{\mathbf{X}}_{f,t} \cap \hat{\mathbf{X}}_{\mathcal{C},t} \quad (4)$$

4.3 COMBINING ACQUISITION FUNCTIONS FOR CBO

Acquisition function for optimizing the objective To optimize the unknown objective f when $\hat{\mathbf{X}}_t$ is established, we can employ the following acquisition function ¹

$$\alpha_{f,t}(\mathbf{x}) \triangleq \begin{cases} \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t,\max} & S_{\mathcal{C},t} \neq \emptyset \\ \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t}(\mathbf{x}) & \text{otherwise} \end{cases} \quad (5)$$

At given t , to efficiently optimize the black-box f we evaluate the point $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \hat{\mathbf{X}}_t} \alpha_{f,t}(\mathbf{x})$. Since at a given t , when $\text{LCB}_{f,t,\max}$ is constant, the acquisition function is equivalent to $\text{UCB}_{f,t}(\mathbf{x})$.

Acquisition function for learning the constraints When we merely focus on identifying the feasible region defined by a certain unknown constraint \mathcal{C}_k , we could apply the following active learning acquisition function.

$$\alpha_{\mathcal{C}_m,t}(\mathbf{x}) \triangleq \text{UCB}_{\mathcal{C}_m,t}(\mathbf{x}) - \text{LCB}_{\mathcal{C}_m,t}(\mathbf{x}) \quad (6)$$

At given t , we evaluate the point $\mathbf{x}_t = \arg \max_{\mathbf{x} \in U_{\mathcal{C}_m,t} \cap \hat{\mathbf{X}}_t} \alpha_{\mathcal{C}_m,t}(\mathbf{x})$ to efficiently identify the feasible region defined by \mathcal{C}_m . Note that the acquisition function $\alpha_{\mathcal{C}_m,t}(\mathbf{x})$ is not maximized on the full $\hat{\mathbf{X}}_{\mathcal{C}_m,t}$, but only on $U_{\mathcal{C}_m,t} \cap \hat{\mathbf{X}}_t$. The active learning on the superlevel-set $S_{\mathcal{C}_m,t} \cap \hat{\mathbf{X}}_t$ doesn't contribute to identifying the corresponding feasible region.

Adaptive Selection Strategy With the acquisitions and ROIs established, we propose the algorithm **CO**nstrained **BO** through **A**daptive **R**egion of Interest Acquisition (COBAR)², with its full procedure detailed in Algorithm 1. To clarify its selection logic, we elaborate on the core mechanism here. At each iteration t , the algorithm first compiles a set \mathcal{G} of candidate functions. This set includes the objective function f and any constraint \mathcal{C}_m that still has an associated region of uncertainty (i.e., its undecided set $U_{\mathcal{C}_m,t}$ is non-empty). For each function $g \in \mathcal{G}$, the algorithm then finds the best candidate point by maximizing its own acquisition function: $\mathbf{x}_{f,t}$ for the objective (line 9) and $\mathbf{x}_{\mathcal{C}_m,t}$ for each uncertain constraint (line 7). The adaptive trade-off occurs in line 11, which implements a "winner-takes-all" strategy. The algorithm compares the acquisition values of all candidates (e.g., $\alpha_{f,t}(\mathbf{x}_{f,t})$ vs. all relevant $\alpha_{\mathcal{C}_m,t}(\mathbf{x}_{\mathcal{C}_m,t})$) and selects the function g_t that offers the maximum value. The next point to query, \mathbf{x}_t , is simply the candidate associated with this winning function (line 12). This process

¹Same criterion has been studied under the unconstrained setting [Zhang et al., 2023].

²We briefly discuss the possible extension to decoupled setting, where the objective and constraints may be evaluated independently, of COBAR in Appendix D.

allows COBAR to pivot dynamically between optimizing the objective and reducing constraint uncertainty. For a more detailed walkthrough, please see Appendix I.1 and the illustration in figure 2.

Implementation Details Algorithm 1 is presented for a general search space \mathbf{X} . For practical implementation, as in our experiments, we operate on a large, finite discretization of the space, denoted by $\tilde{D} \subset \mathbf{X}$. Consequently, all identification steps (line 4) and maximization steps (lines 7 and 9) are performed over the relevant discrete subset of candidate points. For instance, the search domain $\hat{\mathbf{X}}_t$ is replaced by its discrete counterpart, $\tilde{D}_{\hat{\mathbf{X}}_t} \triangleq \hat{\mathbf{X}}_t \cap \tilde{D}$. The membership of each point in \tilde{D} to the various ROIs and undecided sets can be checked in a pointwise fashion, making the identification steps (line 4) computationally straightforward. While this work focuses on a discrete search space for theoretical clarity, we discuss the extension to continuous domains, where acquisition functions could be optimized using standard gradient-based methods, in Appendix C.

We also illustrate the detailed procedure on a 1D toy example in figure 2. We construct the example to demonstrate that the explicit, active learning of the constraint doesn't necessarily hurt the optimization but could contribute directly to the simple regret improvement.

Algorithm 1 **CO**nstrained **BO** through **A**daptive **R**egion of Interest Acquisition (COBAR).

- 1: **Input:** Search space \mathbf{X} , initial observation \mathbf{S}_0 , horizon T , confidence factor δ , confidence coefficient β ;
 - 2: **for** $t = 1$ **to** T **do**
 - 3: Update the posteriors of $\mathcal{GP}_{f,t}$ and $\mathcal{GP}_{\mathcal{C}_m,t}$ according to equation 1 and 2
 - 4: Identify ROIs $\hat{\mathbf{X}}_t$, and undecided sets $U_{\mathcal{C}_m,t}$
 - 5: **for** $m \in \mathbf{M}$ **do**
 - 6: **if** $U_{\mathcal{C}_m,t} \neq \emptyset$ **then**
 - 7: Candidate for learning of each constraint:
 $\mathbf{x}_{\mathcal{C}_m,t} \leftarrow \arg \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{X}}_t} \cap U_{\mathcal{C}_m,t}} \alpha_{\mathcal{C}_m,t}(\mathbf{x})$ (6)
 - 8: $\mathcal{G} \leftarrow \mathcal{G} \cup \mathcal{C}_m,t$
 - 9: Candidate for optimizing the objective:
 $\mathbf{x}_{f,t} \leftarrow \arg \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{X}}_t}} \alpha_{f,t}(\mathbf{x})$ as in equation 5
 - 10: $\mathcal{G} \leftarrow \mathcal{G} \cup f$
 - 11: Maximize the acquisition from different aspects:
 $g_t \leftarrow \arg \max_{g \in \mathcal{G}} \alpha_{g,t}(\mathbf{x}_{g,t})$
 - 12: Pick the candidate to evaluate: $\mathbf{x}_t \leftarrow \mathbf{x}_{g_t,t}$
 - 13: Update the observation set
 $\mathbf{S}_t \leftarrow \mathbf{S}_{t-1} \cup \{(\mathbf{x}_t, y_{f,t}, \{y_{\mathcal{C}_m,t}\}_{m \in \mathbf{M}})\}$
-

Exploration and Multiple Feasible Regions A key challenge in CBO is ensuring adequate exploration, particularly when the feasible space is non-convex or consists of multiple disjoint regions. COBAR addresses this through its acquisition functions (equation 5 and equation 6). By design, the

acquisition drives the exploration: regions with high model uncertainty (i.e., large confidence intervals) yield high acquisition values, which naturally encourages the algorithm to sample in less-explored areas. This mechanism is crucial for discovering initially unknown feasible regions. While our Region of Interest ($\hat{\mathbf{X}}_t$) focuses the search on promising areas for efficiency, it is dynamic and evolves as the GP models are updated, allowing the search to expand into new areas as uncertainty dictates. We acknowledge that guaranteeing the discovery of all disconnected feasible regions while maintaining rapid convergence is a difficult trade-off, a known challenge in global optimization. Nevertheless, we provide empirical evidence of COBAR’s robustness in such a scenario in our Rastrigin-1D-1C experiment (see figure 3), where our method successfully navigates a search space with multiple feasible regions to find the global optimum.

5 THEORETICAL ANALYSIS

We first state a few assumptions that provide insights into the convergence properties of COBAR. The first one follows Srinivas et al. [2009] as a standard assumption for BO.

Assumption 1. *The objective and constraints are sampled from independent Gaussian processes. Formally, for all $t < T$ and $\mathbf{x} \in \mathbf{X}$, $f(\mathbf{x})$ is a sample from $\mathcal{GP}_{f,t}$, and $C_m(\mathbf{x})$ is a sample from $\mathcal{GP}_{C_m,t}$, for all $m \in \mathbf{M}$.*

The second one assumes that the global optimum lies inside the feasible region due to the reason discussed in Remark 1.

Assumption 2. *A global optimum exists within the feasible region. The distance between this global optimum and the boundaries of the feasible regions is uniformly bounded below by ϵ_C . More specifically, for all $m \in \mathbf{M}$, $\exists \epsilon_m > 0$ such that $C_m(\mathbf{x}^*) > \epsilon_m$, then it holds that $C_m(\mathbf{x}^*) > \epsilon_C = \min_{m \in \mathbf{M}} \epsilon_m$.*

We will also show that without Assumption 2, it is possible to bound both the constraint violations and the regret—defined independently of feasibility—with minor adjustments as discussed in Remark 2.

Assumption 3. *Given a proper choice of β_t that is non-increasing, the confidence intervals are consistent. Concretely, $\forall t_1 < t_2 < T$ and $\mathbf{x} \in \mathbf{X}$, if $\beta_{t_1} \geq \beta_{t_2}$, then $UCB_{t_1}(\mathbf{x}) \geq UCB_{t_2}(\mathbf{x})$ and $LCB_{t_1}(\mathbf{x}) \leq LCB_{t_2}(\mathbf{x})$.*

This is a mild assumption as long as β_t is non-increasing, given recent work by Koepf and Pfaff [2021] showing that if the kernel is continuous and the sequence of sampling points lies sufficiently dense, the variance of the posterior \mathcal{GP} converges to zero almost surely monotonically if the function is in metric space. If the assumption is violated, the technique of taking the intersection of all historical confidence intervals introduced by Gotovos et al. [2013] could

similarly guarantee a monotonically shrinking confidence interval. That is, when $\exists t_1 < t_2 < T$, $\mathbf{x} \in \mathbf{X}$, if we have $UCB_{t_1}(\mathbf{x}) < UCB_{t_2}(\mathbf{x})$ or $LCB_{t_1}(\mathbf{x}) > LCB_{t_2}(\mathbf{x})$, we let $UCB_{t_2}(\mathbf{x}) = UCB_{t_1}(\mathbf{x})$ or $LCB_{t_2}(\mathbf{x}) = LCB_{t_1}(\mathbf{x})$ to guarantee the monotonicity. To allow for a plug-in of the intersection technique, and without loss of accuracy, we keep using the notation UCB and LCB without further parsing the value in the following discussion of algorithm design and theoretical analysis. The cost of violating the Assumption 3 has been studied in corollary 3 by Zhang et al. [2023]. We refrain from repeating the analysis here.

The following lemma justifies the definition of the region(s) of interest $\hat{\mathbf{X}}_t$ defined in equation 4. For clarity, we denote $\tilde{D}_{\hat{\mathbf{X}}_t} = \tilde{D} \cap \hat{\mathbf{X}}_t$, and $CI_{f^*,t} = [\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{X}}_t}} LCB_t(\mathbf{x}), \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{X}}_t}} UCB_t(\mathbf{x})]$.

Lemma 1. *Under the assumptions above, the regions of interest $\hat{\mathbf{X}}_t$, as defined in equation 4, contain the global optimum with high probability. Formally, for all $\delta \in (0, 1)$, $T \geq t \geq 1$, and any finite discretization \tilde{D} of \mathbf{X} that contains the optimum $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x})$ where $C_m(\mathbf{x}^*) > \epsilon_C$ for all $m \in \mathbf{M}$ and $\beta_t = 2 \log(2(M+1)|\tilde{D}| \pi_t / \delta)$ with $\sum_{T \geq t \geq 1} \pi_t^{-1} = 1$, we have $\mathbb{P}[\mathbf{x}^* \in \tilde{D}_{\hat{\mathbf{X}}_t}] \geq 1 - \delta$.*

To guarantee β_t to be non-increasing, we could let $\pi_t = T$ and therefore $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ is a constant. The lemma shows that with proper choice of prior and β , the $\hat{\mathbf{X}}_{f,t}$ remains nonempty during optimization.

Subsequently, let’s define the maximum information gain about function f after T rounds: $\gamma_{f,T} = \max_{A \subset \tilde{D}: |A|=T} \mathbb{I}(y_A; f_A)$ and

$$\widehat{\gamma}_T = \sum_{g \in \{f\} \cup \{C_m\}_{m \in \mathbf{M}}} \gamma_{g,T} \quad (7)$$

In the following, we show that we could bound the simple regret of COBAR after sufficient rounds. Concretely, in Theorem 1, we provide an upper bound on the width of the confidence interval for the global optimum $f^* = f(\mathbf{x}^*)$.

Theorem 1. *Under the aforementioned assumptions, with a constant $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ and the acquisition function from Algorithm 1, there exists an $\epsilon_f \leq \epsilon_C$, such that after at most $T \geq \frac{\beta \widehat{\gamma}_T C_1}{\epsilon_f^2}$ iterations, we have $\mathbb{P}[|CI_{f^*,T}| \leq \epsilon_f, f^* \in CI_{f^*,T}] \geq 1 - \delta$. Here, $C_1 = 8 / \log(1 + \sigma^{-2})$.*

Note $\beta \widehat{\gamma}_T C_1$ is sublinear with respect to T . One direct result of Theorem 1 is that if any point belongs to \tilde{D} that lies in the feasible set defined by the unknown constraints bears a suboptimal gap on the reward except for the global optimum, then after sufficient query, the algorithm will identify \mathbf{x}^* as the only point in the ROI. In that case, COBAR will only query \mathbf{x}^* and achieve zero regret afterward.

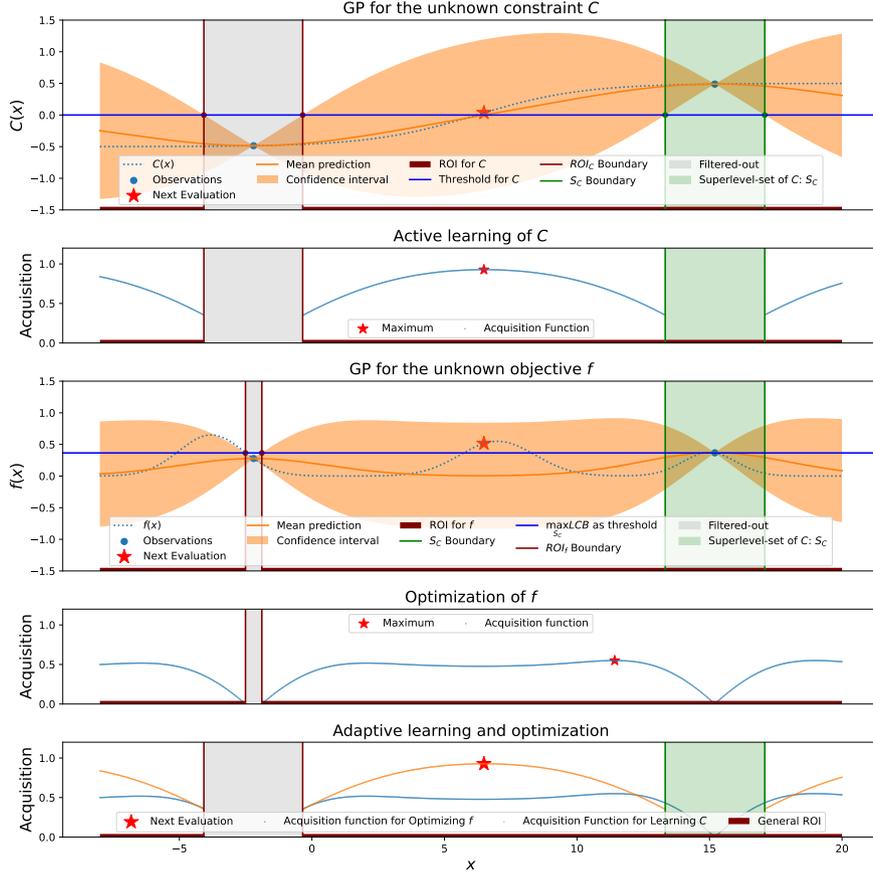


Figure 2: Illustration of COBAR on a synthetic noise-free 1D example. The first two rows show the GP for the \mathcal{C} , the superlevel-set $S_{\mathcal{C}}$, the region of interest $\hat{\mathbf{X}}_{\mathcal{C}}$ and the corresponding acquisition function $\alpha_{\mathcal{C},t}(\mathbf{x})$ as defined in equation 6. The following two rows show the GP for f , the region of interest $\hat{\mathbf{X}}_f$, and the corresponding acquisition function $\alpha_{f,t}(\mathbf{x})$ defined in equation 5. We show that after identifying $S_{\mathcal{C}}$, we could define the threshold for ROI identification of f accordingly. The bottom row demonstrates that the general ROI $\hat{\mathbf{X}}$ as defined in equation 4 is identified by taking the intersection ROI for f and \mathcal{C} . The general acquisition function is defined as the maximum of the acquisition for f and \mathcal{C} and is maximized on the $\hat{\mathbf{X}}$. The scaling and length scale of the GPs are learned via maximum likelihood estimation.

Corollary 1. We assume the aforementioned conditions hold, and $\forall \mathbf{x} \in \tilde{D}$, when $\forall m \in \mathbf{M}$, $C_m(\mathbf{x}) > 0$, $\mathbf{x} \neq \mathbf{x}^*$, it holds that $\exists \epsilon_{\mathcal{C}} \geq 2\epsilon_f > 0$, $f^* - f(\mathbf{x}) > 2\epsilon_f$. In addition, we use $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ and the acquisition function from Algorithm 1. After at most $t \geq \frac{\beta \hat{\gamma}_t C_1}{\epsilon_f^2}$ iterations, we have $\mathbb{P}[\mathbf{R}_t = 0] \geq 1 - \delta$. Here, $C_1 = 8/\log(1 + \sigma^{-2})$ and $t \leq T$.

Similarly, if a group of suboptimal candidates lies in the feasible area and is sufficiently close to \mathbf{x}^* , then Assumption 2 also holds for those suboptimal points. In this condition, the algorithm achieves a sublinear cumulative regret after identifying this near-optimal region.

Corollary 2. We assume the aforementioned conditions hold, and $\forall \mathbf{x} \in \tilde{D}$, when $\forall m \in \mathbf{M}$, $C_m(\mathbf{x}) > 0$, $\mathbf{x} \neq \mathbf{x}^*$, $\exists \epsilon_{\mathcal{C}} \geq \epsilon_f > 0$, $f^* - f(\mathbf{x}) \leq 2\epsilon_f$, it holds that $\forall m \in \mathbf{M}$, $C_m(\mathbf{x}) \geq \epsilon_{\mathcal{C}}$. In addition, we use $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ and the acquisition function from Algorithm 1. After at most $t' \geq \frac{\beta \hat{\gamma}_{t'} C_1}{\epsilon_f^2}$ iterations, we have,

$$\mathbb{P}\left[\sum_{t=t'}^T r(\mathbf{x}^*) - r(\mathbf{x}_t) \leq \sqrt{(T-t')\beta\gamma_T C_1}\right] \geq 1 - \delta.$$

Here, $C_1 = 8/\log(1 + \sigma^{-2})$ and $t' \leq T$.

Following the path of proof for Theorem 1, with Lemma 1, we can show that the algorithm can identify infeasibility when all points in the search space violate at least one of the constraints at least $\epsilon'_{\mathcal{C}}$. Concretely, $\forall \mathbf{x} \in \mathbf{X}$, if it holds that $\exists m \in \mathbf{M}$, $C_m(x) < -\epsilon'_{\mathcal{C}}$, with high probability the identified $\tilde{D}_{\hat{\mathbf{X}}_T} = \emptyset$.

Corollary 3. When the assumptions except for Assumption 2 hold, $\forall \mathbf{x} \in \mathbf{X}$, if $\exists m \in \mathbf{M}$, $C_m(x) < -\epsilon'_{\mathcal{C}}$, then with a constant $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ and the acquisition function from Algorithm 1, after at most $T \geq \frac{\beta \hat{\gamma}_T C_1}{\epsilon'^2_{\mathcal{C}}}$ iterations, we have $\mathbb{P}[\tilde{D}_{\hat{\mathbf{X}}_T} = \emptyset] \geq 1 - \delta$. Here, $C_1 = 8/\log(1 + \sigma^{-2})$.

The above algorithm and theoretical results assume that a discretization \tilde{D} is given but is compatible with any density of the discretization. This means that with additional

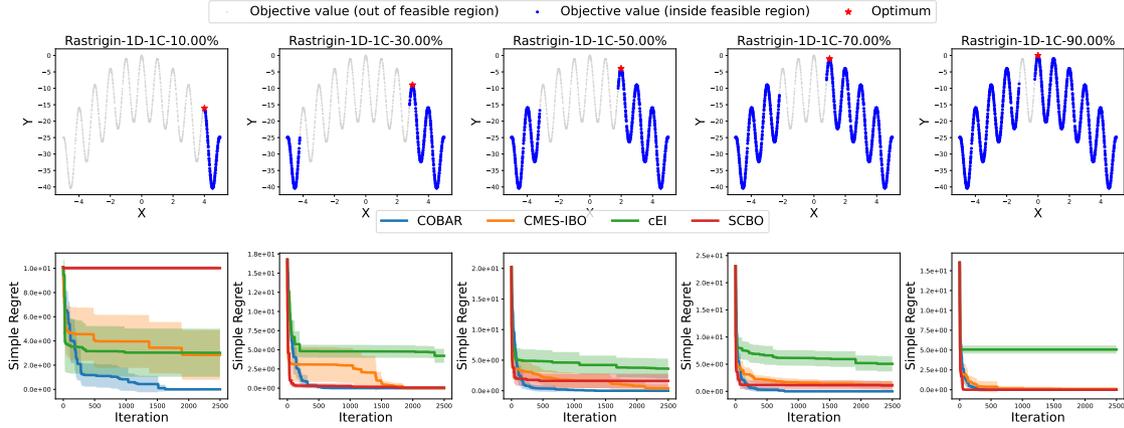


Figure 3: We use black dots and purple dots to show the infeasible region and feasible region in the first row correspondingly. Each column corresponds to a certain threshold choice for the single constraint $c(\mathbf{x}) = |\mathbf{x} + 0.7|^{1/2}$ in the Rastrigin-1D-1C task. The search space contains a certain portion of the feasible region, denoted on each figure and title. The first row shows the distribution of 1000 samples from the noise-free distribution objective function, and the figures are differentiated with different feasible regions. The second row shows corresponding simple regret curves. We test each method with 15 independent trails and impose observation noises sampled from $\mathcal{N}(0, 0.1)$ not shown in the first row. The scaling and length scale of the GPs are learned via maximum likelihood estimation.

assumptions on the underlying functions, we could adapt the algorithm to a continuous setting by taking a sufficiently dense discretization on a proper embedding space.³

Remark 2. *If the goal is to find the boundary optimum despite the feasibility concerns highlighted in Remark 1, a practical approach is to uniformly shift the constraints by a small amount ϵ_c to satisfy Assumption 2 with the modified constraints. Formally, $\forall m \in \mathbf{M}$, $C'_m(\mathbf{x}) = C_m(\mathbf{x}) + \epsilon_c$. Then, running COBAR with these adjusted constraints, C'_m , instead of the original C_m , yields similar guarantees as those in Theorem 1 and Corollary 2, with a high probability that any instantaneous violations of the original constraints are uniformly bounded by ϵ_c . Further details are discussed in Appendix E.*

6 EXPERIMENTS

In this section, we empirically study the performance of COBAR against three baselines, including (1) cEI, the extension of EI into CBO from Gelbart et al. [2014], (2) cMES-IBO, a state-of-the-art information-based approach by Takeno et al. [2022], and (3) SCBO, a recent Thompson Sampling (TS) method tailored for scalable CBO from Eriksson and Poloczek [2021]. We abstain from comparison against Augmented-Lagrangian methods, following the practice of Takeno et al. [2022], as past studies have illustrated its inferior performance against sampling methods [Eriksson and Poloczek, 2021] or information-based methods [Takeno et al., 2022, Hernández-Lobato et al., 2014]. We defer the comparison against CONFIG Xu et al. [2023] to

³With additional assumptions on the regularization of the underlying function, we derive the analogous analysis on continuous search space in Appendix C.

Appendix H.3, due to the difference in objective and a resulting instability on our benchmarks. We begin by describing the optimization tasks, and then discuss the performances. To guarantee a fair comparison across all methods, results are averaged over multiple independent trials. For each trial, the random seed is set universally based only on the trial number, ensuring that every algorithm is evaluated under identical stochastic conditions.

6.1 CBO TASKS

We compare COBAR against the aforementioned baselines across six CBO tasks. The first two synthetic CBO tasks are constructed from conventional BO benchmark tasks [Balandat et al., 2020]. Among the other four real-world CBO tasks, the first three are extracted from Tanabe and Ishibuchi [2020], offering a broad selection of multi-objective multi-constraints optimization tasks. The fourth one is a 32-dimensional optimization task extracted from the UCI Machine Learning repository [mis, 2019]. Further details about the datasets are available in Appendix G.

- The *Rastrigin function* is a non-convex function used as a performance test problem for optimization algorithms. It was first proposed by Rastrigin [1974] and used as a popular benchmark dataset [Pohlheim, 2006]. The feasible region takes up approximately 60% of the search space, which we construct by sampling $|\tilde{D}| = 20000$ and reuse for all 15 trials. We also vary the threshold to control the portion of the feasible region to study the robustness of COBAR. Figure 3 shows the distribution of the objective function and feasible regions.
- The *Ackley function* is another commonly used optimization benchmark. We construct two constraints to enforce

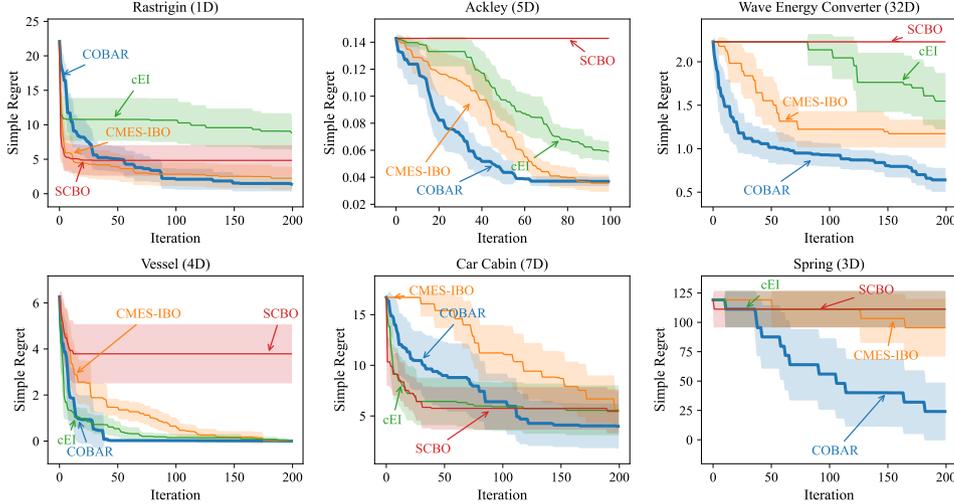


Figure 4: The input dimensionality, the number of constraints, and the approximate portion of the feasible region in the whole search space for each task are denoted on the titles. We run the algorithms on each task for at least 15 independent trials. The curves show the average simple regret after standardization, while the shaded area denotes the 95% confidence interval through the optimization.

a feasible area of 14% of the search space, which we construct by sampling $|\bar{D}| = 20000$ and reuse for all 15 trials. We also include another experiment on continuous search space of Ackley-10D-2C from Eriksson and Poloczek [2021] in Appendix E addressing boundary optimum.

- The *pressure vessel design problem* aims at optimizing the total cost of a cylindrical pressure vessel. The feasible regions take up around 78% of the whole search space.
- The *coil compression spring design problem* aims to optimize the volume of spring steel wire, which is used to manufacture the spring [Lampinen and Zelinka, 1999] under static loading. The feasible regions take up approximately 0.38% of the whole search space.
- The *car cab design problem* includes seven input variables and eight constraints. The feasible region takes up approximately 13% of the whole search space.
- This *UCI water converter problem* consists of positions and absorbed power outputs of wave energy converters (WECs) from the southern coast of Sydney [mis, 2019]. The feasible region takes up approximately 27% of the whole search space.

6.2 RESULTS

We study the robustness of the algorithms with varying feasible region sizes on the Rastrigin-1D-1C task. Results are demonstrated in figure 3. Note that the discrete search space consists of the 1000 points shown in the first row of figure 3, and with the observation noises, only COBAR consistently reaches the global optimum within 2000 iterations. The convergence highlights the essential role of the active learning of the constraint in achieving robust optimization when unknown constraints are present.

We further study COBAR on the aforementioned optimization tasks, with simple regret curves shown in figure 4. On the Rastrigin-1D-1C and Car-Cabin-7D-8C tasks, COBAR initially lags behind the baselines. This is likely because the computational demands of actively learning the constraints temporarily hinder optimization progress. However, the steady improvement of COBAR leads to consistently superior performance after sufficient iterations, whereas the baselines become trapped in local optima. These results demonstrate that COBAR is efficient and effective across various input dimensionalities, constraint numbers, and constraint correlations. For further details, Appendix H includes Table 4, which presents the mean simple regret at specific budget points, and Table 5, which summarizes performance ranks within 100 iterations. A comparison of wall-clock times is provided in table 3.

7 CONCLUSION

Bayesian optimization with unknown constraints poses challenges in the adaptive tradeoff between optimizing the unknown objective and learning the constraints. We introduce COBAR, which is backed by rigorous theoretical guarantees, to efficiently address constrained Bayesian optimization. Our key insights include: (1) the ROIs determined through adaptive level-set estimation can congregate and contribute to the overall Bayesian optimization task; (2) acquisition functions based on independent GPs can be unified in a principled way. Through extensive experiments, we validate the efficacy and robustness of our proposed method across various tasks finding the interior optimum.

ACKNOWLEDGMENT

This work was supported in part by the National Science Foundation under Grant No. IIS 2313131, IIS 2332475, and CMMI 2037026. The authors acknowledge the University of Chicago’s Research Computing Center for their support of this work. We would also like to extend our sincere gratitude to the NSF-FMRG project team for their insightful discussions. We also thank Thomas Desautels, Yisong Yue, Christopher Yeh, Ladd Alexander, Raul Astudillo, Jialin Song, James Bowden, and the anonymous reviewers for their engagement and feedback.

References

- Wave Energy Converters. UCI Machine Learning Repository, 2019. DOI: <https://doi.org/10.24432/C5QS4V>.
- Candelieri Antonio. Sequential model based optimization of partially defined functions under unknown constraints. *Journal of Global Optimization*, 79(2):281–303, 2021.
- Setareh Ariafar, Jaume Coll-Font, Dana H Brooks, and Jennifer G Dy. Admmbo: Bayesian optimization with unknown constraints using admm. *J. Mach. Learn. Res.*, 20(123):1–26, 2019.
- François Bachoc, Céline Helbert, and Victor Picheny. Gaussian process optimization with failures: classification and convergence proof. *Journal of Global Optimization*, 78(3):483–506, 2020.
- Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. Botorch: A framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33:21524–21538, 2020.
- Julien Bect, David Ginsbourger, Ling Li, Victor Picheny, and Emmanuel Vazquez. Sequential design of computer experiments for the estimation of a probability of failure. *Statistics and Computing*, 22:773–793, 2012.
- Felix Berkenkamp, Angela P Schoellig, and Andreas Krause. No-regret bayesian optimization with unknown hyperparameters. *Journal of Machine Learning Research*, 20(50):1–24, 2019.
- J Bernardo, MJ Bayarri, JO Berger, AP Dawid, D Heckerman, AFM Smith, and M West. Optimization under unknown constraints. *Bayesian Statistics*, 9(9):229, 2011.
- Mickael Binois and Nathan WycOFF. A survey on high-dimensional gaussian process modeling with application to bayesian optimization. *ACM Transactions on Evolutionary Learning and Optimization*, 2(2):1–26, 2022.
- Ilija Bogunovic, Jonathan Scarlett, Andreas Krause, and Volkan Cevher. Truncated variance reduction: A unified approach to bayesian optimization and level-set estimation. *Advances in neural information processing systems*, 29, 2016.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pages 844–853. PMLR, 2017.
- Sanjoy Dasgupta. Learning mixtures of gaussians. In *40th Annual Symposium on Foundations of Computer Science (Cat. No. 99CB37039)*, pages 634–644. IEEE, 1999.
- Kalyanmoy Deb and Himanshu Jain. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part i: solving problems with box constraints. *IEEE transactions on evolutionary computation*, 18(4):577–601, 2013.
- Sébastien Le Digabel and Stefan M. Wild. A taxonomy of constraints in simulation-based optimization, 2015. URL <https://arxiv.org/abs/1505.07881>.
- VI Donskoi. Partially defined optimization problems: An approach to a solution that is based on pattern recognition theory. *Journal of Soviet Mathematics*, 65:1664–1668, 1993.
- David Eriksson and Martin Jankowiak. High-dimensional bayesian optimization with sparse axis-aligned subspaces. In *Uncertainty in Artificial Intelligence*, pages 493–503. PMLR, 2021.
- David Eriksson and Matthias Poloczek. Scalable constrained bayesian optimization. In *International Conference on Artificial Intelligence and Statistics*, pages 730–738. PMLR, 2021.
- David Eriksson, Michael Pearce, Jacob Gardner, Ryan D Turner, and Matthias Poloczek. Scalable global optimization via local bayesian optimization. *Advances in neural information processing systems*, 32, 2019.
- Paul Feliot, Julien Bect, and Emmanuel Vazquez. A bayesian approach to constrained single-and multi-objective optimization. *Journal of Global Optimization*, 67(1-2):97–133, 2017.
- Peter I Frazier. A tutorial on bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.
- Jacob R Gardner, Matt J Kusner, Zhixiang Eddie Xu, Kilian Q Weinberger, and John P Cunningham. Bayesian optimization with inequality constraints. In *ICML*, volume 2014, pages 937–945, 2014.
- Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2023.

- Michael A Gelbart, Jasper Snoek, and Ryan P Adams. Bayesian optimization with unknown constraints. *arXiv preprint arXiv:1403.5607*, 2014.
- Alkis Gotovos, Nathalie Casati, Gregory Hitz, and Andreas Krause. Active learning for level set estimation. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 1344–1350, 2013.
- Robert B Gramacy. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. Chapman and Hall/CRC, 2020.
- Robert B Gramacy, Genetha A Gray, Sébastien Le Digabel, Herbert KH Lee, Pritam Ranjan, Garth Wells, and Stefan M Wild. Modeling an augmented lagrangian for blackbox constrained optimization. *Technometrics*, 58(1): 1–11, 2016.
- Hengquan Guo, Zhu Qi, and Xin Liu. Rectified pessimistic-optimistic learning for stochastic continuum-armed bandit with constraints. In Nikolai Matni, Manfred Morari, and George J. Pappas, editors, *Proceedings of The 5th Annual Learning for Dynamics and Control Conference*, volume 211 of *Proceedings of Machine Learning Research*, pages 1333–1344. PMLR, 15–16 Jun 2023. URL <https://proceedings.mlr.press/v211/guo23a.html>.
- Erik Orm Hellsten, Carl Hvarfner, Leonard Papenmeier, and Luigi Nardi. High-dimensional bayesian optimization with group testing. *arXiv preprint arXiv:2310.03515*, 2023.
- José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. *Advances in neural information processing systems*, 27, 2014.
- José Miguel Hernández-Lobato, Michael Gelbart, Matthew Hoffman, Ryan Adams, and Zoubin Ghahramani. Predictive entropy search for bayesian optimization with unknown constraints. In *International conference on machine learning*, pages 1699–1707. PMLR, 2015.
- Carl Hvarfner, Erik Hellsten, Frank Hutter, and Luigi Nardi. Self-correcting bayesian optimization through bayesian active learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- BK Kannan and SN Kramer. An augmented lagrange multiplier based method for mixed integer discrete continuous optimization and its applications to mechanical design. *Journal of Mechanical Design*, 116(2):405–411, 1994.
- Peter Koepf and Florian Pfaff. Consistency of gaussian process regression in metric spaces. *The Journal of Machine Learning Research*, 22(1):11066–11092, 2021.
- Junpei Komiyama, Gustavo Malkomes, Bolong Cheng, and Michael McCourt. Bridging offline and online experimentation: Constraint active search for deployed performance optimization. *Transactions on Machine Learning Research*, 2022.
- Jouni Lampinen and Ivan Zelinka. Mixed integer-discrete-continuous optimization by differential evolution. In *Proceedings of the 5th international conference on soft computing*, pages 71–76. Citeseer, 1999.
- Ben Letham, Roberto Calandra, Akshara Rai, and Eytan Bakshy. Re-examining linear embeddings for high-dimensional bayesian optimization. *Advances in neural information processing systems*, 33:1546–1558, 2020.
- Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained bayesian optimization with noisy experiments. *Bayesian Analysis*, 14(2):495–519, 2019.
- Congwen Lu and Joel A Paulson. No-regret bayesian optimization with unknown equality and inequality constraints using exact penalty functions. *IFAC-PapersOnLine*, 55(7):895–902, 2022.
- Congwen Lu and Joel A Paulson. No-regret constrained bayesian optimization of noisy and expensive hybrid models using differentiable quantile function approximations. *arXiv preprint arXiv:2305.03824*, 2023.
- David JC MacKay. Information-based objective functions for active data selection. *Neural computation*, 4(4):590–604, 1992.
- Gustavo Malkomes, Bolong Cheng, Eric H Lee, and Mike Mccourt. Beyond the pareto efficient frontier: Constraint active search for multiobjective experimental design. In *International Conference on Machine Learning*, pages 7423–7434. PMLR, 2021.
- Amin Nayebi, Alexander Munteanu, and Matthias Poloczek. A framework for bayesian optimization in embedded subspaces. In *International Conference on Machine Learning*, pages 4752–4761. PMLR, 2019.
- Quoc Phong Nguyen, Bryan Kian Hsiang Low, and Patrick Jaillet. An information-theoretic framework for unifying active learning problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9126–9134, 2021.
- Valerio Perrone, Iaroslav Shcherbatyi, Rodolphe Jenatton, Cedric Archambeau, and Matthias Seeger. Constrained bayesian optimization with max-value entropy search. *arXiv preprint arXiv:1910.07003*, 2019.
- Victor Picheny, Robert B Gramacy, Stefan Wild, and Sébastien Le Digabel. Bayesian optimization under mixed

- constraints with a slack-variable augmented lagrangian. *Advances in neural information processing systems*, 29, 2016.
- H Pohlheim. Geatbx examples examples of objective functions. URL http://www.geatbx.com/download/GEATbx_ObjFunExpl_v37.pdf, 2006.
- C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Leonard Andreevič Rastrigin. Systems of extremal control. *Nauka*, 1974.
- LI Rudenko. Objective functional approximation in a partially defined optimization problem. *Journal of Mathematical Sciences*, 72:3359–3363, 1994.
- Matthieu Sacher, Régis Duvigneau, Olivier Le Maitre, Mathieu Durand, Elisa Berrini, Frédéric Hauville, and Jacques-André Astolfi. A classification approach to efficient global optimization in presence of non-computable domains. *Structural and Multidisciplinary Optimization*, 58:1537–1557, 2018.
- Matthias Schonlau, William J Welch, and Donald R Jones. Global versus local search in constrained optimization of computer models. *Lecture notes-monograph series*, pages 11–25, 1998.
- Yaroslav D Sergeyev, Dmitri E Kvasov, and Falah MH Khalaf. A one-dimensional local tuning algorithm for solving go problems with partially defined constraints. *Optimization Letters*, 1(1):85–99, 2007.
- Niranjan Srinivas, Andreas Krause, Sham M Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.
- Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. In *International conference on machine learning*, pages 997–1005. PMLR, 2015.
- Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Sequential and parallel constrained max-value entropy search via information lower bound. In *International Conference on Machine Learning*, pages 20960–20986. PMLR, 2022.
- Ryoji Tanabe and Hisao Ishibuchi. An easy-to-use real-world multi-objective optimization problem suite. *Applied Soft Computing*, 89:106078, 2020.
- Jingyi Wang, Cosmin G Petra, and J Luc Peterson. Constrained bayesian optimization with merit functions. *arXiv preprint arXiv:2403.13140*, 2024.
- Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient bayesian optimization. In *International Conference on Machine Learning*, pages 3627–3635. PMLR, 2017.
- Ziyu Wang, Frank Hutter, Masrour Zoghi, David Matheson, and Nando De Freitas. Bayesian optimization in a billion dimensions via random embeddings. *Journal of Artificial Intelligence Research*, 55:361–387, 2016.
- Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P Xing. Deep kernel learning. In *Artificial intelligence and statistics*, pages 370–378. PMLR, 2016.
- Wenjie Xu, Yuning Jiang, Bratislav Svetozarevic, and Colin Jones. Constrained efficient global optimization of expensive black-box functions. In *International Conference on Machine Learning*, pages 38485–38498. PMLR, 2023.
- Zeji Yi, Yunyue Wei, Chu Xin Cheng, Kaibo He, and Yanan Sui. Improving sample efficiency of high dimensional bayesian optimization with mcmc. *arXiv preprint arXiv:2401.02650*, 2024.
- Fengxue Zhang, Jialin Song, James Bowden, Alexander Ladd, Yisong Yue, Thomas A. Desautels, and Yuxin Chen. Learning regions of interest for bayesian optimization with adaptive level-set estimation, 2023.
- Yunxiang Zhang, Xiangyu Zhang, and Peter Frazier. Constrained two-step look-ahead bayesian optimization. *Advances in Neural Information Processing Systems*, 34:12563–12575, 2021.
- Xingyu Zhou and Bo Ji. On kernelized multi-armed bandits with constraints. *Advances in Neural Information Processing Systems*, 35:14–26, 2022.

A IMPACT STATEMENTS

Bayesian optimization with unknown constraints has emerged as a powerful tool in diverse fields, including scientific experimental design and engineering optimization tasks. This new algorithm provides a principled solution for this challenging problem, offering several potential benefits, including improved efficiency and enhanced robustness. No major ethical concerns are anticipated, given the algorithm’s generality and focus on solving practical problems.

B PROOFS

B.1 PROOF OF LEMMA 1

Lemma 1. Under the assumptions above, the regions of interest $\hat{\mathbf{X}}_t$, as defined in equation 4, contain the global optimum with high probability. Formally, for all $\delta \in (0, 1)$, $T \geq t \geq 1$, and any finite discretization \tilde{D} of \mathbf{X} that contains the optimum $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x})$ where $\mathcal{C}_m(\mathbf{x}^*) > \epsilon_c$ for all $m \in \mathbf{M}$ and $\beta_t = 2 \log(2(M+1)|\tilde{D}|/\pi_t/\delta)$ with $\sum_{t \geq 1}^T \pi_t^{-1} = 1$, we have $\mathbb{P}[\mathbf{x}^* \in \tilde{D}_{\hat{\mathbf{X}}_t}] \geq 1 - \delta$.

Proof. With probability at least $1 - 1/2\delta$, $\forall \mathbf{x} \in \tilde{D}, \forall T \geq t \geq 1, \forall g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$,

$$|g(\mathbf{x}) - \mu_{g,t-1}(\mathbf{x})| \leq \beta_t^{1/2} \sigma_{g,t-1}(\mathbf{x})$$

Note that we also take the union bound on $g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$.

This is similarly derived as lemma 5.1 of Srinivas et al. [2009] or lemma 1 of Zhang et al. [2023]. *Different from previous proofs*, we do not require the lemma to hold for $\forall t \geq 1$. Instead, we require it to hold for $\forall T \geq t \geq 1$. This alleviates the need of the convergence of the series $\sum_{t \geq 1} \pi_t^{-1} = 1$ to $\sum_{t \geq 1}^T \pi_t^{-1} = 1$ when taking the union bound. Specifically, we could set $\pi_t = T$, which essentially makes $\beta_t = 2 \log(\frac{2(M+1)|\tilde{D}|}{\delta})$ a constant. Hence, we use the β in the following instead of β_t as traditionally used to highlight this difference.

First, by definition $S_{\mathcal{C},t} \triangleq \bigcap_m^{\mathbf{M}} S_{\mathcal{C}_m,t}$, we have $\forall t \leq T, \mathbf{x} \in \tilde{D} \cap S_{\mathcal{C},t}, \forall m \in \mathbf{M}$

$$\mathbb{P}[\mathcal{C}_m(\mathbf{x}) \geq \text{LCB}_{\mathcal{C}_m,t}(\mathbf{x}) > 0] \geq 1 - 1/2\delta$$

meaning with probability at $1 - \delta$, \mathbf{x} lies in the feasible region. At the same time, we have, $\forall t \leq T, \forall m \in \mathbf{M}$, given $\mathcal{C}_m(\mathbf{x}) > 0$

$$\mathbb{P}[\text{UCB}_{f,t}(\mathbf{x}^*) \geq f(\mathbf{x}^*) \geq f(\mathbf{x}) \geq \text{LCB}_{f,t}(\mathbf{x})] \geq 1 - 1/2\delta$$

Given the mutual independency between the objective f and the constraints \mathcal{C}_m , and by the definition of the threshold $\text{LCB}_{f,t,\max}$, we have $\forall t \leq T$, when $\exists \mathbf{x} \in \tilde{D} \cap S_{\mathcal{C},t}$,

$$\mathbb{P}[\text{UCB}_{f,t}(\mathbf{x}^*) > \text{LCB}_{f,t,\max}] \geq 1 - \delta$$

Note when $\tilde{D} \cap S_{\mathcal{C},t} = \emptyset$, $\text{LCB}_{f,t,\max} = -\infty$, we have $\mathbb{P}[\text{UCB}_{f,t}(\mathbf{x}^*) > \text{LCB}_{f,t,\max}] = 1$.

In summary, we’ve shown that with probability at least $1 - \delta$, $\mathbf{x}^* \in \tilde{D} \cap \hat{\mathbf{X}}_{f,t}$.

Next, by the definition of $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x})$ s.t. $\mathcal{C}_m(\mathbf{x}^*) > \epsilon_c$ we have $\forall t \leq T, \forall m \in \mathbf{M}$

$$\mathbb{P}[\text{UCB}_{\mathcal{C}_m,t}(\mathbf{x}^*) \geq \mathcal{C}_m(\mathbf{x}^*) > 0] \geq 1 - 1/2\delta$$

meaning with probability at least $1 - 1/2\delta$, $\mathbf{x}^* \in \tilde{D} \cap \hat{\mathbf{X}}_{\mathcal{C}_m,t}$. And in general, we have $\forall t \leq T, \forall m \in \mathbf{M}$

$$\mathbb{P}[\mathbf{x}^* \in \tilde{D} \cap \hat{\mathbf{X}}_t] \geq 1 - \delta$$

□

B.2 PROOF OF THEOREM 1

The following lemmas show that the maximum of the acquisition functions equation 5 and 6 are both bounded after sufficient evaluations.

Lemma 1. *Under the conditions assumed in Theorem 1 except for Assumption 2, let $\alpha_t = \max_{g \in \mathcal{G}} \alpha_{g,t}(\mathbf{x}_{g,t})$ as in Algorithm 1, with $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ that is a constant, after at most $T \geq \frac{\beta \widehat{\gamma}_T C_1}{\epsilon_f^2}$ iterations, $\alpha_T \leq \epsilon_f$. Here $C_1 = 8/\log(1 + \sigma^{-2})$.*

The inequation $T \geq \frac{\beta \widehat{\gamma}_T C_1}{\epsilon_f^2}$ has T on both side, which follows the convention in Gotovos et al. [2013].

Proof. We first unify the notation in the acquisition functions.

$\forall T \geq t \geq 1, \forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}$, when $\tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t} \neq \emptyset$,

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t}} \text{UCB}_{g,t}(\mathbf{x}) - \text{LCB}_{g,t}(\mathbf{x}) \leq \alpha_t \quad (8)$$

$\forall T \geq t \geq 1, \forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}$, when $\tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{\mathcal{C}_m,t} = \emptyset$, let

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t}} \text{UCB}_{g,t}(\mathbf{x}) - \text{LCB}_{g,t}(\mathbf{x}) = 0 \leq \alpha_t \quad (9)$$

$\forall T \geq t \geq 1, g = f$, when $S_{\mathcal{C},t} = \emptyset$, we have

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}} \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t}(\mathbf{x}) \leq \alpha_t \quad (10)$$

$\forall T \geq t \geq 1, g = f$, when $S_{\mathcal{C},t} \neq \emptyset$, we have

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}} \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t,\max} \leq \alpha_t \quad (11)$$

By lemma 5.1, 5.2 and 5.4 of Srinivas et al. [2009], with $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$, $\forall g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$ and $\forall x_t \in \tilde{D}_{\hat{\mathbf{x}}_t} \subseteq \tilde{D}$, we have $\sum_{t=1}^T (2\beta^{1/2} \sigma_{g,t-1}(\mathbf{x}_t))^2 \leq C_1 \beta \gamma_{g,T}$. By definition of α_t , we have the following

$$\begin{aligned} \sum_{t=1}^T \alpha_t^2 &\leq \sum_{t=1}^T \max_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}} (2\beta^{1/2} \sigma_{g,t-1}(\mathbf{x}_{g,t}))^2 \\ &\leq \sum_{t=1}^T \sum_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}} (2\beta^{1/2} \sigma_{g,t-1}(\mathbf{x}_t))^2 \\ &\leq \sum_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}} C_1 \beta \gamma_{g,T} \\ &= C_1 \beta \widehat{\gamma}_T \end{aligned}$$

The last line holds due to the definition in equation 7. By Cauchy-Schwarz, we have

$$\frac{1}{T} \left(\sum_{t=1}^T \alpha_t \right)^2 \leq C_1 \beta \widehat{\gamma}_T$$

With Assumption 3, $\forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}, \forall 1 \leq t_1 < t_2 \leq T, \forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}$, we have $U_{g,t_2} \subseteq U_{g,t_1}$ and $\hat{\mathbf{x}}_{t_2} \subseteq \hat{\mathbf{x}}_{t_1}$, and most importantly, $\alpha_{t_2} \leq \alpha_{t_1}$. Therefore

$$\alpha_T \leq \frac{1}{T} \sum_{t=1}^T \alpha_t \leq \sqrt{\frac{C_1 \beta \widehat{\gamma}_T}{T}}$$

As a result, after at most $T \geq \frac{\beta \widehat{\gamma}_T C_1}{\epsilon_f^2}$ iterations, we have $\alpha_T \leq \epsilon_f$. \square

With Lemma 1, we could first prove that after adequately T rounds of evaluations such that $\epsilon_f \leq \min_{m \in \mathbf{M}} \epsilon_m$ is sufficiently small, with certain probability, $\mathbf{x}^* \in S_{C,T}$. Then $\text{LCB}_{f,t,\max} \neq -\infty$, and therefore the width of $[\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}} \text{LCB}_{f,T}(\mathbf{x}), \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}} \text{UCB}_{f,T}(\mathbf{x})]$, which is a the high confidence interval of f^* , is bounded by ϵ_f .

Proof. We first prove that after at most $T \geq \frac{\beta \hat{\gamma}_T C_1}{\epsilon_f^2}$ iterations, $\mathbb{P}[\mathbf{x}^* \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap S_{C,T}] \geq 1 - 1/2\delta$. Given equation 8 and 9 and Lemma 1, we have $\forall g \in \{C_m\}_{m \in \mathbf{M}}$,

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap U_{g,T}} \text{UCB}_{g,T}(\mathbf{x}) - \text{LCB}_{g,T}(\mathbf{x}) \leq \epsilon_f \leq \min_{m \in \mathbf{M}} \epsilon_m$$

According to the definition of $U_{g,T}$, $\forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap U_{g,T}$, $\forall g \in \{C_m\}_{m \in \mathbf{M}}$

$$\text{UCB}_{g,T}(\mathbf{x}) \leq \epsilon_f + \text{LCB}_{g,T}(\mathbf{x}) \leq \epsilon_f \leq \min_{m \in \mathbf{M}} \epsilon_m$$

According to Assumption 2, and Lemma 1, $\forall m \in \mathbf{M}$, we have

$$\mathbb{P}\left[\text{UCB}_{C_m,T}(\mathbf{x}^*) > \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{C_m,t}} \text{UCB}_{C_m,T}(\mathbf{x})\right] \geq 1 - 1/2\delta$$

Given $\tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{C,T} = \tilde{D}_{\hat{\mathbf{x}}_t} \cap \hat{\mathbf{X}}_{C,T} \setminus \cup_{m \in \mathbf{M}} U_{C_m,T}$, when $t = T$, we have

$$\mathbb{P}[\mathbf{x}^* \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{C,T}] \geq 1 - 1/2\delta \quad (12)$$

As a result

$$\mathbb{P}[\text{LCB}_{f,T,\max} \neq -\infty] \geq 1 - 1/2\delta$$

Next, we prove the upper bound for the width of the high-confidence interval of f^* . Given that $\text{LCB}_{f,T,\max} \neq -\infty$, we have

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{C,T}} \text{UCB}_{f,T}(\mathbf{x}) - \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{C,T}} \text{LCB}_{f,T}(\mathbf{x}) \leq \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{C,T}} \text{UCB}_{f,T}(\mathbf{x}) - \text{LCB}_{f,T,\max} \leq \alpha_T \leq \epsilon_f$$

Combining it with the observation that with probability $1 - 1/2\delta$,

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{C,T}} \text{LCB}_{f,T}(\mathbf{x}) < f(\mathbf{x}^*) \leq \max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{C,T}} \text{UCB}_{f,T}(\mathbf{x})$$

we attain the final result that after $T \geq \frac{\beta \hat{\gamma}_T C_1}{\epsilon_f^2}$ iterations,

$$\mathbb{P}[|CI_{f^*,T}| \leq \epsilon, f^* \in CI_{f^*,T}] \geq 1 - \delta$$

□

B.3 PROOF OF COROLLARY 1

Proof. We simply need to show that after $t \geq \frac{\beta \hat{\gamma}_t C_1}{\epsilon_f^2}$ iterations, with probability at least $1 - \delta$, \mathbf{x}^* is the only member in $\tilde{D}_{\hat{\mathbf{x}}_t}$.

Similar to Theorem 1, we have $\mathbb{P}[|CI_{f^*,t}| \leq \epsilon_f, f^* \in CI_{f^*,t}] \geq 1 - \delta$. At the same time, given the proof of Lemma 1, we have $\forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}$, $2\beta^{1/2}\sigma_{f,t-1}(\mathbf{x}) \leq \epsilon_f$.

Then if $\exists \mathbf{x} \neq \mathbf{x}^*$ and $\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}$, we have $f^* - f(\mathbf{x}) > 2\epsilon_f$, while

$$\mathbb{P}\left[f^* - f(\mathbf{x}) \leq |CI_{f^*,t}| + \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t}(\mathbf{x}) \leq 2\beta^{1/2}\sigma_{f,t-1}(\mathbf{x}) + \epsilon_f \leq 2\epsilon_f\right] \geq 1 - \delta$$

This contradiction means with probability at least $1 - \delta$, \mathbf{x}^* is the only member in $\tilde{D}_{\hat{\mathbf{x}}_t}$, and $\mathbf{x}_t = \mathbf{x}^*$. As a result, $\mathbb{P}[\mathbf{R}_t = 0] \geq 1 - \delta$, when $T \geq t \geq \frac{\beta \hat{\gamma}_t C_1}{\epsilon_f^2}$. □

B.4 PROOF OF COROLLARY 2

Proof. We follow the same path as the proof of Corollary 1.

Similar to Theorem 1, we have $\mathbb{P}[|CI_{f^*,t}| \leq \alpha_t \leq \epsilon_f, f^* \in CI_{f^*,t}] \geq 1 - \delta$. At the same time, given the proof of Lemma 1, we have $\forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}, 2\beta^{1/2}\sigma_{f,t-1}(\mathbf{x}) \leq \alpha_t \leq \epsilon_f$.

Then $\forall \mathbf{x} \neq \mathbf{x}^*$ and $\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}$, we have

$$\mathbb{P}[f^* - f(\mathbf{x}) \leq |CI_{f^*,t}| + \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t}(\mathbf{x}) \leq 2\alpha_t \leq 2\epsilon_f] \geq 1 - \delta$$

Then by assumption, $\forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}, \forall m \in \mathbf{M}$, we have probability at least $1 - \delta$, $\mathcal{C}_m(\mathbf{x}) \geq \epsilon_c$, and hence $\mathbf{x} \notin U_{\mathcal{C}_m,t}$. According to the algorithm, it regresses to GP-UCB by Srinivas et al. [2009] between t' and T .

$$\begin{aligned} \sum_{t=t'}^T (r(\mathbf{x}^*) - r(\mathbf{x}_t))^2 &\leq \beta C_1 (\gamma_T - \gamma_{t'}) \\ &\leq \beta C_1 \gamma_T (1 - t'/T) \end{aligned}$$

By Cauchy-Schwarz, we have

$$\begin{aligned} \sum_{t=t'}^T (r(\mathbf{x}^*) - r(\mathbf{x}_t)) &\leq \sqrt{(T - t') \sum_{t=t'}^T (r(\mathbf{x}^*) - r(\mathbf{x}_t))^2} \\ &\leq \sqrt{\frac{(T - t')^2}{T} \beta C_1 \gamma_T} \\ &\leq \sqrt{(T - t') \beta C_1 \gamma_T} \end{aligned}$$

□

B.5 PROOF OF COROLLARY 3

Proof. We assume $\tilde{D}_{\hat{\mathbf{x}}_T} \neq \emptyset$ and prove by contradiction. Given equation 8 and 9 and Lemma 1, we have $\forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}$,

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap U_{g,T}} \text{UCB}_{g,T}(\mathbf{x}) - \text{LCB}_{g,T}(\mathbf{x}) \leq \epsilon'_c$$

According to the definition of $U_{g,T}, \forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap U_{g,T}, \forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}$, with probability at least $1 - 1/2\delta$, we have

$$\mathcal{C}_m(\mathbf{x}) \leq \text{UCB}_{\mathcal{C}_m,T}(\mathbf{x}) \leq \epsilon'_c + \text{LCB}_{g,T}(\mathbf{x}) \leq \epsilon'_c + \mathcal{C}_m(\mathbf{x})$$

Then we have $\forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap U_{g,T}, \exists m \in \mathbf{M}$

$$\mathbb{P}[\mathcal{C}_m(\mathbf{x}) \leq \epsilon'_c + \mathcal{C}_m(\mathbf{x}) < 0] \geq 1 - 1/2\delta$$

This contradiction means $\forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}, \tilde{D}_{\hat{\mathbf{x}}_T} \cap U_{g,T} = \emptyset$ with probability as least $1 - 1/2\delta$.

According to the definition of $S_{g,T}, \forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{g,T}, \forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}$

$$\text{LCB}_{g,T}(\mathbf{x}) \geq \epsilon'_c$$

Then we have $\forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{g,T}, \exists g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}$

$$\mathbb{P}[-\epsilon'_c \geq \mathcal{C}_m(\mathbf{x}) \geq \text{LCB}_{g,T}(\mathbf{x}) \geq \epsilon'_c] \geq 1 - 1/2\delta$$

This contradiction means $\forall g \in \{\mathcal{C}_m\}_{m \in \mathbf{M}}, \tilde{D}_{\hat{\mathbf{x}}_T} \cap S_{g,T} = \emptyset$ with probability as least $1 - 1/2\delta$.

Combining the above contradictions, we have at least when $t = T$,

$$\mathbb{P}[\tilde{D}_{\hat{\mathbf{x}}_T} = \emptyset] \geq 1 - \delta$$

□

C CONTINUOUS SEARCH SPACE

C.1 THEORETICAL RESULTS

In the following, we introduce the additional assumption that bounds the unknown functions' complexity when they are members of an RKHS space and enables the performance analysis when applying COBAR on continuous search space \mathbf{X} instead of \tilde{D} .

Assumption 4. *The objective and constraints all lie in the RKHS \mathcal{H}_k corresponding to the kernel $k(\mathbf{x}, \mathbf{x}')$, and the corresponding norm is bounded by \mathcal{B} . Formally, $f : \mathbf{X} \rightarrow \mathbb{R}$ is a member of the RKHS of real-valued functions on \mathbf{X} with kernel k , with RKHS norm $\|f\|_k \leq \mathcal{B}$. Similarly, $\mathcal{C}_m : \mathbf{X} \rightarrow \mathbb{R}$ is a member of the RKHS of real-valued functions on \mathbf{X} with kernel k , with RKHS norm $\|\mathcal{C}_m\|_k \leq \mathcal{B}$, for all $m \in \mathbf{M}$.*

Then, we could derive similar results mapping from Lemma 1.

Lemma 2. Under the assumptions above, the regions of interest $\hat{\mathbf{X}}_t$, as defined in equation 4, contain the global optimum with high probability. Formally, for all $\delta \in (0, 1)$, $T \geq t \geq 1$, and the search space \mathbf{X} that contains the optimum $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x})$ where $\mathcal{C}_m(\mathbf{x}^*) > \epsilon_C$ for all $m \in \mathbf{M}$ and $\beta_t^{1/2} = B + \sigma \sqrt{2(\widehat{\gamma}_T + 1 + \ln(2(M+1)/\delta))}$, we have $\mathbb{P}[\mathbf{x}^* \in \hat{\mathbf{X}}_t] \geq 1 - \delta$.

Proof. Similar to theorem 2 of Chowdhury and Gopalan [2017], with probability at least $1 - 1/2\delta$, $\forall \mathbf{x} \in \tilde{D}, \forall T \geq t \geq 1, \forall g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$,

$$|g(\mathbf{x}) - \mu_{g,t-1}(\mathbf{x})| \leq \beta_t^{1/2} \sigma_{g,t-1}(\mathbf{x})$$

Note that we also take the union bound on $g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$.

First, by definition $S_{\mathcal{C},t} \triangleq \bigcap_m^{\mathbf{M}} S_{\mathcal{C}_m,t}$, we have $\forall t \leq T, \mathbf{x} \in S_{\mathcal{C},t}, \forall m \in \mathbf{M}$

$$\mathbb{P}[\mathcal{C}_m(\mathbf{x}) \geq \text{LCB}_{\mathcal{C}_m,t}(\mathbf{x}) > 0] \geq 1 - 1/2\delta$$

meaning with probability at $1 - \delta$, \mathbf{x} lies in the feasible region. At the same time, we have, $\forall t \leq T, \forall m \in \mathbf{M}$, given $\mathcal{C}_m(\mathbf{x}) > 0$

$$\mathbb{P}[\text{UCB}_{f,t}(\mathbf{x}^*) \geq f(\mathbf{x}^*) \geq f(\mathbf{x}) \geq \text{LCB}_{f,t}(\mathbf{x})] \geq 1 - 1/2\delta$$

Given the mutual independency between the objective f and the constraints \mathcal{C}_m , and by the definition of the threshold $\text{LCB}_{f,t,\max}$, we have $\forall t \leq T$, when $\exists \mathbf{x} \in S_{\mathcal{C},t}$,

$$\mathbb{P}[\text{UCB}_{f,t}(\mathbf{x}^*) > \text{LCB}_{f,t,\max}] \geq 1 - \delta$$

Note when $S_{\mathcal{C},t} = \emptyset$, $\text{LCB}_{f,t,\max} = -\infty$, we have $\mathbb{P}[\text{UCB}_{f,t}(\mathbf{x}^*) > \text{LCB}_{f,t,\max}] = 1$.

In summary, we've shown that with probability at least $1 - \delta$, $\mathbf{x}^* \in \hat{\mathbf{X}}_{f,t}$.

Next, by the definition of $\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x})$ s.t. $\mathcal{C}_m(\mathbf{x}^*) > \epsilon_C$ we have $\forall t \leq T, \forall m \in \mathbf{M}$

$$\mathbb{P}[\text{UCB}_{\mathcal{C}_m,t}(\mathbf{x}^*) \geq \mathcal{C}_m(\mathbf{x}^*) > 0] \geq 1 - 1/2\delta$$

meaning with probability at least $1 - 1/2\delta$, $\mathbf{x}^* \in \hat{\mathbf{X}}_{\mathcal{C}_m,t}$. And in general, we have $\forall t \leq T, \forall m \in \mathbf{M}$

$$\mathbb{P}[\mathbf{x}^* \in \hat{\mathbf{X}}_t] \geq 1 - \delta$$

□

Remark 2. The proof of Lemma 2 substitutes the β in the proof of Lemma 1 and alleviates the need for a discretization \tilde{D} with the additional assumption on the complexity of the unknown functions in Assumption 4. Note the $\beta_t^{1/2} = B + \sigma \sqrt{2(\widehat{\gamma}_T + 1 + \ln(2(M+1)/\delta))}$, we have $\mathbb{P}[\mathbf{x}^* \in \hat{\mathbf{X}}_t] \geq 1 - \delta$ is larger than the original value in the theorem 2 of Chowdhury and Gopalan [2017] to make sure β_t is the same for all $\forall T \geq t \geq 1$ and to guarantee a union bound on $g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$. In the following, since β_t is constant, we substitute it with β .

Then, we could trivially map the Theorem 1 when maximizing the acquisition functions on $\hat{\mathbf{X}}_t$ instead of $\tilde{D}_{\hat{\mathbf{X}}_t}$ as in line 9 of Algorithm 1 and on $\hat{\mathbf{X}}_t \cap U_{C_m,t}$ instead of $\tilde{D}_{\hat{\mathbf{X}}_t} \cap U_{C_m,t}$ as in line 8 of Algorithm 1. The proof would be identical to Appendix B except for the different β and search space.

Algorithm 2 COntstrained BO through Adaptive Region of Interest Acquisition on Continuous Space(COBAR-CS)

```

1: Input: Search space  $\mathbf{X}$ , initial observation  $\mathbf{S}_0$ , horizon  $T$ , confidence factor  $\delta$ , confidence coefficient  $\beta$ ;
2: for  $t = 1$  to  $T$  do
3:   Update the posteriors of  $\mathcal{GP}_{f,t}$  and  $\mathcal{GP}_{C_m,t}$  according to equation 1 and 2
4:   Identify ROIs  $\hat{\mathbf{X}}_t$ , and undecided sets  $U_{C_m,t}$ 
5:   for  $m \in \mathbf{M}$  do
6:     if  $U_{C_m,t} \neq \emptyset$  then
7:       Candidate for learning of each constraint:
7:        $\mathbf{x}_{C_m,t} \leftarrow \arg \max_{\mathbf{x} \in \hat{\mathbf{X}}_t \cap U_{C_m,t}} \alpha_{C_m,t}(\mathbf{x})$  (6)
8:        $\mathcal{G} \leftarrow \mathcal{G} \cup C_{m,t}$ 
9:       Candidate for optimizing the objective:
9:        $\mathbf{x}_{f,t} \leftarrow \arg \max_{\mathbf{x} \in \hat{\mathbf{X}}_t} \alpha_{f,t}(\mathbf{x})$  as in equation 5
10:       $\mathcal{G} \leftarrow \mathcal{G} \cup f$ 
11:      Maximize the acquisition from different aspects:
11:       $g_t \leftarrow \arg \max_{g \in \mathcal{G}} \alpha_{g,t}(\mathbf{x}_{g,t})$ 
12:      Pick the candidate to evaluate:  $\mathbf{x}_t \leftarrow \mathbf{x}_{g_t,t}$ 
13:      Update the observation set
13:       $\mathbf{S}_t \leftarrow \mathbf{S}_{t-1} \cup \{(\mathbf{x}_t, y_{f,t}, \{y_{C_m,t}\}_{m \in \mathbf{M}})\}$ 

```

Theorem 3. Under the aforementioned assumptions, with a constant $\beta_t^{1/2} \triangleq \beta^{1/2} = B + \sigma \sqrt{2(\widehat{\gamma}_T + 1 + \ln(2(M+1)/\delta))}$ and the acquisition function from Algorithm 2, there exists an $\epsilon_f \leq \epsilon_C$, such that after at most $T \geq \frac{\beta \widehat{\gamma}_T C_1}{\epsilon_f^2}$ iterations, we have $\mathbb{P}[|CI_{f^*,T}| \leq \epsilon_f, f^* \in CI_{f^*,T}] \geq 1 - \delta$ Here, $C_1 = 8/\log(1 + \sigma^{-2})$.

C.2 EFFICIENT DISCRETIZATION

Aiming at a continuous search space demands additional consideration when implementing a practical ROI identification on the continuous search space or requires a better coverage by the discretization of the dense search space for COBAR in practice. This problem is more outstanding in high-dimensional tasks. Here, we briefly discuss potential remedies if we still resort to an efficient discretization. The random linear projection has been used for discretizing the search space to mitigate the dependency on the dimensionality while, with high probability, preserving the original geometry [Dasgupta, 1999, Nayebi et al., 2019]. To efficiently discretize the dense search space for COBAR (in high-dimensional applications), one option is to apply the random projection and its reverse studied by Nayebi et al. [2019], which shows strong empirical performance when combined with other BO algorithms and offers the following theoretical guarantee.

Definition 1. (ϵ -subspace embedding [Nayebi et al., 2019]) Given a matrix $V \in \mathbb{R}^{D \times d}$ with orthonormal columns, an integer $d \leq D$ and an approximation parameter $\epsilon \in (0, 1)$, an ϵ -subspace embedding for V is a map $H : \mathbb{R}^d \rightarrow \mathbb{R}^D$ such that $\forall \mathbf{x} \in \mathbb{R}^d$:

$$(1 - \epsilon)\|V\mathbf{x}\|_2^2 \leq \|HV\mathbf{x}\|_2^2 \leq (1 + \epsilon)\|V\mathbf{x}\|_2^2$$

Theorem 4. (Theorem 2 of Nayebi et al. [2019]) Consider a Gaussian process that acts directly in the unknown active subspace of dimension d_e with mean and variance functions $\mu(\cdot), \sigma^2(\cdot)$. Let $\hat{\mu}(\cdot), \hat{\sigma}^2(\cdot)$ be their approximations using an ϵ -subspace embedding for the active subspace. Then we have for every $\mathbf{x} \in X$

1. $|\mu(\mathbf{x}) - \hat{\mu}(\mathbf{x})| \leq 5\epsilon\|\mathbf{x}\|\|X - \hat{f}\|$
2. $\sigma^2(\mathbf{x}) - \hat{\sigma}^2(\mathbf{x}) \leq 12\epsilon\|\mathbf{x}\|^2$

For a comprehensive survey on the treatments of high-dimensional search space for BO, we refer to the recent survey by Binois and Wycoff [2022]. Besides the random projection [Nayebi et al., 2019, Wang et al., 2016, Letham et al., 2020], variable selection Hellsten et al. [2023], tree-structure partition [Eriksson and Jankowiak, 2021], and Markov Chain Monte

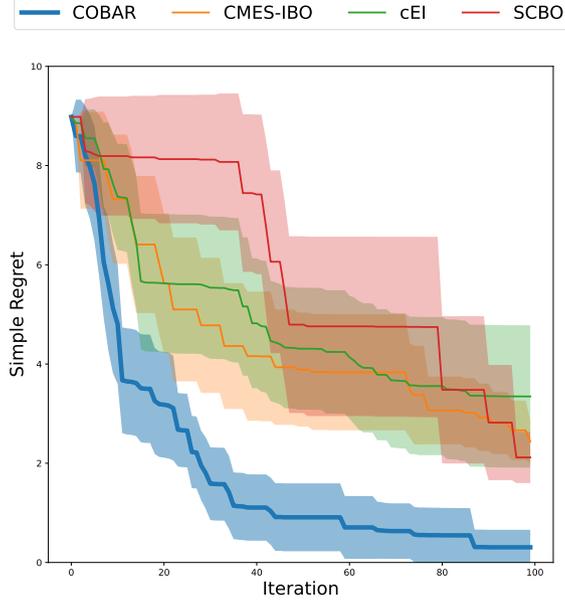


Figure 5: The figure illustrates the simple regret for Ackley-40D-2C. All the tested algorithms rely on the low-dimensional embedding in HeSBO. The results are collected from 15 independent trials. The shaded area denotes the standard error.

Problem	COBAR	CMES-IBO	SCBO	cEI
Ackley-40D-2C	75.87	281.47	34.33	185.07

Table 2: Average wall time (sec) of different CBO Methods collected from 15 independent trials.

Carlo sampling [Yi et al., 2024] on the search space could all be applied as plugins to improve the discretization efficiency for COBAR.

C.3 CASE STUDY

We illustrate the effectiveness of integrating Hashing-enhanced Subspace BO (HeSBO) [Nayebi et al., 2019] into CBO algorithms. We construct the following 40-dimensional CBO task that makes any grid discretization containing a feasible candidate on the original embedding intractable.

Ackley-40D-2C This Ackley-40D-2C function is a variant for the task we study in section 6. $f(\mathbf{x}) = 20 \exp(-0.2\sqrt{1/d \sum_i^d x_i^2}) + \exp(1/d \sum_i^d \cos(2\pi x_i)) + 20 + \exp(1)$, $d = 5$ where $\mathbf{x} \in [-5, 10]^{40}$. We construct two constraints to enforce a feasible area taking up less than 0.6% of the search space. The first constraint $\mathcal{C}_1 = 1 - (\sum_i^5 x_i)$. The second constraint $\mathcal{C}_2 = 6 - (\sum_i^5 x_i^2)$.

We find that COBAR fails in the original search space $[-5, 10]^{40}$ due to the intractability of any discretization containing sufficient feasible candidates. We integrate HeSBO into all the tested CBO algorithms to allow the algorithms to process on a 5-dimensional embedding space $[-1, 1]^5$. COBAR relies on a random sampling containing 200000 candidates. The simple regret curves are shown in figure 5. Though the point-wise comparison of COBAR is not tractable in the original 40-dimensional search space, integrating HeSBO allows COBAR to optimize the high-dimensional CBO toy problem efficiently. Table 2 shows that COBAR could efficiently optimize the embedding space, benefited from the reduced dimensionality and the ROI identification that further reduces computation need dynamically.

D DECOUPLED SETTING

In the main paper, we assume both objective f and the constraints $\{\mathcal{C}_m\}_{m \in \mathbf{M}}$ are revealed upon querying an input point. The setting is regarded as a coupling of the objective and constraints to differentiate from the decoupled setting, where the objective and constraints may be evaluated independently. In the decoupled setting, acquisition functions need to explicitly tradeoff the evaluation of the different aspects and, in addition to helping to pick the candidate $\mathbf{x}_t \in \mathbf{X}$, suggest $g_t \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$ for evaluation each time. This typically requires different acquisition from coupled setting [Gelbart et al., 2014]. However, we will show that our acquisition function and COBAR require minimum adaptation to the decoupled setting while bearing a similar performance guarantee.

D.1 ALGORITHM FOR DECOUPLED SETTING

When taking the $g_t \leftarrow \arg \max_{g \in \mathcal{G}} \alpha_{g,t}(\mathbf{x}_{g,t})$ in Algorithm 1, we explicitly choose the aspect that matters most at a certain iteration. Naturally, we could adapt COBAR to the decoupled setting by querying $\mathbf{x}_{g,t}$ on this unknown function $g_t \in \mathcal{G} \subseteq \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$ at iteration t . The modified algorithm is shown below.

Algorithm 3 Decoupled Constrained BO through Adaptive Region of interest Acquisition (COBAR-Decoupled)

- 1: **Input:** Search space \mathbf{X} , initial observation \mathbf{S}_0 , horizon T , confidence factor δ , estimated ϵ_C ;
 - 2: **for** $t = 1$ to T **do**
 - 3: Update the posteriors of $\mathcal{GP}_{f,t}$ and $\mathcal{GP}_{\mathcal{C}_m,t}$ according to equation 1 and 2
 - 4: Identify ROIs $\tilde{\mathbf{X}}_t$, and undecided sets $U_{\mathcal{C}_m,t}$
 - 5: **for** $m \in \mathbf{M}$ **do**
 - 6: **if** $U_{\mathcal{C}_m,t} \neq \emptyset$ **then**
 - 7: Candidate for active Learning of each constraint:
 $\mathbf{x}_{\mathcal{C}_m,t} \leftarrow \arg \max_{\mathbf{x} \in \tilde{D}_{\tilde{\mathbf{X}}_t} \cap U_{\mathcal{C}_m,t}} \alpha_{\mathcal{C}_m,t}(\mathbf{x})$ as in equation 6
 - 8: $\mathcal{G} \leftarrow \mathcal{G} \cup \mathcal{C}_{k,t}$
 - 9: Candidate for optimizing the objective:
 $\mathbf{x}_{f,t} \leftarrow \arg \max_{\mathbf{x} \in \tilde{D}_{\tilde{\mathbf{X}}_t}} \alpha_{f,t}(\mathbf{x})$ as in equation 5
 - 10: $\mathcal{G} \leftarrow \mathcal{G} \cup f$
 - 11: Maximize the acquisition values from different aspects:
 $g_t \leftarrow \arg \max_{g \in \mathcal{G}} \alpha_{g,t}(\mathbf{x}_{g,t})$
 - 12: Pick the candidate to evaluate: $\mathbf{x}_t \leftarrow \mathbf{x}_{g_t,t}$
 - 13: Update the observation set with the candidate and corresponding new observations on g_t
 $\mathbf{S}_t \leftarrow \mathbf{S}_{t-1} \cup \{(\mathbf{x}_t, y_{g_t,t})\}$
-

D.2 THEORETICAL GUARANTEE AND PROOF

We first denote the maximum mutual information gain after T rounds of evaluations as

$$\tilde{\gamma}_T = \sum_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}} \gamma_{g,T_g} \quad (13)$$

Where T_g denotes the number of evaluations for $g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}$ before T . Therefore we have

$$T = \sum_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathbf{M}}} T_g$$

Then, we have the following guarantee for the performance of COBAR-Decoupled.

Theorem 5. *The width of the resulting confidence interval of the global optimum $f^* = f(\mathbf{x}^*)$ has an upper bound. That is, under the same assumptions in Theorem 1, with $\beta = 2 \log(2(M+1)|\tilde{D}_{\tilde{\mathbf{X}}_t}| \pi_t / \delta)$ that is constant, and acquisition function in Algorithm 3, $\exists \epsilon_f \leq \epsilon_C$, after at most $T \geq \frac{\beta \tilde{\gamma}_T C_1}{\epsilon_f^2}$ iterations, we have $\mathbb{P}[|CI_{f^*,T}| \leq \epsilon_f, f^* \in CI_{f^*,T}] \geq 1 - \delta$ Here $C_1 = 8 / \log(1 + \sigma^{-2})$.*

Lemma 2. Under the conditions assumed in Theorem 5 except for Assumption 2, let $\alpha_t = \max_{g \in \mathcal{G}} \alpha_{g,t}(\mathbf{x}_{g,t})$ as in Algorithm 3, with $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ that is a constant, after at most $T \geq \frac{\beta \tilde{\gamma}_T C_1}{\epsilon_f^2}$ iterations, $\alpha_T \leq \epsilon_f$. Here $C_1 = 8/\log(1 + \sigma^{-2})$.

Here is the critical difference to the proof of Theorem 1.

Proof. We first unify the notation in the acquisition functions.

$\forall T \geq t \geq 1, \forall g \in \{\mathcal{C}_m\}_{m \in \mathcal{M}}$, when $\tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t} \neq \emptyset$,

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t}} \text{UCB}_{g,t}(\mathbf{x}) - \text{LCB}_{g,t}(\mathbf{x}) \leq \alpha_t \quad (14)$$

$\forall T \geq t \geq 1, \forall g \in \{\mathcal{C}_m\}_{m \in \mathcal{M}}$, when $\tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{\mathcal{C}_m,t} = \emptyset$, let

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t}} \text{UCB}_{g,t}(\mathbf{x}) - \text{LCB}_{g,t}(\mathbf{x}) = 0 \leq \alpha_t \quad (15)$$

$\forall T \geq t \geq 1, g = f$, when $S_{\mathcal{C},t} = \emptyset$, we have

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}} \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t}(\mathbf{x}) \leq \alpha_t \quad (16)$$

$\forall T \geq t \geq 1, g = f$, when $S_{\mathcal{C},t} \neq \emptyset$, we have

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t}} \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t,\max} \leq \alpha_t \quad (17)$$

By lemma 5.1, 5.2 and 5.4 of Srinivas et al. [2009], with $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$, $\forall g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathcal{M}}$ and $\forall \mathbf{x}_t \in \tilde{D}_{\hat{\mathbf{x}}_t} \subseteq \tilde{D}$, we have $\sum_{t=1}^T (2\beta^{1/2} \sigma_{g,t-1}(\mathbf{x}_t))^2 \mathbf{1}(g_t = g) \leq C_1 \beta \gamma_{g,T_g}$. By definition of α_t , we have the following

$$\begin{aligned} \sum_{t=1}^T \alpha_t^2 &= \sum_{t=1}^T \max_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathcal{M}}} \alpha_{g_t,t}^2(\mathbf{x}_{g_t,t}) \\ &\leq \sum_{t=1}^T \max_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathcal{M}}} (2\beta^{1/2} \sigma_{g_t,t-1}(\mathbf{x}_{g_t,t}))^2 \\ &\leq \sum_{g \in \{f\} \cup \{\mathcal{C}_m\}_{m \in \mathcal{M}}} C_1 \beta \gamma_{g,T_g} \\ &= C_1 \beta \tilde{\gamma}_T \end{aligned}$$

By Cauchy-Schwarz, we have

$$\frac{1}{T} \left(\sum_{t=1}^T \alpha_t \right)^2 \leq C_1 \beta \tilde{\gamma}_T$$

By the monotonicity assumed in Assumption 3, $\forall g \in \{\mathcal{C}_m\}_{m \in \mathcal{M}}, \forall 1 \leq t_1 < t_2 \leq T, \forall g \in \{\mathcal{C}_m\}_{m \in \mathcal{M}}$, we have $U_{g,t_2} \subseteq U_{g,t_1}$ and $\hat{\mathbf{X}}_{t_2} \subseteq \hat{\mathbf{X}}_{t_1}$, and most importantly, $\alpha_{t_2} \leq \alpha_{t_1}$. Therefore

$$\alpha_T \leq \frac{1}{T} \sum_{t=1}^T \alpha_t \leq \sqrt{\frac{C_1 \beta \tilde{\gamma}_T}{T}}$$

As a result, after at most $T \geq \frac{\beta \tilde{\gamma}_T C_1}{\epsilon_f^2}$ iterations, we have $\alpha_T \leq \epsilon_f$. \square

The rest of the proof for Theorem 5 is essentially the same as proof for Theorem 1 except for substituting Lemma 1 with Lemma 2.

E DEALING WITH BOUNDARY OPTIMUM

Here, we discuss the treatment and theoretical behavior when dealing with the boundary optimum as mentioned in Remark 2. First, we extend the results in Theorem 1, when not assuming the Assumption 2 hold. We uniformly shift the constraints by a small amount ϵ_C to satisfy Assumption 2 with the modified constraints. Formally, $\forall m \in \mathbf{M}$, $C'_m(\mathbf{x}) = C_m(\mathbf{x}) + \epsilon_C$. Then, running COBAR with these adjusted constraints, C'_m , instead of the original C_m , we have the following guarantee, which is a direct extension of Theorem 1. We denote the $\tilde{f}^* = f(\tilde{\mathbf{x}}^*)$, Here $\tilde{\mathbf{x}}^* = \arg \max_{\mathbf{x} \in \mathbf{X}, \forall m \in \mathbf{M}, C'_m(\mathbf{x}) > 0} f(\mathbf{x})$.

Corollary 4. *Under the aforementioned assumptions and modifications, with a constant $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ and the acquisition function from Algorithm 1, there exists an $\epsilon_f \leq \epsilon_C$, such that after at most $T \geq \frac{\beta \gamma_T C_1}{\epsilon_f^2}$ iterations, we have*

$$\mathbb{P} \left[|CI_{\tilde{f}^*, T}| \leq \epsilon_f, \tilde{f}^* \in CI_{\tilde{f}^*, T} \right] \geq 1 - \delta. \text{ Here, } C_1 = 8/\log(1 + \sigma^{-2}).$$

This Corollary 4 allows us to depict the width of the global optimum defined in the enlarged feasible region similarly. Since the optimum is defined in the enlarged area, it could be an upper bound of the global optimum defined in the original feasible region, including the feasible region boundaries. That is

$$\tilde{f}^* \geq \arg \max_{\mathbf{x} \in \mathbf{X}, \forall m \in \mathbf{M}, C_m(\mathbf{x}) \geq 0} f(\mathbf{x})$$

This allows us to extend further the Corollary 2 that depicts the partial cumulative regret after sufficient iterations and the upper bound of the violations.

Corollary 5. *Under the aforementioned assumptions and modifications, when $\forall m \in \mathbf{M}$, $C'_m(\mathbf{x}) > 0$, $\mathbf{x} \neq \mathbf{x}^*$, $\exists \epsilon_C > \epsilon'_C \geq \epsilon_f > 0$, $f^* - f(\mathbf{x}) \leq 2\epsilon_f$, it holds that $\forall m \in \mathbf{M}$, $C'_m(\mathbf{x}) \geq \epsilon'_C$. We use $\beta = 2 \log(\frac{2(M+1)|\tilde{D}|T}{\delta})$ and the acquisition function from Algorithm 1. After at most $t' \geq \frac{\beta \gamma_T C_1}{\epsilon_f^2}$ iterations, we have,*

$$\mathbb{P} \left[\sum_{t=t'}^T r(\mathbf{x}^*) - r(\mathbf{x}_t) \leq \sqrt{(T-t')\beta\gamma_T C_1}, \forall \mathbf{x} \in \tilde{D}, C_m(\mathbf{x}_t) \geq -\epsilon_C \right] > 1 - \delta. \text{ Here, } C_1 = 8/\log(1 + \sigma^{-2}) \text{ and } t' \leq T.$$

Proof. First, we have $\forall \mathbf{x} \in \tilde{D}_{\tilde{\mathbf{x}}_t}$, $2\beta^{1/2}\sigma_{f,t-1}(\mathbf{x}) \leq \alpha_t \leq \epsilon_f$.

Then $\forall \mathbf{x} \neq \tilde{\mathbf{x}}^*$ and $\mathbf{x} \in \tilde{D}_{\tilde{\mathbf{x}}_t}$, we have

$$\mathbb{P} \left[\tilde{f}^* - f(\mathbf{x}) \leq |CI_{\tilde{f}^*, t}| + \text{UCB}_{f,t}(\mathbf{x}) - \text{LCB}_{f,t}(\mathbf{x}) \leq 2\alpha_t \leq 2\epsilon_f \right] \geq 1 - \delta$$

Then by assumption, $\forall \mathbf{x} \in \tilde{D}_{\tilde{\mathbf{x}}_t}$, $\forall m \in \mathbf{M}$, we have probability at least $1 - \delta$, $C'_m(\mathbf{x}) \geq \epsilon'_C$. Hence we have both $C_m(\mathbf{x}) \geq \epsilon_C - \epsilon'_C$ and $\mathbf{x} \notin U_{C'_m, t}$. According to the algorithm, it regresses to GP-UCB by Srinivas et al. [2009] between t' and T .

$$\begin{aligned} \sum_{t=t'}^T (r(\mathbf{x}^*) - r(\mathbf{x}_t)) &\leq \sum_{t=t'}^T (r(\tilde{\mathbf{x}}^*) - r(\mathbf{x}_t))^2 \\ &\leq \beta C_1 (\gamma_T - \gamma_{t'}) \\ &\leq \beta C_1 \gamma_T (1 - t'/T) \end{aligned}$$

By Cauchy-Schwarz, we have

$$\begin{aligned} \sum_{t=t'}^T (r(\mathbf{x}^*) - r(\mathbf{x}_t)) &\leq \sqrt{(T-t') \sum_{t=t'}^T (r(\mathbf{x}^*) - r(\mathbf{x}_t))^2} \\ &\leq \sqrt{\frac{(T-t')^2}{T} \beta C_1 \gamma_T} \\ &\leq \sqrt{(T-t') \beta C_1 \gamma_T} \end{aligned}$$

□

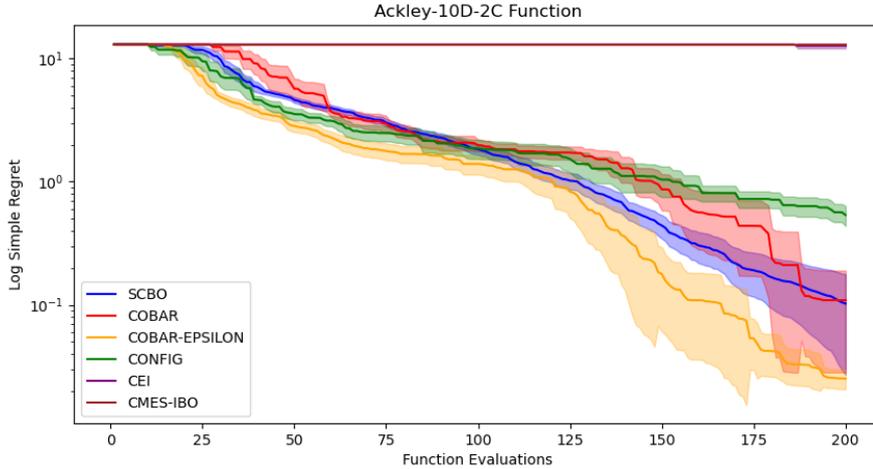


Figure 6: The figure illustrates the simple regret for Ackley-10D-2C. The results are collected from 15 independent trials. The shaded area denotes the 98% confidence interval. We reproduce the reported performance of SCBO using the corresponding Botorch tutorial. Then, we fix the kernel choices and other hyperparameters to make a fair comparison. For COBAR-EPSILON, we set $\epsilon_C = 1.2$.

Note that by enlarging the feasible region with ϵ_C , we don't risk losing the feasible region to enable COBAR to identify both interior and boundary optimum. We don't change the definition of f^* . Instead, we only leverage the modified running constraints and verify the feasibility with the original constraints. For the algorithms aiming at violation tolerant objectives like CONFIG [Xu et al., 2023], there is no similar guarantee with straightforward modification, e.g., adding small ϵ_C to the threshold while not risking losing the feasible region. We denote the modified COBAR as COBAR-EPSILON. We include the corresponding comparison on the noise-free Ackley-10D-2C as exactly defined in Eriksson and Poloczek [2021], where the feasible region is less than $2.2 * 10^{-3}\%$ of the whole search space, and the optimum lies on the boundary of feasible region by construction.

As is shown in figure 6, COBAR is initially outperformed by SCBO while converging to the near-optimal area after the sufficient budget as SCBO. COBAR-EPSILON archives the best convergence throughout the optimization with the proposed minor tweak. In contrast, CONFIG fails to converge to the global optimum, possibly due to its tolerance of the constraints violation, and by definition, no reward is incurred for a point outside the feasible region.

F REWARD FUNCTION

F.1 REWARD CHOICE 1: PRODUCT OF REWARD AND FEASIBILITY

The definition of reward plays an important role in online machine learning performance analysis. In the CBO setting, one possible definition of constrained reward derived from the constraint nature is $r(\mathbf{x}) = f(\mathbf{x}) \prod_m \mathbb{I}_{C_m(\mathbf{x}) \geq h_m}$ when assuming the $f(\mathbf{x}) > 0$. Considering both the aleatoric and epistemic uncertainty on the constraints, we could transform the problem into finding the maximizer

$$\arg \max_{\mathbf{x} \in \mathbf{X}} r(\mathbf{x}) = \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}) \prod_m \mathbb{P}[Y_{C_m}(\mathbf{x}) \geq h_m]$$

Here $Y_{C_m}(\mathbf{x})$ denotes the observation of the constraint C_m at \mathbf{x} .

The problem with this product reward, on the one hand, is that it is likely to incur a Pareto front if we regard the problem as a multi-objective optimization where the objectives are composed of $f(\mathbf{x})$ and $\mathbb{P}[Y_{C_m}(\mathbf{x}) \geq h_m]$. The multi-objective nature and resulting Pareto front indicate that the optimization could be more challenging to converge than the single-objective unconstrained BO problem, though the unique global optimum is not always expected there either. More critically, when the feasibility of reaching a certain threshold, we prefer to focus on optimizing the objective value rather than the product for the following reasons.

Firstly, the marginal gain on improving feasibility by increasing the value of the constraint function drops after the feasibility reaches 0.5 assuming it follows a Gaussian. Especially in the tail region, improving the feasibility and then the product of feasibility and objective value by optimizing the constraint function is prohibitively difficult.

Secondly, in most real-world scenarios except for certain applications that focus on feasibility (where the feasibility should be treated as another objective and make it in nature a multi-objective optimization), the actual marginal gain, in general, increases the feasibility decay faster than the increase of objective value. (e.g., when choosing between doubling the feasibility from 0.25 to 0.5 or doubling the objective value drop from 25 to 50, we probably favor the former as 0.25, meaning it is unlikely to happen. However, when choosing between increasing feasibility from .8 to .9 or increasing the objective drop from 80 to 90, there would be no such clear preference.) Then, the user would possibly favor the gain on the objective function after the feasibility reaches a certain level. Therefore, we propose the following reward for constrained optimization tasks according to this insight.

E.2 REWARD CHOICE 2: OBJECTIVE FUNCTION AFTER THE FEASIBILITY REACHING CERTAIN THRESHOLD

Instead of defining the reward as the product of the objective value and feasibility, we have to look into the probabilistic constraints and distinguish the epistemic uncertainty and aleatoric uncertainty. First, when assuming the observation on the constraints are noise-free, namely $Y_{C_m}(\mathbf{x}) = C_m(\mathbf{x})$, we could simply use the indicator function μ_m for each constraint to turn the feasibility function into an indicator function. This definition accommodates the scenarios where the infeasible region does not incur credible reward as discussed by Sacher et al. [2018], Bachoc et al. [2020] due to simulation failures

$$r(\mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } \mathbb{I}(C_m(\mathbf{x}) \geq h_m) \quad \forall m \in \mathbf{M} \\ -inf & \text{o.w} \end{cases} \quad (18)$$

Next, if the observation on the constraints is perturbed with a known Gaussian noise, namely $Y_{C_m}(\mathbf{x}) \sim \mathcal{N}(C_m(\mathbf{x}), \sigma)$, we could deal with the aleatoric uncertainty with a user-specific confidence level for each constraint $\chi_m \in (0, 1), \forall m \in \mathbf{M}$. Then we could turn $\mathbb{I}(Y_{C_m}(\mathbf{x}) \geq h_m)$ into probabilistic constraints following the definition proposed by Gelbart et al. [2014] and

$$\mathbb{P}[Y_{C_m}(\mathbf{x}) \geq h_m] \geq \chi_m$$

to explicitly deal with the aleatoric uncertainty. With the percentage point function (PPF), we could transform the probabilistic constraints into a deterministic constraint $\mathbb{I}(C_m(\mathbf{x}) \geq \hat{h}_m)$ with $\hat{h}_m = PPF(h_m, \sigma, \mu_m)$, meaning \hat{h} is the χ_m percent point of a Gaussian distribution with h_m and σ as its mean and standard deviation. Hence, we could unify the form of rewards of noise-free and noisy observation on the constraints with the user-specified confidence levels. For simplicity and without loss of generalization, we stick to the definition in equation 3 and let all $\hat{h}_m = 0$.

G ADDITIONAL EXPERIMENT DETAILS

In the benchmarks, we applied the deep Gaussian process [Wilson et al., 2016]. Here, we offer a more detailed discussion of the construction of the six CBO tasks studied in section 6.

G.1 SYNTHETIC TASKS

We study two synthetic CBO tasks constructed from conventional BO benchmark tasks. Here, we rely on the implementation contained in BoTorch’s [Balandat et al., 2020] test function module.

Rastrigin-1D-1C The Rastrigin function is a non-convex function used as a performance test problem for optimization algorithms. It was first proposed by Rastrigin [1974] and used as a popular benchmark dataset [Pohlheim, 2006]. It is constructed to be highly multimodal, with local optima being regularly distributed to trap optimization algorithms. Concretely, we negate the 1D Rastrigin function and try to find its maximum: $f(\mathbf{x}) = -10d - \sum_{i=1}^d (x_i^2 - 10 \cos(2\pi x_i))$, $d = 1$. The range of \mathbf{x} is $[-5, 5]$, and we construct the constraint to be $c(\mathbf{x}) = |\mathbf{x} + 0.7|^{1/2}$. When setting the threshold as $\sqrt{2}$, we essentially exclude the global optimum from the feasible area. The constraint enforces the optimization algorithm to explore

feasibility rather than allowing algorithms to improve the reward by merely optimizing the objective. Then, the feasible region takes up approximately 60% of the search space. This one-dimensional task is designed to illustrate the necessity of adaptively trade-off learning of constraints and optimization of the objective.

We also vary the threshold to control the portion of the feasible region to study the robustness of COBAR. Figure 3 shows the distribution of the objective function and feasible regions on the samples.

Ackley-5D-2C The Ackley function is also a popular benchmark for optimization algorithms. Compared with the Rastrigin function, it is similarly highly multimodal, while the region near the center is growingly steep. Same as what is done for Rastrigin, we negate the 5D Ackley function and try to find its maximum: $f(\mathbf{x}) = 20 \exp(-0.2\sqrt{1/d \sum_i^d x_i^2}) + \exp(1/d \sum_i^d \cos(2\pi x_i)) + 20 + \exp(1)$, $d = 5$. The search space is restricted to $[-5, 3]^5$. We construct two constraints to enforce a feasible area approximately taking up 14% of the search space. The first constraint $(\|x - \mathbf{1}\|_2 - 5.5)^2 - 1$ constructs two feasible regions. One of them lies in the center, and the other is close to the boundary of the search space. The second constraint $-\|x\|_\infty^2 + 9$ allows one hypercube feasible region in the center.

G.2 REAL-WORLD TASKS

We study four real-world CBO tasks. The first three are extracted from Tanabe and Ishibuchi [2020], which offers a broad selection of real-world multi-objective multi-constraints optimization tasks. The fourth one is a 32-dimensional optimization task extracted from the UCI Machine Learning repository [mis, 2019].

Vessel-4D-3C The pressure vessel design problem aims to optimize the total cost of a cylindrical pressure vessel. The four variables represent the thicknesses of the shell, the head of a pressure vessel, the inner radius, and the length of the cylindrical section. The problem is originally studied in Kannan and Kramer [1994], and we follow the formulation in RE2-4-3 in Tanabe and Ishibuchi [2020]. The feasible regions take up approximately 78% of the whole search space.

Spring-3D-6C The coil compression spring design problem aims to optimize the volume of spring steel wire, which is used to manufacture the spring [Lampinen and Zelinka, 1999] under static loading. The three input variables denote the number of spring coils, the outside diameter of the spring, and the spring wire diameter, respectively. The constraints incorporate the mechanical characteristics of the spring in real-world applications. We follow the formulation in RE2-3-5 in Tanabe and Ishibuchi [2020]. The feasible regions take up approximately 0.38% of the whole search space.

Car-7D-8C The car cab design problem includes seven input variables and eight constraints. The problem is originally studied in Deb and Jain [2013]. We follow the problem formulation in RE9-7-1 in Tanabe and Ishibuchi [2020] and focus on the objective of minimizing the weight of the car while meeting the European enhanced Vehicle-Safety Committee (EEVC) safety performance constraints. The seven variables indicate the thickness of different parts of the car. The feasible region takes up approximately 13% of the whole search space.

Converter-32D-3C This UCI dataset we use consists of positions and absorbed power outputs of wave energy converters (WECs) from the southern coast of Sydney. The applied converter model is a fully submerged three-tether converter called CETO. 16 WECs 2D-coordinates are placed and optimized in a size-constrained environment [mis, 2019]. The input is, therefore, 32 dimensional. We place three constraints on the tasks, including the absorbed power of the first two converters being above a certain threshold of 96000 and the general position being not too distant with the two-norm below 2000. The feasible region takes up approximately 27% of the whole search space.

H ADDITIONAL EXPERIMENTS

Here, we provide additional experiment results on COBAR.

H.1 ROBUSTNESS TO CHOICES OF β

As is shown in figure 7, the algorithm is robust to moderate values of β . Except from the Ackley $\beta = 0.1$ where the filtering of ROI is over-aggressive and traps the model on a certain locality when a very small number of candidates remain in ROI. We observe that certain β choices could be slightly better but don't impact the convergence and lack statistical significance.

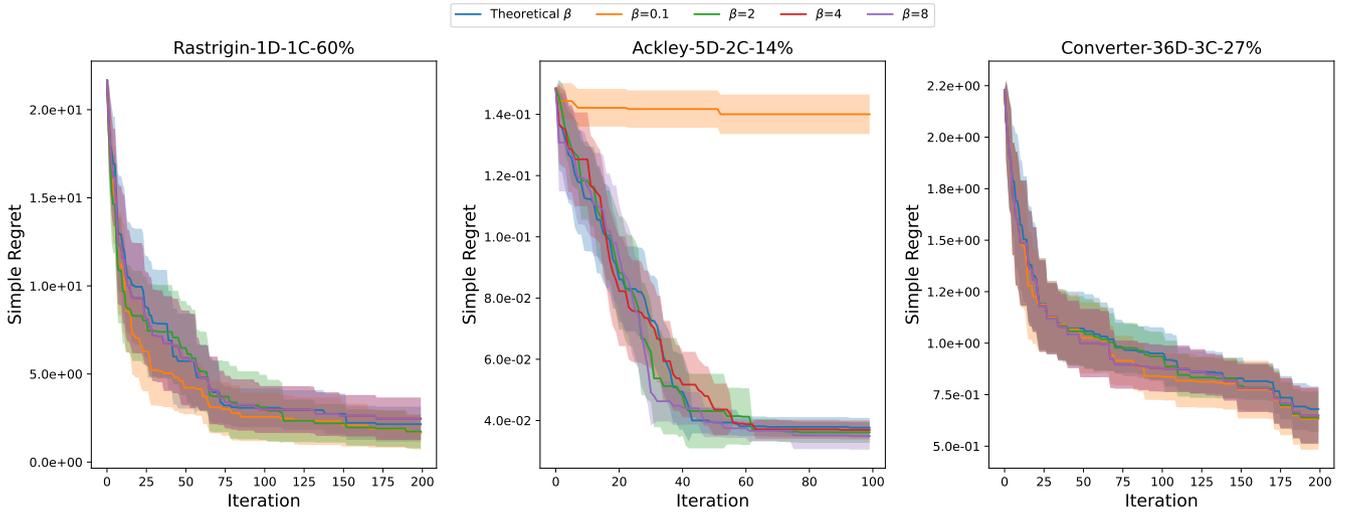


Figure 7: The figure illustrates the simple regret for a different choice of constant β for COBAR. Here the theoretical β are 6.51 for Rastrigin-1D-1C, 6.47 for Ackley-5D-2C, and 6.51 for Converter-36D-3C. The results are collected from 15 independent trials.

Problem	COBAR	CMES-IBO	SCBO	cEI
Rastrigin-1D-1C	144.29	545.83	32.39	231.12
Ackley-5D-2C	96.19	565.10	25.43	180.39
Converter-36D-3C	190.05	660.27	31.73	267.36

Table 3: Average wall time (sec) of different CBO Methods collected from 15 independent trials.

We believe the acquisitions in Eq. (6) and Eq. (5), together with the \hat{X} identification when the models are well-fitted, contribute to this robustness. Different from conventional GP-UCB [Srinivas et al., 2009], the acquisition functions are standardized with the (maximum) lower confidence bound. The search domains are filtered when historical observations suggest poor performance in nearby areas.

H.2 WALL TIME

We show the wall time of COBAR compared with the baselines in table 3. The results demonstrate the efficiency of COBAR due to the ROI filtering reducing the search space, though the ROI identification incurs additional cost for membership check.

H.3 ADDITIONAL COMPARISON WITH CONFIG

Though the objective is defined differently, we add additional baseline CONFIG from Xu et al. [2023]. The results are shown in figure 8 and figure 9. We observe that COBAR outperforms or at least matches CONFIG in all the problems in our setting. Specifically, in the early stage of the Rastrigin-1D-1C task and through the Ackley-5D-2C, where the underlying objective is highly fluctuating, as is shown in figure 8 for Rastrigin-1D-1C, CONFIG fails to enter the feasible region consistently even after exhausting sufficient budget and gets stuck in learning the constraints passively.

At the same time, we observe that on Converter-32D-3C, Vessel-4D-3C, and Spring-3D-6C, CONFIG generally matches the performance of COBAR. We hypothesize that in these applications, the constraint learning part of COBAR is not as beneficial as directly optimizing the underlying function is possibly feasible regions, as the unknown feasibility coincides with the optimality of the underlying objectives. Still, COBAR bears higher consistency in all the benchmarks, highlighting the efficiency and necessity of the adaptive trade-off of active learning and optimization in COBAR when assuming no reward is incurred outside the feasible region. This difference also highlights the necessity of actively learning the complex underlying constraints to guarantee a stable convergence to a feasible optimum.

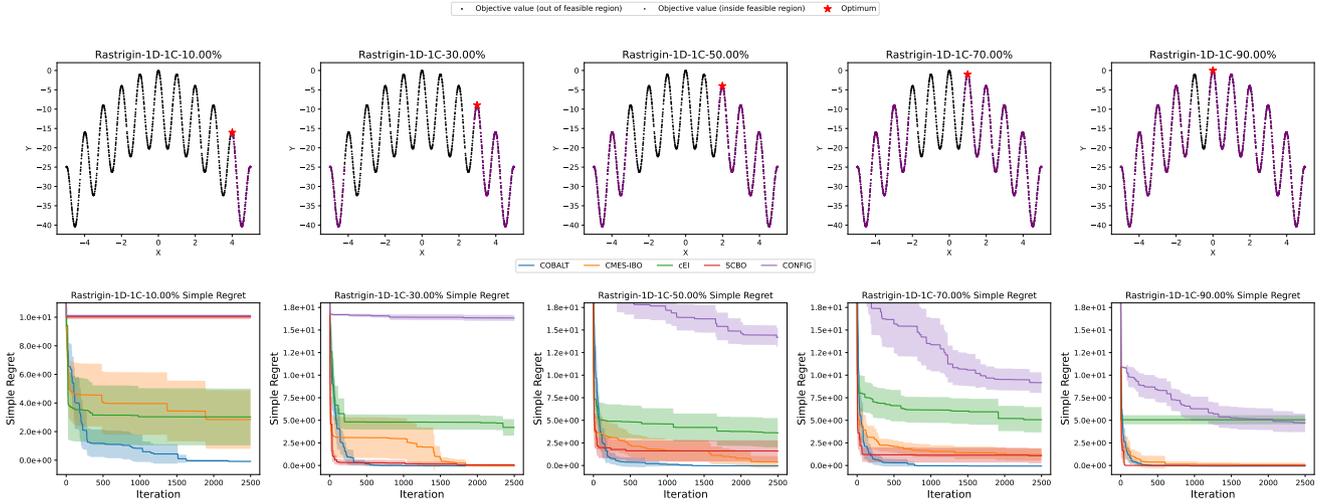


Figure 8: We use black dots and purple dots to show the infeasible region and feasible region in the first row correspondingly. Each column corresponds to a certain threshold choice for the single constraint $c(\mathbf{x}) = |\mathbf{x} + 0.7|^{1/2}$ in the Rastrigin-1D-1C task. The search space contains a certain portion of the feasible region, denoted on each figure and title. The first row shows the distribution of 1000 samples from the noise-free distribution objective function, and the figures are differentiated with different feasible regions. The second row shows corresponding simple regret curves. We test each method with 15 independent trails and impose observation noises sampled from $\mathcal{N}(0, 0.1)$ not shown in the first row. The scaling and length scale of the GPs are learned via maximum likelihood estimation.

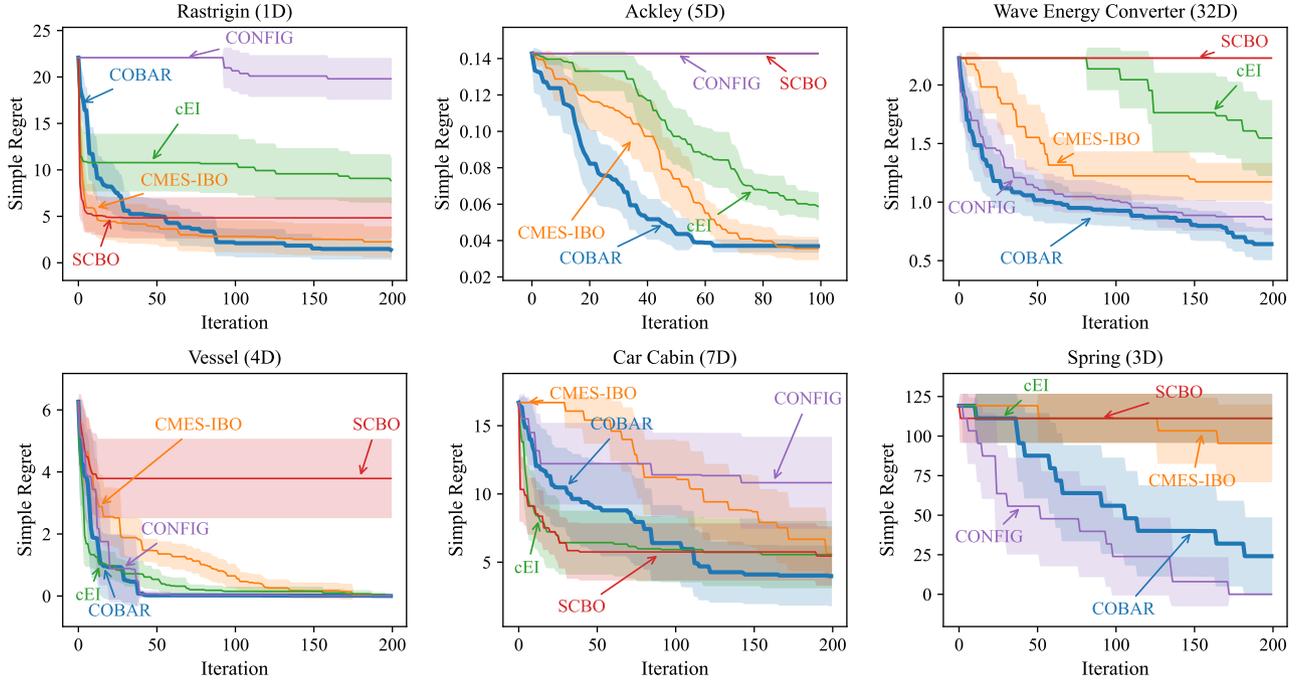


Figure 9: The input dimensionality, the number of constraints, and the approximate portion of the feasible region in the whole search space for each task are denoted on the titles. The curves show the average simple regret after standardization, while the shaded area denotes the 95% confidence interval through the optimization.

H.4 ADDITIONAL COMPARISON WITH SVM-CBO

SVM-CBO [Antonio, 2021] offers a practicality-oriented solution. It uses the SVM to learn the feasibility bound to estimate the decision boundary efficiently. The challenge of analyzing the learning of SVM combined with the coverage-oriented first-phase acquisition function poses a challenge to regret analysis. In addition, SVM-CBO requires a specific split of feasibility identification and optimization within the feasible region and demands different performance metrics for evaluation. This split makes direct comparisons with COBAR, which is somewhat challenging and does not explicitly split the two processes. Nonetheless, we follow the practice in the paper that uses a 10:60:30 split for the random sampling, phase 1 and phase 2 of SVM-CBO, and report the simple regret.

Experiment	COBAR-70	CMES-IBO-70	cEI-70	SCBO-70	SVM-CBO-70
Rastrigin-1D-1C-60%	3.80e+00 (1.85e+00)	3.00e+00 (1.84e+00)	1.08e+01 (3.02e+00)	4.83e+00 (2.12e+00)	4.83e+00 (1.43e+00)
Ackley-5D-2C-14%	3.71e-02 (7.19e-04)	4.41e-02 (6.67e-03)	7.94e-02 (1.46e-02)	1.43e-01 (0.00e+00)	1.11e-01 (2.45e-02)
Converter-36D-3C-27%	9.50e-01 (1.77e-01)	1.32e+00 (2.31e-01)	2.23e+00 (0.00e+00)	2.23e+00 (0.00e+00)	1.09e+00 (1.52e-01)
Vessel-4D-3C-78%	2.06e-02 (1.39e-02)	1.20e+00 (3.90e-01)	2.44e-01 (2.48e-01)	3.79e+00 (1.25e+00)	2.59e-02 (3.36e-02)
Car_Cabin-7D-8C-13%	8.62e+00 (3.30e+00)	1.40e+01 (2.14e+00)	6.15e+00 (2.28e+00)	5.75e+00 (2.03e+00)	6.84e+00 (3.24e+00)
Spring-3D-6C-0.38%	6.40e+01 (2.83e+01)	1.11e+02 (1.50e+01)	1.11e+02 (1.50e+01)	1.11e+02 (1.50e+01)	8.35e+01 (2.76e+01)
Experiment	COBAR-100	CMES-IBO-100	cEI-100	SCBO-100	SVM-CBO-100
Rastrigin-1D-1C-60%	2.21e+00 (1.41e+00)	2.84e+00 (1.62e+00)	1.07e+01 (3.07e+00)	4.83e+00 (2.12e+00)	2.67e+00 (8.14e-01)
Ackley-5D-2C-14%	3.69e-02 (3.08e-03)	3.56e-02 (5.93e-03)	5.88e-02 (7.24e-03)	1.43e-01 (0.00e+00)	1.09e-01 (2.64e-02)
Converter-36D-3C-27%	9.29e-01 (1.27e-01)	1.22e+00 (2.02e-01)	2.14e+00 (1.75e-01)	2.23e+00 (0.00e+00)	9.73e-01 (1.45e-01)
Vessel-4D-3C-78%	1.94e-02 (1.43e-02)	6.48e-01 (3.60e-01)	1.51e-01 (1.25e-01)	3.79e+00 (1.25e+00)	2.23e-02 (8.96e-04)
Car_Cabin-7D-8C-13%	6.40e+00 (2.72e+00)	1.12e+01 (2.71e+00)	5.92e+00 (2.34e+00)	5.75e+00 (2.03e+00)	6.03e+00 (2.40e+00)
Spring-3D-6C-0.38%	5.60e+01 (2.99e+01)	1.11e+02 (1.50e+01)	1.11e+02 (1.50e+01)	1.11e+02 (1.50e+01)	8.35e+01 (2.76e+01)

Table 4: Comparison of different methods’ simple regrets across experiments. The table shows the updated experiment results after incorporating the SVM-CBO as an additional baseline. The upper block shows the simple regret at 70 iterations, while the lower shows the simple regret at 100 iterations. The standard error is shown in parentheses.

Table 4 shows the simple regret of the end of both phases of SVM-CBO. We emphasize the best simple regret achieved. The results demonstrate that COBAR ultimately outperforms or matches the best baseline in the end.

To provide a clearer, high-level quantitative summary, we aggregate the final performance of the core methods from Table 4. Table 5 shows the average rank of each method based on its mean simple regret at the final iteration (T=100) across all six problems. As the summary shows, COBAR (COBAR) achieves the best overall rank, confirming its strong and consistent performance.

Table 5: Average method ranking (lower is better) at the final iteration (T=100). Ranks are calculated for each of the 6 problems based on the mean simple regret data in Table 4 and then averaged. SVM-CBO is excluded from the ranking to ensure a fair comparison with the core baselines from the main paper.

Method	Average Rank	Overall Rank
COBAR	~ 1.17	1st
CMES-IBO	~ 2.33	2nd
cEI	~ 3.17	3rd
SCBO	~ 3.33	4th

I DISCUSSIONS

Here, we offer additional explanation and discussion over COBAR.

I.1 ADDITIONAL EXPLANATION OF COBAR

For algorithm 1, $\{\mathbf{x}_{g_t,t}\}$ in line 11 are acquired in line 7 as $\mathbf{x}_{C_{m,t}}$ or line 9 as $\mathbf{x}_{f,t}$, since \mathcal{G} is composed of C_m and f . Roughly speaking, we are taking $\arg \max_{g,x}$, yet we avoid using such notation for two reasons. (1) the domain where equation 5 and equation 6 are maximized are different; (2) the domain for equation 6 could even be empty. Therefore, we

are currently taking the $\arg \max$ of equation 5 and equation 6 over different domains (if not empty) separately and then taking the $\arg \max$ of the corresponding acquisition function values as in line 11.

I.2 ON THE COMPARABILITY OF ACQUISITION FUNCTIONS OVER DIFFERENT UNDERLYING FUNCTIONS

Both the acquisition functions for optimizing the objective and active learning are confidence interval-based, which reflects the uncertainty and is intrinsically comparable. With Assumption 1 that the black-box underlying functions are samples from the corresponding GPs specified by the kernels, we use the kernels to capture the scaling of the different unknowns. Our analysis does not assume that the kernels are the same, meaning that the theoretical results hold when the objective and constraints are of different scales. This analysis converts the algorithm’s sensitivity to scale to the sensitivity to hyper-parameter misspecification. In our experiments, we report the results when following the standard practice of kernel learning [Rasmussen and Williams, 2006] for both the proposed algorithm and baselines, as is stated at the end of the caption of figure 3. In summary, the compatibility is guaranteed by the properly specified kernel. Recent advancements in self-correcting BO [Hvarfner et al., 2024] or BO with unknown hyperparameters [Berkenkamp et al., 2019] propose various methods to address the challenge.

With regard to the practical concern over why the analysis does not require normalizing the different acquisition functions, the answer is threefold. First, since the correlation between the constraints and the objective is unknown, it is possible that the objective, in general, is of a smaller scale but bears the highest gradient near the boundary, meaning that the general scale of functions does not offer a guarantee to normalize the near-boundary uncertainties. Second, using the ROI to constrain the acquisition helps exclude the useless uncertainty reduction as the ROI considers both the objective and the constraints. If constraints dominate the COBAR acquisition, it suggests that the selected points remain likely to contain the global optimum as its objective does not have a high probability of being suboptimal. Such a query won’t be wasted. A concrete example is illustrated in figure 2. Third, since we are assuming a universal upper bound for each constraint, the scale difference could make certain constraints dominant in the ϵ_m . This could be addressed through the normalization of observations given prior knowledge.

I.3 DIFFERENCE FROM OTHER EXISTING CBO METHODS WITH NO-REGRET GUARANTEE

We briefly discuss the differences between COBAR and the previous theoretical results in CBO. Lu and Paulson [2022] addresses equality constraints for instantaneous penalty-based regret. However, the reward formulation is different. Lu and Paulson [2023] offers theoretical results on cumulative regret and violations. Yet, they assume querying points out of the feasible region still yields rewards and consider the violation separately.

In general, we are unaware that the existing CBO analysis results lead to a similar guarantee as in our work when assuming querying infeasible points does not yield a reward. One key difference is that with the active learning component and feasibility assumption, we could guarantee to query a feasible point that bears a reward converging to optimal value with the desired confidence. In our specific reward formulation, we regard such a guarantee and, therefore, the contribution in algorithm design and analysis as sufficiently different from the previous work, even when only focusing on the coupled setting.

I.4 EMPTY SUBSETS OF SEARCH SPACE

It is possible that certain subsets discussed in section 4 could be empty at a certain t as a result of intersections. However, according to the assumptions in section 5 and Lemma 1, the properly chosen β does not result in over-aggressive filtering with high probability. From this perspective, ROI \tilde{X} is soundly defined. COBAR is also robust to empty $U_{C_m,t}$. As shown in algorithm 1, the domain where the acquisition functions defined in equation 6 and equation 5 are maximized allow empty $U_{C_m,t}$ for COBAR to proceed.

I.5 LIMITATIONS AND FUTURE WORK

The limitation of COBAR includes (1) the inefficiency of identifying the ROIs due to the pointwise comparison in current implementation relying on discretization; (2) the lack of discussion over correlated unknowns, which are common in practice (e.g., two constraints are actually lower bound and upper bound of the same value). Though we briefly discuss

and study corresponding scenarios, we expect the following work could improve the algorithm's effectiveness and the comprehensiveness of corresponding analysis accordingly.