OPTIMISTIC BAYESIAN OPTIMIZATION WITH UNKNOWN CONSTRAINTS

Quoc Phong Nguyen[1], Wan Theng Ruth Chew[2], Le Song[2],
Bryan Kian Hsiang Low[2] & Patrick Jaillet[1]
[1]LIDS and EECS, Massachusetts Institute of Technology, USA
[2]School of Computing, National University of Singapore, Singapore
qphongmp@gmail.com, chew.ruth@u.nus.edu, le.song@u.nus.edu,
lowkh@comp.nus.edu.sg, jaillet@mit.edu

ABSTRACT

Though some research efforts have been dedicated to constrained *Bayesian optimization* (BO), there remains a notable absence of a principled approach with a theoretical performance guarantee in the *decoupled setting*. Such a setting involves independent evaluations of the objective function and constraints at different inputs, and is hence a relaxation of the commonly-studied *coupled setting* where functions must be evaluated together. As a result, the decoupled setting requires an adaptive selection between evaluating either the objective function or a constraint, in addition to selecting an input (in the coupled setting). This paper presents a novel constrained BO algorithm with a provable performance guarantee that can address the above relaxed setting. Specifically, it considers the fundamental trade-off between exploration and exploitation in constrained BO, and, interestingly, affords a noteworthy connection to active learning. The performance of our proposed algorithms is also empirically evaluated using several synthetic and real-world optimization problems.

1 INTRODUCTION

In real-world applications, we often encounter expensive-to-evaluate *black-box* objective functions that can only be assessed through simulations or experimentation. For example, problems involve optimizing the hyperparameters of a machine learning model (Wistuba et al., 2018; Perrone et al., 2020), or choosing experiments in the fields of material and drug design (Schweidtmann et al., 2018). To address these problems, *Bayesian optimization* (BO) has gained prominence as a widely adopted approach (Brochu et al., 2010; Frazier, 2018; Garnett, 2022). It is an iterative model-based approach that employs a probabilistic model, e.g., a *Gaussian process* (GP), to estimate the unknown objective function. At each iteration, BO searches for the optimal solution by strategically selecting an *input query* to evaluate the objective function, maintaining a balance between exploiting promising areas and exploring poorly-estimated regions. BO encompasses many well-established techniques such as the probability of improvement (Kushner, 1964), expected improvement (EI) (Mockus et al., 1978), Gaussian process upper confidence bound (GP-UCB) (Srinivas et al., 2010), the knowledge-gradient based approach (Frazier et al., 2008), and information-theoretic approaches: entropy search (Hennig and Schuler, 2012), predictive entropy search (PES) (Hernández-Lobato et al., 2014), and max-value entropy search (MES) (Wang and Jegelka, 2017).

Beyond the black-box objective function, recent advancements in BO have focused on addressing the prevalent presence of *black-box constraints*. For example, there often exist prediction time constraints and class-wise performance constraints when tuning machine learning models (Hernández-Lobato et al., 2016; Takeno et al., 2022). They are just as costly to evaluate as the objective function. Constrained BO has led to many BO extensions such as EIC (an EI-based method) (Gardner et al., 2014), a knowledge gradient-based method (Chen et al., 2021), CMES-IBO (an MES-based method) (Takeno et al., 2022), augmented Lagrangian approaches (Gramacy et al., 2016; Picheny et al., 2016), and upper trust bound (UTB) (a GP-UCB-based method) (Priem et al., 2020).

However, existing approaches primarily concentrate on empirical performance and lack a theoretical analysis to ensure consistent performance. Recently, theoretical studies for constrained BO have gained attention via the works of Lu and Paulson (2022) and Xu et al. (2023). While Lu and Paulson (2022) perform the analysis by proposing a penalty-based regret, Xu et al. (2023) analyse the cumulative regret due to the objective function and the constraint violation separately.

Besides the above largely unexplored theoretical analysis, existing works often overlook the potential for evaluating the objective function and constraints independently at different inputs, known as *decoupled queries*. Specifically, the above works, including the theoretical studies by Lu and Paulson (2022) and Xu et al. (2023), require simultaneous evaluations of the objective function and constraints at an input query, known as *coupled queries*. The distinction between coupled and decoupled queries was first mentioned in the work of Gelbart et al. (2014). They discuss a chicken-and-egg pathology which prevents extending a *myopic* BO approach such as EIC to the decoupled setting. Later, Hernández-Lobato et al. (2016) introduce a principled approach based on PES, namely PESC, to address constrained BO with decoupled queries. While PESC is functionally equivalent to a lookhead approach, it leverages the symmetric property of mutual information to avoid performing the actual lookahead computation. Regrettably, its implementation is fairly complex, making it less accessible to practitioners. Besides, it lacks a theoretical performance guarantee. While ADMMBO (Ariafar et al., 2019) deals with decoupled queries, it deterministically alternates the evaluations of the objective function and constraints, without exploiting the benefits of an adaptive selection approach. Hence, the question of devising an approach that offers a theoretical performance guarantee, is adaptable to decoupled queries, and can be readily implemented by practitioners remains unanswered.

In this paper, we address this question by proposing a simple algorithm with a theoretical performance guarantee, especially in the decoupled setting. Notably, our algorithm is myopic without expensive lookaheads. In Sec. 2, we introduce a regret that does not require any penalty parameter, unlike that in the work of Lu and Paulson (2022). Then, we discuss the exploration-exploitation trade-off in constrained BO in Sec. 3. Specifically, we introduce a new form of exploration, namely *horizontal exploration*, resulted from the presence of black-box constraints. It is to differentiate from the *vertical exploration* in unconstrained BO (Srinivas et al., 2010). Then, we design a unified approach that handles coupled and decoupled queries from this perspective. More importantly, our algorithms are shown to be no-regret in Theorem 3.3 and App. B. While the viewpoint of balancing exploration and exploitation is inherently grounded in BO, especially in the bandit setting like GP-UCB (Srinivas et al., 2010), Sec. 3.3 shows that the choice of the function to query can also be framed within the well-known uncertainty sampling paradigm in the active learning literature (Settles, 2009). In Sec. 3.4, we propose an estimator for approximating the optimal solution at each BO iteration with a theoretical performance guarantee. To empirically demonstrate the performance of our algorithms, we presents several experiments using both synthetic and real-world optimization problems in Sec. 4.

2 Constrained Bayesian Optimization and Regret Definition

Let f be a real-valued black-box objective function and $\mathcal{C}$ be a finite set of real-valued black-box constraints. Let the compact subset $\mathcal{X} \subset \mathbb{R}^d$ be the input domain and $d \in \mathbb{N}_+$ be the input dimension. We consider the following constrained optimization problem

$$\max_{\mathbf{x}\in\mathcal{S}} f(\mathbf{x}) \text{ where the } \textit{feasible region } \mathcal{S} \triangleq \{\mathbf{x} \in \mathcal{X} |\ c(\mathbf{x}) \geq \lambda_c\ \forall c \in \mathcal{C}\} \text{ and } \lambda_c \in \mathbb{R}\ \forall c \in \mathcal{C}\ . \quad (1)$$

An equality constraint can be transformed into two inequality constraints. Let us denote the set of the objective function and constraints as $\mathcal{F} \triangleq \{f\} \cup \mathcal{C}$ and denote a function in $\mathcal{F}$ as h.

To identify the optimal solution $\mathbf{x}^* \triangleq \arg\max_{\mathbf{x}\in\mathcal{S}} f(\mathbf{x})$, we employ BO which is an algorithm operating in a sequential manner. In the *decoupled query* setting, at iteration t, we gather a noisy observation of the *function query* $h_t \in \mathcal{F}$ evaluated at an *input query* $\mathbf{x}_t \in \mathcal{X}$

$$y_{h_t}(\mathbf{x}_t) \triangleq h_t(\mathbf{x}_t) + \epsilon_{h_t}(\mathbf{x}_t), \quad \epsilon_{h_t}(\mathbf{x}_t) \sim \mathcal{N}(0, \sigma^2_{h_t})\ . \quad (2)$$

For instance, at iteration t, the algorithm may decide to query the objective function (i.e., $h_t = f$) while it may query a constraint (i.e., $h_{t'} = c$ for some $c \in \mathcal{C}$) at a different iteration $t' \neq t$. On the contrary, in the *coupled query* setting, at iteration t, we query for observations $\{y_h(\mathbf{x}_t)\}_{h\in\mathcal{F}}$ of all functions $h \in \mathcal{F}$ evaluated at the same *input query* $\mathbf{x}_t \in \mathcal{X}$. The coupled setting is less challenging since it does not require specifying the function query.

Let $\mathcal{D}_{h,t}$ denote the set of observed inputs of h until iteration t (including $\mathbf{x}_t$) and $\mathcal{D}_t \triangleq \cup_{h\in\mathcal{F}}\mathcal{D}_{h,t}$. Then, $\mathcal{D}_0$ consists of initial observed inputs. Let $\mathbf{y}_h(\mathcal{D}_t) = \mathbf{y}_h(\mathcal{D}_{h,t}) \triangleq \{y_h(\mathbf{x})\}_{\mathbf{x}\in\mathcal{D}_{h,t}}$ denote the set of observations from h at $\mathcal{D}_{h,t}$. BO effectively utilizes the acquired observations $\{\mathbf{y}_h(\mathcal{D}_{h,t-1})\}_{h\in\mathcal{F}}$ in the previous $t-1$ iterations to formulate a strategy for determining the next input query $\mathbf{x}_t$, the next function query h_t, and an estimator, denoted as $\tilde{\mathbf{x}}^*_t$, for approximating the optimal solution $\mathbf{x}^*$. We will discuss the probabilistic model of $h \in \mathcal{F}$ given the acquired observations and a performance metric of BO, called the *regret*, in the rest of this section. Then, we will elaborate on our strategy of choosing $\mathbf{x}_t$, h_t, and $\tilde{\mathbf{x}}^*_t$ in the following sections.

Gaussian process. For each $h \in \mathcal{F}$, we model h with a *Gaussian process* (GP). It implies that every finite subset of $\{h(\mathbf{x})\}_{\mathbf{x}\in\mathcal{X}}$ follows a multivariate Gaussian distribution. The GP is fully specified by its prior mean $m_h(\mathbf{x})$ and its kernel $k_h(\mathbf{x},\mathbf{x}') \triangleq \mathrm{cov}(h(\mathbf{x}), h(\mathbf{x}'))$. We employ the commonly-used *squared exponential* (SE) kernel. At iteration t, given the observations $\mathbf{y}_h(\mathcal{D}_{t-1})$ in the previous $t-1$ iterations, the posterior distribution of $h(\mathbf{x})$ follows a Gaussian distribution with a closed-form posterior mean and variance, denoted as $\mu_{h,t-1}(\mathbf{x})$ and $\sigma^2_{h,t-1}(\mathbf{x})$, respectively.[1]

Regrets. To analyse the theoretical performance of constrained BO, we propose the following *instantaneous regret* r including that of the objective function r_f and the constraints r_c.

$$r(\mathbf{x}_t) \triangleq \max_{h\in\mathcal{F}} r_h(\mathbf{x}_t) \quad \text{where} \quad \begin{array}{c} r_f(\mathbf{x}_t) \triangleq \max(0, f(\mathbf{x}^*) - f(\mathbf{x}_t)) \\ \forall c \in \mathcal{C},\ r_c(\mathbf{x}_t) \triangleq \max(0, \lambda_c - c(\mathbf{x}_t)) \,. \end{array} \tag{3}$$

Then, our goal is to design BO algorithms that achieve a sublinear cumulative regret

$$\lim_{T\to\infty} \frac{1}{T} R_T \triangleq \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} r(\mathbf{x}_t) = 0 \,. \tag{4}$$

as it implies that $\min_{\mathbf{x}\in\{\mathbf{x}_t\}_{t=1}^T} r(\mathbf{x}_t) \leq \frac{1}{T}\sum_{t=1}^T r(\mathbf{x}_t)$ approaches 0 as T approaches ∞.

Remark 2.1 (Instantaneous regret as a sum). Alternatively, the instantaneous regret can be defined as a sum of instantaneous regrets of the objective function and the constraints

$$s(\mathbf{x}_t) \triangleq \sum_{h\in\mathcal{F}} r_h(\mathbf{x}_t) \,. \tag{5}$$

Let us consider the case of a single constraint $\mathcal{C} = \{c_0\}$. For $\mathbf{x}_t \neq \mathbf{x}'_t$, if $r_f(\mathbf{x}_t) = r_f(\mathbf{x}'_t) = 1$ and $r_{c_0}(\mathbf{x}_t) = 0$ while $r_{c_0}(\mathbf{x}'_t) = 1$, then $r(\mathbf{x}_t) = r(\mathbf{x}'_t) = 1$ while $s(\mathbf{x}_t) = 1 < 2 = s(\mathbf{x}'_t)$. As a result, $s(\mathbf{x})$ is more effective than $r(\mathbf{x})$ at measuring the suboptimality of a solution. Nevertheless, a sublinear cumulative regret w.r.t. r implies a sublinear cumulative regret w.r.t. s and vice versa since $r(\mathbf{x}_t) \leq s(\mathbf{x}_t) \leq |\mathcal{F}|\, r(\mathbf{x}_t)$. We revisit $s(\mathbf{x}_t)$ in Sec. 3.4 when discussing the estimator $\tilde{\mathbf{x}}^*_t$.

3 Optimistic Bayesian Optimization with Unknown Constraints

To simplify the derivation, we consider the case of finite input domain $\mathcal{X}$ and utilize the following Lemma 5.1 of Srinivas et al. (2010) with a modification by applying the union bound for all functions in $\mathcal{F}$ (Lu and Paulson, 2022).

Lemma 3.1. *Pick $\delta \in [0,1]$ and set $\beta_t = 2\log(|\mathcal{F}||\mathcal{X}|t^2\pi^2/6\delta)$. Then,*

$$|h(\mathbf{x}) - \mu_{h,t-1}(\mathbf{x})| \leq \beta_t^{1/2}\sigma_{h,t-1}(\mathbf{x}) \quad \forall \mathbf{x}\in\mathcal{X}\ \forall t \geq 1\ \forall h \in \mathcal{F} \tag{6}$$

holds with probability $\geq 1-\delta$. It suggests

$$u_{h,t-1}(\mathbf{x}) \triangleq \mu_{h,t-1}(\mathbf{x}) + \beta_t^{1/2}\sigma_{h,t-1}(\mathbf{x}) \ \textit{ and } \ l_{h,t-1}(\mathbf{x}) \triangleq \mu_{h,t-1}(\mathbf{x}) - \beta_t^{1/2}\sigma_{h,t-1}(\mathbf{x}) \tag{7}$$

as the upper *and* lower confidence bounds *of $h(\mathbf{x})$ for all $h \in \mathcal{F}$, $\mathbf{x}\in\mathcal{X}$, and $t \geq 1$, respectively.*

[1] Please refer to Rasmussen and Williams (2006) for the closed-form expressions.

3.1 INPUT QUERY

In the classic unconstrained GP-UCB work of Srinivas et al. (2010) (i.e., $\mathcal{C} = \emptyset$), it balances between exploiting the current posterior belief by selecting those with high posterior mean $\mu_{f,t-1}(\mathbf{x})$, and exploring inputs with highly uncertain evaluations of f by selecting those with high posterior standard deviations $\sigma_{f,t-1}(\mathbf{x})$. Specifically, GP-UCB (Srinivas et al., 2010) selects an input query that maximizes an *optimistic objective function evaluation* $u_{f,t-1}(\mathbf{x})$

$$\mathbf{x}_t^{\text{GP-UCB}} = \arg\max_{\mathbf{x}\in\mathcal{X}} u_{f,t-1}(\mathbf{x}) = \arg\max_{\mathbf{x}\in\mathcal{X}} \Big(\mu_{f,t-1}(\mathbf{x}) + \underbrace{\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x})}_{\text{vertical exploration bonus}}\Big) . \tag{8}$$

Let us call $\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x})$ the *vertical exploration bonus* as it encourages us to be optimistic about the unknown objective function evaluations.[2] It is shown as the blue region in Fig. 1a.

In the presence of unknown constraints $\mathcal{C}$, apart from the above *optimistic objective function evaluation* $u_{f,t-1}(\mathbf{x})$, we additionally consider an *optimistic feasible region*, denoted as $\mathcal{O}_t$, which may contain some infeasible inputs. It represents the exploration of the feasible region which we refer to as *horizontal exploration* as opposed to the vertical exploration in the optimistic objective function evaluation. Ideally, $\mathcal{O}_t$ is a superset of the feasible region $\mathcal{S}$ to ensure that the optimal solution $\mathbf{x}^*$ remains in $\mathcal{O}_t$. We consider the following optimistic feasible region

$$\mathcal{O}_t \triangleq \{\mathbf{x}\in\mathcal{X}|\, u_{c,t-1}(\mathbf{x}) \geq \lambda_c \;\forall c\in\mathcal{C}\} \tag{9}$$

which aligns with our goal that $\mathcal{O}_t \supset \mathcal{S}$ with high probability since $u_{c,t-1}(\mathbf{x}) \geq c(\mathbf{x})$ with high probability (Lemma 3.1). By considering an optimistic feasible region, we avoid the subtle issue that the probability mass of the feasible region is 0 given the GP posterior beliefs of the constraints. This issue affects several existing approaches such as EIC, PESC, and CMES-IBO as discussed in App. A.

Combining the vertical and horizontal explorations, we select the input query $\mathbf{x}_t$ that maximizes the optimistic objective function $u_{f,t-1}$ restricted to the optimistic feasible region $\mathcal{O}_t$

$$\mathbf{x}_t \triangleq \arg\max_{\mathbf{x}\in\mathcal{O}_t} u_{f,t-1}(\mathbf{x}) . \tag{10}$$

We derive an upper confidence bound of $r_f(\mathbf{x}_t)$ similar to that in GP-UCB, and an additional upper confidence bound of $r_c(\mathbf{x}_t)$ which holds with probability $\geq 1-\delta$ (see App. B)

$$r_f(\mathbf{x}_t) \leq 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t) \quad \text{and} \quad r_c(\mathbf{x}_t) \leq 2\beta_t^{1/2}\sigma_{c,t-1}(\mathbf{x}_t) . \tag{11}$$

If the queries are *coupled*, we can achieve a no-regret BO algorithm by simply obtaining observations $\{y_h(\mathbf{x}_t)\}_{h\in\mathcal{F}}$ at iteration t as elaborated in App. B. This algorithm for the coupled setting is called UCB-C to distinguish it from another algorithm, namely UCB-D, for the decoupled setting described in the next section.

While a condition resembling equation (9) is utilized in the work of Priem et al. (2020), they do not offer any theoretical analysis. Furthermore, they maximize a variant of the EI criterion as opposed to the upper confidence bound $u_{f,t-1}$ in our approach. Besides, the optimization problem in equation (10) can also be framed as the unconstrained penalized acquisition function in the work of Lu and Paulson (2022). It suffers from an extra penalty parameter requiring automated fine-tuning, which is left as a future work by Lu and Paulson (2022). The algorithm most closely related to our UCB-C is the recent CONFIG algorithm (Xu et al., 2023), which is accompanied by theoretical bounds on the cumulative regret from the objective function and the constraints. However, decoupled queries remain unexplored in these studies. In the next section, we address this scenario by adaptively selecting the *function query* h_t while maintaining the theoretical performance guarantee.

In order for the choice of $\mathbf{x}_t$ in equation (10) to exist, $\mathcal{O}_t$ must be non-empty. This holds with probability $\geq 1-\delta$ according to Lemma 3.1 if the optimization problem is feasible. To avoid the subtle case of $\mathcal{O}_t = \emptyset$, we recommend setting the GP prior mean of constraint c to λ_c in practice.

3.2 FUNCTION QUERY

To begin with, we provide insights into a rational function query selection strategy in the decoupled setting. Then, we will rigorously translate them into a concrete strategy.

[2]The term "vertical" refers to the output of f often plotted as the (vertical) y-axis, as opposed to the term "horizontal" which refers to the input of f often plotted as the (horizontal) x-axis, e.g., in Fig. 1a.

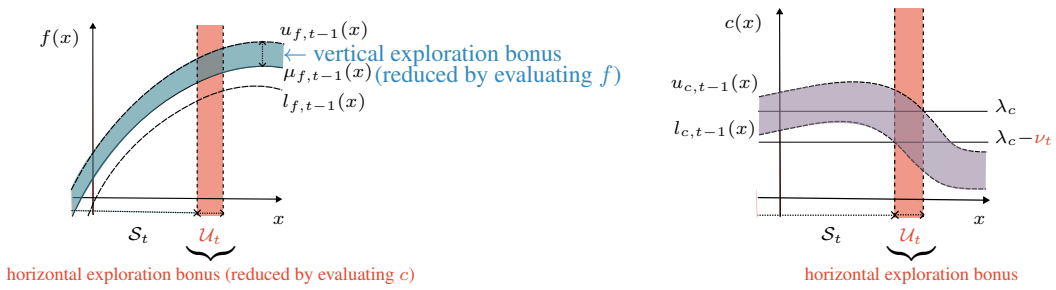

(a) (b)

Figure 1: (a) Plot of the vertical and horizontal exploration bonuses in the space of the objective function; and (b) Plot of the horizontal exploration bonus in the space of the constraint function.

Remark 3.2 (On a rational function query selection strategy). In the decoupled setting, it can be inefficient by querying all functions in $\mathcal{F}$ at every iteration since the possibility of querying/evaluating functions in $\mathcal{F}$ independently is left unexploited. On one hand, it is unnecessary to evaluate any constraint at an input query that is likely to be feasible, which is illustrated in Fig. 2a in Sec. 4. On the other hand, if there is a significant risk of a constraint violation at $\mathbf{x}_t$, it is likely that querying the "most-violated" constraint eliminates $\mathbf{x}_t$ from the feasible region (hence, from being $\mathbf{x}^*$), in which case querying the objective function at $\mathbf{x}_t$ is redundant. It is illustrated in Figs. 2b-c in Sec. 4.

From the regret analysis perspective, the decoupled setting poses a challenge in bounding the instantaneous regret $r(\mathbf{x}_t)$ at each iteration. It is because $r(\mathbf{x}_t)$ depends on all instantaneous regrets $\{r_h(\mathbf{x}_t)\}_{h\in\mathcal{F}}$ while we do not evaluate all functions in $\mathcal{F}$ at each iteration, unlike in the coupled setting. Therefore, it is necessary to establish a connection across $\{r_h(\mathbf{x}_t)\}_{h\in\mathcal{F}}$.

To motivate our choice of h_t, we partition the optimistic feasible region $\mathcal{O}_t$ into a ν_t-*relaxed feasible confidence region* $\mathcal{S}_t$ and an *uncharted region* $\mathcal{U}_t$.

$$\mathcal{O}_t = \mathcal{S}_t \cup \underbrace{\mathcal{U}_t}_{\text{horizontal exploration bonus}} \tag{12}$$

$$\mathcal{S}_t \triangleq \{\mathbf{x}\in\mathcal{X}|\, l_{c,t-1}(\mathbf{x}) \geq \lambda_c - \nu_t \;\forall c\in\mathcal{C}\} \cap \mathcal{O}_t \qquad \text{and} \qquad \mathcal{U}_t \triangleq \mathcal{O}_t \setminus \mathcal{S}_t \tag{13}$$

where $\nu_t \geq 0$ is a constraint-relaxation parameter ($\mathcal{S}_t$ and $\mathcal{U}_t$ are illustrated in Fig. 1b). Recall that $l_{c,t-1}(\mathbf{x}) \leq c(\mathbf{x})$ holds with high probability (Lemma 3.1), so does $\mathcal{S}_t \subset \widetilde{\mathcal{S}}_{\nu_t}$ where $\widetilde{\mathcal{S}}_{\nu_t} \triangleq \{\mathbf{x}\in\mathcal{X}|\, c(\mathbf{x}) \geq \lambda_c - \nu_t \;\forall c\in\mathcal{C}\}$ is a ν_t-relaxation of $\mathcal{S}$. Therefore, $\mathcal{S}_t$ consists of feasible inputs w.r.t. $\widetilde{\mathcal{S}}_{\nu_t}$ with high probability. Furthermore, any input $\mathbf{x}$ in the uncharted region $\mathcal{U}_t$ satisfies

$$\exists c\in\mathcal{C},\; u_{c,t-1}(\mathbf{x}) \geq \lambda_c \wedge \lambda_c - \nu_t > l_{c,t-1}(\mathbf{x})\;. \tag{14}$$

Hence, the uncharted region $\mathcal{U}_t$ consists inputs whose feasibilities w.r.t. $\mathcal{S}$ are unknown (because $u_{c,t-1}(\mathbf{x}) \geq \lambda_c > l_{c,t-1}(\mathbf{x})$) and whose risks of a constraint violation are sufficiently high (because $\lambda_c - l_{c,t-1}(\mathbf{x}) \geq \nu_t$ for some $c\in\mathcal{C}$) as illustrated in Fig. 1b.

We interpret $\mathcal{U}_t$ as the *horizontal exploration bonus*. Its role in the optimistic feasible region $\mathcal{O}_t$ is analogous to the *vertical exploration bonus* $\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x})$ in the optimistic objective function evaluation $u_{f,t-1}(\mathbf{x})$ (illustrated in Fig. 1a and equation (8) vs. equation (12)). Let us translate the 2 intuitive cases in Remark 3.2 into the following concrete conditions.

Querying a constraint when $\mathbf{x}_t \in \mathcal{U}_t$. When $\mathbf{x}_t \in \mathcal{U}_t$, the risk of a constraint violation at $\mathbf{x}_t$ is sufficiently high ($\lambda_c - l_{c,t-1}(\mathbf{x}_t) \geq \nu_t$, illustrated in Fig. 1b). From Remark 3.2, we query the "most-violated" constraint defined as the one with the *highest risk of a constraint violation*, i.e., $\arg\max_{c\in\mathcal{C}} \lambda_c - l_{c,t-1}(\mathbf{x}_t)$. It is noted that as ν_t decreases, the size of the uncharted region $\mathcal{U}_t$ is non-decreasing (see Fig. 1b). Hence, we use ν_t to control the size of $\mathcal{U}_t$, i.e., controlling the horizontal exploration bonus. However, assigning *a small value* to ν_t is risky because we may excessively query a constraint, i.e., *excessive horizontal exploration*. Let us consider an extreme scenario: $\nu_t = 0$ and there exists an iteration t such that a constraint $c\in\mathcal{C}$ is active at $\mathbf{x}_t$, i.e., $c(\mathbf{x}_t) = \lambda_c$. In this case, $\lambda_c - l_{c,t-1}(\mathbf{x}_t) = c(\mathbf{x}_t) - l_{c,t-1}(\mathbf{x}_t) \geq 0$, so the algorithm will keep querying a constraint without querying the objective function.

Algorithm 1 UCB-D

Require: $\mathcal{X}, \mathcal{D}_0$
Update GP posterior beliefs: $\{(\mu_{h,0}, \sigma_{h,0})\}_{h\in\mathcal{F}}$
for $t \leftarrow 1; t \leftarrow t+1; t \leq T$ **do**
  $\mathbf{x}_t \leftarrow \arg\max_{\mathbf{x}\in\mathcal{O}_t} u_{f,t-1}(\mathbf{x})$
  $c_t \leftarrow \arg\max_{c\in\mathcal{C}} \lambda_c - l_{c,t-1}(\mathbf{x}_t)$ `// most-violated constraint`
  if $\lambda_{c_t} - l_{c_t,t-1}(\mathbf{x}_t) > 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$ **then** `//` $\mathbf{x}_t \in \mathcal{U}_t$
    $h_t \leftarrow c_t$ `// query most-violated constraint`
  else `//` $\mathbf{x}_t \in \mathcal{S}_t$
    $h_t \leftarrow f$ `// query objective function`
  end if
  $\mathbf{y}_{h_t}(\mathcal{D}_{h_t,t}) \leftarrow \mathbf{y}_{h_t}(\mathcal{D}_{h_t,t-1}) \cup \{y_{h_t}(\mathbf{x}_t)\}$
  Update GP posterior belief: $\mu_{h_t,t}, \sigma_{h_t,t}$
end for

Querying the objective function when $\mathbf{x}_t \in \mathcal{S}_t$. When $\mathbf{x}_t \in \mathcal{S}_t$, it is likely that $\mathbf{x}_t$ is a feasible solution (relaxed by ν_t), illustrated in Fig. 1b. From Remark 3.2, we prefer querying the objective function. However, assigning *a large value* to ν_t is risky because we may excessively query the objective function, i.e., *insufficient horizontal exploration*. In particular, if $\nu_t > \max_{c\in\mathcal{C}} \max_{\mathbf{x}\in\mathcal{X}} \lambda_c - l_{c,t-1}(\mathbf{x})$, then $\mathcal{U}_t = \emptyset$ and $\mathcal{O}_t = \mathcal{S}_t$. It means $\mathbf{x}_t \in \mathcal{S}_t$ and the algorithm queries the objective function.

To resolve the dilemma of setting ν_t too large (excessively querying the objective function) or too small (excessively querying a constraint), we let ν_t to be "self-tuned" by tying its value with the vertical exploration bonus, i.e., setting $\nu_t = 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$. The broad intuition is that if ν_t is too large, the algorithm repeats querying the objective function which reduces the vertical exploration bonus $\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$. It, in turn, reduces ν_t as $\nu_t = 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$. On the contrary, ν_t is too small only if $2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$ is too small. From equation (11), it implies that $r_f(\mathbf{x}_t)$ is small, so it is justifiable to refrain from querying the objective function. The resulting algorithm, called UCB-D, for the decoupled setting is described in Algorithm 1. In App. C, we prove the following Theorem 3.3 on the cumulative regret of Algorithm 1.

Theorem 3.3. *The cumulative regret R_T of Algorithm 1 is bounded by*

$$Pr\Big\{R_T \leq \sqrt{|\mathcal{F}|T\beta_T \max_{h\in\mathcal{F}} C_h\gamma_{h,T}}\ \forall T \geq 1\Big\} \geq 1-\delta \tag{15}$$

where $C_h \triangleq 8/\log(1+\sigma_h^{-2})$ and $\gamma_{h,T}$ adopted from the work of Srinivas et al. (2010) is the maximum information gain from observing T noisy evaluations of h.

Srinivas et al. (2010) show that $\gamma_{h,T}$ is sublinear for some commonly used kernels including SE and Matérn kernels. Hence, Theorem 3.3 suggests that Algorithm 1 results in a cumulative regret that grows sublinearly when employing GPs with these kernels.

3.3 Function Query From Active Learning Perspective

The choice of the input query $\mathbf{x}_t$ from GP-UCB exhibits an intriguing link to the concept of information gain found in the active learning literature, where one seeks the "most informative data point" or its approximate equivalent, the "most uncertain data point" as discussed in the work of Srinivas et al. (2010). Interestingly, one can view the choice of the function query h_t as an uncertainty sampling strategy as well (i.e., seeking the "most uncertain data point") (Settles, 2009).

Let us denote the upper confidence bounds of the instantaneous regrets w.r.t the objective function f and constraint c in equation (11) as

$$u_{r_f,t-1}(\mathbf{x}_t) \triangleq 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t) \geq r_f(\mathbf{x}_t) \tag{16}$$

$$u_{r_c,t-1}(\mathbf{x}_t) \triangleq \max(0, \lambda_c - l_{c,t-1}(\mathbf{x}_t)) \geq r_c(\mathbf{x}_t)\ . \tag{17}$$

While the nature of the uncertainty in active learning differs from that of the instantaneous regret in our problem, we are interested in minimizing their values in both scenarios. Hence, we employ the

uncertainty sampling paradigm to choose the function query by setting

$$h_t = \arg\max_{h\in\mathcal{F}} u_{r_h,t-1}(\mathbf{x}_t)\ . \tag{18}$$

This resulting strategy, interestingly, coincides with the choice of h_t in Algorithm 1. It is noted that the uncertainty sampling approach may not be effective when there exists uncertainty that cannot be reduced with observations, referred to as *aleatoric uncertainty* in the work of Hüllermeier and Waegeman (2021). Fortunately, our instantaneous regrets are analogous to the *epistemic uncertainty* that can be reduced with observations. Specifically, the more observations we obtain, the better the objective function and constraints (hence, $\mathbf{x}^*$) are estimated. Thus, a smaller regret can be achieved.

Remark 3.4. If the evaluation of $h \in \mathcal{F}$ incurs a cost $l(h) > 0$, then we would like to choose the function query by maximizing a *cost-aware* upper confidence bound of the instantaneous regret. Specifically, the function query is chosen as $h_t = \arg\max_{h\in\mathcal{F}} u_{r_h,t-1}(\mathbf{x}_t)/l(h)$. It is interpreted as the upper confidence bound of the instantaneous regret per unit cost (Swersky et al., 2013).

3.4 Estimator of the Optimal Solution

Remark 2.1 states that $s(\mathbf{x})$ is more effective than $r(\mathbf{x})$ in assessing the suboptimality of a solution. Hence, we would like to use $s(\mathbf{x})$ to propose an estimator $\tilde{\mathbf{x}}_t^*$ for approximating the optimal solution $\mathbf{x}^*$. From equation (11), we obtain an upper confidence bound of $s(\mathbf{x})$ at iteration t

$$s(\mathbf{x}) \triangleq \sum_{h\in\mathcal{F}} r_h(\mathbf{x}) \leq \sum_{h\in\mathcal{F}} u_{r_h,t-1}(\mathbf{x}) \tag{19}$$

where $u_{r_h,t-1}$ is defined in equation (16) and equation (17). We would like to select the input with the lowest upper confidence bound of $s(\mathbf{x})$ as the estimator by considering all previous $t-1$ iterations

$$\tilde{\mathbf{x}}_t^* = \tilde{\mathbf{x}}_{\kappa(t)}\ , \tag{20}$$

where

$$\tilde{\mathbf{x}}_{t'} \triangleq \arg\min_{\mathbf{x}\in\mathcal{X}} \sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x}) \quad \text{and} \quad \kappa(t) \triangleq \arg\min_{t'=1,\ldots,t} \sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\tilde{\mathbf{x}}_{t'})\ . \tag{21}$$

Then, App. D proves the following lemma.

Lemma 3.5. *By picking the estimator in equation (20), it holds with probability* $\geq 1-\delta$ *that*

$$\forall t \geq 1,\ s(\tilde{\mathbf{x}}_t^*) \leq |\mathcal{F}|\sqrt{|\mathcal{F}|\beta_t \max_{h\in\mathcal{F}} C_h\gamma_{h,t}/t} \tag{22}$$

where $C_h \triangleq 8/\log(1+\sigma_h^{-2})$ *and* $\gamma_{h,T}$ *adopted from the work of Srinivas et al. (2010) is the maximum information gain from observing* T *noisy evaluations of* h*.*

Hence, when $\gamma_{h,t}$ is sublinear (e.g., when the kernel is SE or Matérn (Srinivas et al., 2010)), the sum $s(\tilde{x}_t^*)$ of instantaneous regrets at the estimator approaches 0 as $t \to \infty$.

4 Experiments

This section validates the empirical performance of our algorithms (UCB-C in the coupled setting and UCB-D in the decoupled setting) by comparing with EIC (Gardner et al., 2014), ADMMBO (Ariafar et al., 2019), and the state-of-the-art CMES-IBO which significantly outperforms other existing approaches including EIC and PESC in the work of Takeno et al. (2022). We did not conduct a comparison with PESC due to the challenge of maintaining a consistent initial configuration for PESC, as emphasized by Takeno et al. (2022). Moreover, in the decoupled setting, PESC is not currently available in the primary branch of the Spearmint tool at `https://github.com/HIPS/Spearmint`. This also highlights the complexity involved in implementing PESC for decoupled queries and its limited accessibility to practitioners. For these baselines, we select the estimator $\tilde{\mathbf{x}}_t^* = \arg\max_{\mathbf{x}\in\mathcal{X}} \mu_{f,t-1}(\mathbf{x})$ such that $\forall c \in \mathcal{C}, \Pr(c(\mathbf{x}) \geq \lambda_c) \geq \sqrt[|\mathcal{C}|]{0.95}$ as suggested by Takeno et al. (2022). This definition may be undefined in the presence of an equality constraint, so we do not consider equality constraints in our experiments. The estimator in our algorithms is described in equation (20). To illustrate both the instantaneous regrets of the objective function and constraints, we plot the average and standard error (over 10 repeated experiments) of the sum $s(\tilde{\mathbf{x}}_t^*)$ (equation (5)) of these regrets at the estimator against the number of queries $|\mathcal{D}_t|$. The noise's standard deviation is set at $\sigma_h = 0.01\ \forall h \in \mathcal{F}$. Additional details are described in App. E.

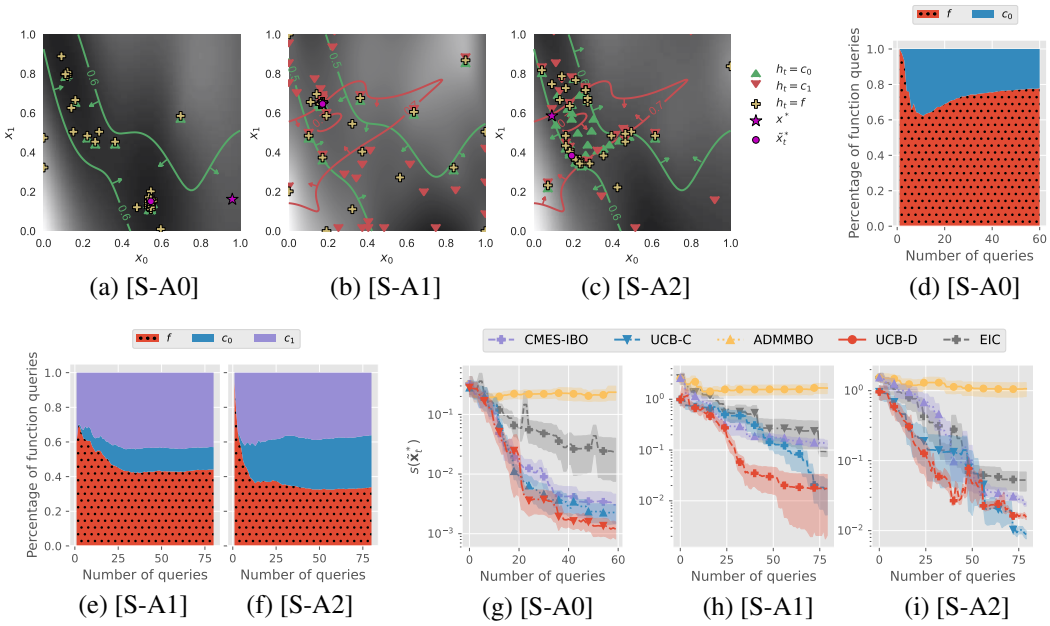

Figure 2: Synthetic problems: (a-c) Plots of the upper confidence bound $u_{f,t-1}$ (plotted as the gray heatmap), constraints (plotted as contour lines with arrows showing the side of the feasible region), the optimal x^*, the estimator $\tilde{\mathbf{x}}_t^*$, and input queries of UCB-D; (d-f) Plots of the percentage of the objective function (the dotted area) and constraints selected as h_t by UCB-D; (g-i) Plots of the instantaneous regret $s(\tilde{\mathbf{x}}_t^*)$ against the number of queries.

4.1 Synthetic Problems

The experiments are conducted on 3 synthetic constrained optimization problems each labeled in the format [S-A{#_of_active_constraints_at_$\mathbf{x}^*$}]: [S-A0], [S-A1], and [S-A2] with 0, 1, and 2 active constraints, respectively (the formulations are described in App. E.1). In these problems, there are 2 input dimensions so we can visualize the constraints and the input queries as shown in Figs. 2a-c. **[S-A0]:** The constraint is inactive at $\mathbf{x}^*$ which is located distant from the boundary of the feasible region $\mathcal{S}$ (Fig. 2a). Thus, UCB-D does not require precise boundary estimations of $\mathcal{S}$ (i.e., of c_0) to pinpoint $\mathbf{x}^*$. This results in a sparse allocation of input queries around the boundary of $\mathcal{S}$. Furthermore, around the estimator $\tilde{\mathbf{x}}_t^*$, a substantial number of queries are evaluated at the objective function (i.e., $h_t = f$, as plotted by the yellow pluses) due to the high certainty that the input query is feasible, which aligns with Remark 3.2. Specifically, Fig. 2d shows that more than 70% of 60 input queries are evaluated at the objective function f, as plotted by the dotted area. Despite the distance between the estimator $\tilde{\mathbf{x}}_t^*$ and the optimal $\mathbf{x}^*$, the difference between $f(\tilde{\mathbf{x}}_t^*)$ and $f(\mathbf{x}^*)$ is minimal because $s(\tilde{\mathbf{x}}_t^*)$ is small in Fig. 2g for UCB-D. **[S-A1]:** At $\mathbf{x}^*$, the constraint c_0 is inactive, but unlike [S-A0], the constraint c_1 is active (Fig. 2b). Thus, it requires precise boundary estimations of c_1 around $\mathbf{x}^*$ to pinpoint $\mathbf{x}^*$, but does not require precise boundary estimations of c_0. This results in a sparse allocation of input queries around the boundary of c_0 and a denser allocation of input queries around the boundary of c_1, especially near $\tilde{\mathbf{x}}_t^*$ in Fig. 2b. Specifically, Fig. 2e shows that only a small number of the input queries are evaluated at the inactive c_0. **[S-A2]:** Both constraints c_0 and c_1 are active at $\mathbf{x}^*$ (Fig. 2c). Thus, UCB-D requires precise boundary estimations of both c_0 and c_1 around $\mathbf{x}^*$ to pinpoint $\mathbf{x}^*$. This results in a dense allocation of input queries around the boundaries of both c_0 and c_1, especially around $\mathbf{x}^*$ in Fig. 2c. Specifically, Fig. 2f shows that the input queries are roughly allocated equally to the objective function and the 2 constraints. Despite the distance between $\tilde{\mathbf{x}}_t^*$ and $\mathbf{x}^*$, the difference between $f(\tilde{\mathbf{x}}_t^*)$ and $f(\mathbf{x}^*)$ is minimal because $s(\tilde{\mathbf{x}}_t^*)$ is small in Fig. 2i for UCB-D. In Figs. 2b-c, we also observe that only a minority of the function queries $h_t = f$ are located far away from the feasible region $\mathcal{S}$, which aligns with Remark 3.2.

Regarding the instantaneous regret, Figs. 2g-i show that our UCB-D converges faster than other algorithms. Hence, UCB-D is more query-efficient, as explained by the above discussion. UCB-C performs competitively compared to the state-of-the-art CMES-IBO as both are designed for the

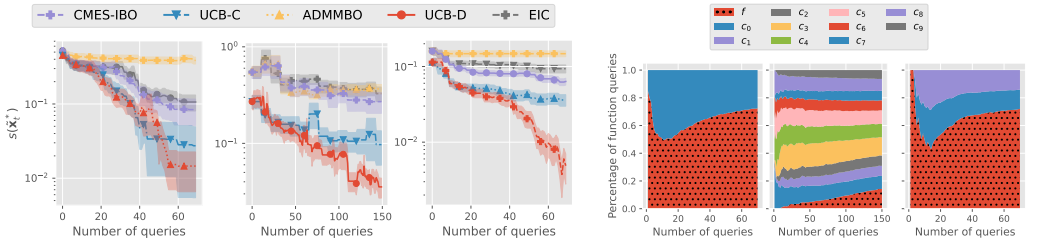

(a) [Gas] (b) [CNN] (c) [QChip] (d) [Gas] (e) [CNN] (f) [QChip]

Figure 3: Real-world problems: Plots of (a-c) the instantaneous regret $s(\tilde{\mathbf{x}}^*)$ and (d-f) the percentage of function queries (by UCB-D) against the number of queries.

coupled setting. ADMMBO does not work well probably because it requires tuning the number of evaluations of f and c at each BO iteration. In our experiments, ADMMBO evaluates f and c once at each BO iteration to be consistent with EIC, CMES-IBO, and UCB-C.

4.2 Real-world Problems

In this section, we introduce 3 optimization problems utilizing real-world objective functions and constraints. These problems serve to assess the effectiveness of our algorithms in practice. We select a real-world problem of optimizing a gas transmission compressor design, referred to as [Gas], from Kumar et al. (2020). It consists of $d = 4$ input dimensions and has $|\mathcal{C}| = 1$ constraint. The problem of tuning hyperparameters of a convolutional neural network (CNN), referred to as [CNN], is taken from the work of Takeno et al. (2022). In the [CNN] problem, a two-layer CNN is trained on a class-imbalanced CIFAR10 dataset. The goal is to maximize the overall accuracy across 10 classes subject to the constraint that the recall of each class is at least 0.5, i.e., $|\mathcal{C}| = 10$. There are $d = 5$ hyperparameters to be optimized. The final experiment, referred to as [QChip], involves maximizing the coupling strength of a synthetic superconducting quantum chip (Yan et al., 2018). While it is a critical aspect for the chip's performance, coupling strength must be carefully controlled within the constraints of the coupling energy to prevent issues like noise, cross-talk between qubits, and poor gate fidelity (Kwon et al., 2021). In particular, we maximize the coupling strength subject to $|\mathcal{C}| = 2$ constraints specifying the desirable range of the coupling energy, by adjusting $d = 11$ geometric features that describe the physical dimensions and arrangement of quantum chip components. To create the groundtruth functions, we obtain a dataset consisting of 393 data points. They are rigorously generated through extensive simulations using electrical simulation software (see App. E.3).

Figs. 3(a-c) show that our UCB-D and UCB-C converge faster than other baseline methods. Therefore, when considering the same number of queries, our algorithms outperform other baseline methods in identifying superior designs for the above real-world problems. In Figs. 3d and 3f, we observe that the number of queries to the objective function dominates that to the constraints in the [Gas] and [QChip] experiments, hinting that the constraints are inactive at the optimal solution, aligning with the groundtruth. Fig. 3e shows that in the [CNN] experiment, UCB-D initially focuses on identifying a feasible input as it allocates few queries to the objective function at the start. It is noted that locating a feasible input is more challenging in this experiment due to the large number of constraints.

5 Conclusion

In this paper, we propose a novel constrained BO algorithm with a provable performance guarantee that adaptively selects not only the input query but also the function query to account for the decoupled query. We formulate the algorithm from the standpoint of the fundamental exploration-exploitation trade-off and, interestingly, cast the proposed algorithm under the uncertainty sampling paradigm in the active learning literature. As our constrained BO solution requires only the confidence bounds of the function evaluations, we believe the approach can be applied to other BO problems such as BO of risk measures (Cakmak et al., 2020; Nguyen et al., 2021b;a) and meta-BO (Nguyen et al., 2023).

REPRODUCIBILITY STATEMENT

We have described detailed proofs for the theoretical results in App. B, C, and D. These proofs utilize an assumption from the work of Srinivas et al. (2010) as elaborated in Sec. 3. Regarding the experimental results, we have included both the code and the datasets in the submission. We have also provided a more detailed description of the experiment settings in App. E.

ACKNOWLEDGEMENT

This research is supported by AI Singapore, under grant AISG2-RP-2020-018.

This research is part of the programme DesCartes and is supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme.

REFERENCES

S. Ariafar, J. Coll-Font, D. H. Brooks, and J. G. Dy. ADMMBO: Bayesian optimization with unknown constraints using ADMM. *JMLR*, 20(123):1–26, 2019.

E. Brochu, V. M. Cora, and N. de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv:1012.2599*, 2010.

S. Cakmak, R. Astudillo, P. I. Frazier, and E. Zhou. Bayesian optimization of risk measures. In *Proc. NeurIPS*, 2020.

W. Chen, S. Liu, and K. Tang. A new knowledge gradient-based method for constrained Bayesian optimization. *arXiv preprint arXiv:2101.08743*, 2021.

P. I. Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018.

P. I. Frazier, W. B. Powell, and S. Dayanik. A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization*, 47(5):2410–2439, 2008.

J. R. Gardner, M. J. Kusner, Z. E. Xu, K. Q. Weinberger, and J. P. Cunningham. Bayesian optimization with inequality constraints. In *Proc. ICML*, volume 2014, pages 937–945, 2014.

Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2022. in preparation.

M. A. Gelbart, J. Snoek, and R. P. Adams. Bayesian optimization with unknown constraints. *arXiv preprint arXiv:1403.5607*, 2014.

J. Ghosh and M. R. Geller. Controlled-not gate with weakly coupled qubits: Dependence of fidelity on the form of interaction. *Physical Review A*, 81(5), 2010.

R. B. Gramacy, G. A. Gray, S. Le Digabel, H. K. Lee, P. Ranjan, G. Wells, and S. M. Wild. Modeling an augmented Lagrangian for blackbox constrained optimization. *Technometrics*, 58(1):1–11, 2016.

M. Gupta, G. V. Graziano, M. Pendharkar, J. T. Dong, C. P. Dempsey, C. Palmstrøm, and V. S. Pribiag. Gate-tunable superconducting diode effect in a three-terminal Josephson device. *Nature Communications*, 14(1), 2023.

T. Hao, R. Shaydulin, M. Pistoia, and J. Larson. Exploiting in-constraint energy in constrained variational quantum optimization. In *2022 IEEE/ACM Third International Workshop on Quantum Computing Software (QCS)*, pages 100–106, 2022.

P. Hennig and C. J. Schuler. Entropy search for information-efficient global optimization. *JMLR*, pages 1809–1837, 2012.

J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In *Proc. NIPS*, pages 918–926, 2014.

J. M. Hernández-Lobato, M. A. Gelbart, R. P. Adams, M. W. Hoffman, and Z. Ghahramani. A general framework for constrained Bayesian optimization using information-based search. *JMLR*, 17(1): 5549–5601, 2016.

E. Hüllermeier and W. Waegeman. Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning*, 110:457–506, 2021.

A. Kumar, G. Wu, M. Z. Ali, R. Mallipeddi, P. N. Suganthan, and S. Das. A test-suite of non-convex constrained optimization problems from the real-world and some baseline results. *Swarm and Evolutionary Computation*, 56:100693, 2020.

H. J. Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of basic engineering*, 86(1):97–106, 1964.

Sangil Kwon, Akiyoshi Tomonaga, Gopika Lakshmi Bhai, Simon J. Devitt, and Jaw-Shen Tsai. Gate-based superconducting quantum computing. *Journal of Applied Physics*, 129(4), 2021.

Fei-Yu Li and Li-Jing Jin. Quantum chip design optimization and automation in superconducting coupler architecture. *Quantum Science and Technology*, 8(4):045015, 2023.

B. Lienhard, J. Braumuller, W. Woods, D. Rosenberg, G. Calusine, S. Weber, A. Vepsalainen, K. O'Brien, T. P. Orlando, S. Gustavsson, and W. D. Oliver. Microwave packaging for superconducting qubits. In *2019 IEEE MTT-S International Microwave Symposium (IMS)*, 2019.

Q. Liu, Y. Huang, Y. Du, Z. Zhao, M. Geng, Z. Zhang, and K. Wei. Advances in chip-based quantum key distribution. *Entropy*, 24:1334, 2022.

C. Lu and J. A. Paulson. No-regret Bayesian optimization with unknown equality and inequality constraints using exact penalty functions. *IFAC-PapersOnLine*, 55(7):895–902, 2022.

X. Y. Lu, S. Ashhab, W. Cui, R. Wu, and F. Nori. Two-qubit gate operations in superconducting circuits with strong coupling and weak anharmonicity. *New Journal of Physics*, 14:073041, 2012.

D. A. B. Miller. Optical physics of quantum wells. Technical report, Stanford University, 1997.

J. Mockus, V. Tiešis, and A. Žilinskas. The application of Bayesian methods for seeking the extremum. In L. C. W. Dixon and G. P. Szegö, editors, *Towards Global Optimization 2*, pages 117–129. North-Holland Publishing Company, 1978.

Q. P. Nguyen, Z. Dai, B. K. H. Low, and P. Jaillet. Optimizing conditional value-at-risk of black-box functions. In *Proc. NeurIPS*, pages 4170–4180, 2021a.

Q. P. Nguyen, Z. Dai, B. K. H. Low, and P. Jaillet. Value-at-risk optimization with Gaussian processes. In *Proc. ICML*, pages 8063–8072, 2021b.

Q. P. Nguyen, B. K. H. Low, and P. Jaillet. Meta-VBO: Utilizing prior tasks in optimizing risk measures with Gaussian processes. In *Proc. ICLR*, 2023.

V. Perrone, H. Shen, A. Zolic, I. Shcherbatyi, A. Ahmed, T. Bansal, M. Donini, F. Winkelmolen, R. Jenatton, J. B. Faddoul, et al. Amazon SageMaker automatic model tuning: Scalable black-box optimization. *arXiv preprint arXiv:2012.08489*, 2020.

V. Picheny, R. B. Gramacy, S. Wild, and S. Le Digabel. Bayesian optimization under mixed constraints with a slack-variable augmented Lagrangian. In *Proc. NeurIPS*, volume 29, 2016.

R. Priem, N. Bartoli, Y. Diouane, and A. Sgueglia. Upper trust bound feasibility criterion for mixed constrained Bayesian optimization with application to aircraft design. *Aerospace Science and Technology*, 105:105980, 2020.

C. E. Rasmussen and C. K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.

M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout. Detecting crosstalk errors in quantum information processors. *Quantum*, 4:321, 2020.

A. M. Schweidtmann, A. D. Clayton, N. Holmes, E. Bradford, R. A. Bourne, and A. A. Lapkin. Machine learning meets continuous flow chemistry: Automated optimization towards the Pareto front of multiple objectives. *Chemical Engineering Journal*, 352:277–282, 2018.

E. A. Sete, N. Didier, A. Q. Chen, S. Kulshreshtha, R. Manenti, and S. Poletto. Parametric-resonance entangling gates with a tunable coupler. *Physical Review Applied*, 16(2), 2021.

B. Settles. Active learning literature survey. 2009.

N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. ICML*, pages 1015–1022, 2010.

K. Swersky, J. Snoek, and R. P. Adams. Multi-task Bayesian optimization. In *Proc. NeurIPS*, volume 26, 2013.

S. Takeno, T. Tamura, K. Shitara, and M. Karasuyama. Sequential and parallel constrained max-value entropy search via information lower bound. In *Proc. ICML*, pages 20960–20986, 2022.

Z. Wang and S. Jegelka. Max-value entropy search for efficient Bayesian optimization. In *Proc. ICML*, pages 3627–3635, 2017.

M. Wistuba, N. Schilling, and L. Schmidt-Thieme. Scalable Gaussian process-based transfer surrogates for hyperparameter optimization. *Machine Learning*, 107(1):43–78, 2018.

Y. Wu, W. S. Bao, S. Cao, F. Chen, M. C. Chen, X. Chen, T. H. Chung, H. Deng, Y. Du, D. Fan, et al. Strong quantum computational advantage using a superconducting quantum processor. *Physical review letters*, 127(18):180501, 2021.

W. Xu, Y. Jiang, B. Svetozarevic, and C. Jones. Constrained efficient global optimization of expensive black-box functions. In *Proc. ICML*, pages 38485–38498, 2023.

F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando, S. Gustavsson, and W. D. Oliver. Tunable coupling scheme for implementing high-fidelity two-qubit gates. *Physical Review Applied*, 10(5), 2018.

C. H. Yang, R. C. C. Leon, J. C. C. Hwang, A. Saraiva, T. Tanttu, W. Huang, J. Camirand Lemyre, K. W. Chan, K. Y. Tan, F. E. Hudson, et al. Operation of a silicon quantum processor unit cell above one kelvin. *Nature*, 580(7803):350–354, 2020.

A On Issue with Equality Constraint

Several existing works including EIC (Gardner et al., 2014) and CMES-IBO (Takeno et al., 2022) rely on the probability mass $P(c(\mathbf{x}) \geq \lambda_c \; \forall c \in \mathcal{C})$ where the unknown constraint c is modeled with a GP. Let us consider an optimization problem where there is only 1 equality constraint $c(\mathbf{x}) = 0$. Then, the above probability mass $P(c(\mathbf{x}) = 0)$ is 0. As a result, existing works such as CMES-IBO require manually adding a tolerance value to ensure the probability mass is strictly positive.

Regarding PESC (Hernández-Lobato et al., 2016), it relies on the sampling of $\mathbf{x}^*$. When there are several equality constraints, the probability of obtaining a feasible solution is small, making the sampling inefficient.

In our approach, we consider an optimistic feasible region (which is partitioned into $\mathcal{S}_t$ and $\mathcal{U}_t$ shown in Fig. 1b) that can handle equality constraints without any modification. Though one may argue that the role of ν_t is similar to the tolerance added to the equality constraint in CMES-IBO, it is noted that ν_t is not fixed to any pre-defined sequence but it is "self-tuned" by the vertical exploration bonus as explained in Sec. 3.2 and has a theoretical justification.

When proposing an estimator for approximating the optimal solution, one may consider a pessimistic feasible region to ensure the feasbility of the estimator with high probability. However, the pessimistic feasible region can be empty if the optimal solution is not an interior point of the feasible region, e.g., when there are equality constraints.

B A Bayesian Optimization Algorithm for Coupled Queries

Algorithm 2 UCB-C

Require: $\mathcal{X}$, $\mathcal{D}_0$
1: Update GP posterior beliefs: $\{(\mu_{h,0}, \sigma_{h,0})\}_{h \in \mathcal{F}}$
2: **for** $t \leftarrow 1; t \leftarrow t+1; t \leq T$ **do**
3: $\quad \mathbf{x}_t \leftarrow \arg\max_{\mathbf{x} \in \mathcal{O}_t} u_{f,t-1}(\mathbf{x})$
4: $\quad$ **for** $h \in \mathcal{F}$ **do** `// coupled query`
5: $\quad\quad \mathbf{y}_h(\mathcal{D}_{h,t}) \leftarrow \mathbf{y}_h(\mathcal{D}_{h,t-1}) \cup \{y_h(\mathbf{x}_t)\}$
6: $\quad\quad$ Update GP posterior belief: $\mu_{h,t}, \sigma_{h,t}$
7: $\quad$ **end for**
8: **end for**

Our proposed algorithm for coupled queries is shown in Algorithm 2. In this section, we assume that $h(\mathbf{x}) \in [l_{h,t-1}(\mathbf{x}), u_{h,t-1}(\mathbf{x})]$ for all $h \in \mathcal{F}$, $\mathbf{x} \in \mathcal{X}$, and $t \geq 1$, which happens with probability $\geq 1 - \delta$ from Lemma 3.1. In order to prove the upper confidence bound of its cumulative regret, we derive the inequalities in equation (11) as follows.

The instantaneous regret w.r.t. the objective function f is bounded by

$$
\begin{aligned}
r_f(\mathbf{x}_t) &\triangleq \max(0, f(\mathbf{x}^*) - f(\mathbf{x}_t)) && (23)\\
&\leq \max(0, u_{f,t-1}(\mathbf{x}^*) - l_{f,t-1}(\mathbf{x}_t)) \quad \text{from Lemma 3.1} && (24)\\
&\leq \max(0, u_{f,t-1}(\mathbf{x}_t) - l_{f,t-1}(\mathbf{x}_t)) && (25)\\
&= u_{f,t-1}(\mathbf{x}_t) - l_{f,t-1}(\mathbf{x}_t) && (26)\\
&= 2\beta_t^{1/2} \sigma_{f,t-1}(\mathbf{x}_t) && (27)
\end{aligned}
$$

where inequality equation (25) holds with probability $\geq 1 - \delta$ as

- $\mathbf{x}^* \in \mathcal{O}_t$ with probability $\geq 1 - \delta$ since $\mathcal{O}_t \supset \mathcal{S}$ with probability $\geq 1 - \delta$.
- $u_{f,t-1}(\mathbf{x}_t) \geq u_{f,t-1}(\mathbf{x})$ for all $\mathbf{x} \in \mathcal{O}_t$ because $\mathbf{x}_t \triangleq \arg\max_{\mathbf{x} \in \mathcal{O}_t} u_{f,t-1}(\mathbf{x})$.

The instantaneous regret w.r.t. the constraint c is bounded by

$$
\begin{aligned}
r_c(\mathbf{x}_t) &\triangleq \max(0, \lambda_c - c(\mathbf{x}_t)) && (28)\\
&\leq \max(0, \lambda_c - l_{c,t-1}(\mathbf{x}_t)) \quad \text{from Lemma 3.1} && (29)\\
&\leq \max(0, u_{c,t-1}(\mathbf{x}_t) - l_{c,t-1}(\mathbf{x}_t)) && (30)\\
&= u_{c,t-1}(\mathbf{x}_t) - l_{c,t-1}(\mathbf{x}_t) && (31)\\
&= 2\beta_t^{1/2}\sigma_{c,t-1}(\mathbf{x}_t) && (32)
\end{aligned}
$$

where inequality equation (30) holds with probability $\geq 1-\delta$ as $u_{c,t-1}(\mathbf{x}_t) \geq \lambda_c$ because $\mathbf{x}_t \in \mathcal{O}_t$. Therefore,

$$
\begin{aligned}
r(\mathbf{x}_t) &\triangleq \max_{h\in\mathcal{F}} r_h(\mathbf{x}_t) \leq \max_{h\in\mathcal{F}} 2\beta_t^{1/2}\sigma_{h,t-1}(\mathbf{x}_t) && (33)\\
R_T &\triangleq \sum_{t=1}^{T} r(\mathbf{x}_t) \leq \sum_{t=1}^{T} \max_{h\in\mathcal{F}} 2\beta_t^{1/2}\sigma_{h,t-1}(\mathbf{x}_t)\,. && (34)
\end{aligned}
$$

Let $T_h \triangleq \sum_{t=1}^{T} \mathbb{1}_{h=\arg\max_{h'\in\mathcal{F}} \sigma_{h',t-1}(\mathbf{x}_t)}$ (breaking the tie arbitrarily if necessary to ensure a unique maximizer) which implies that $\sum_{h\in\mathcal{F}} T_h = T$.

$$
\begin{aligned}
R_T &\leq \sum_{t=1}^{T} \max_{h\in\mathcal{F}} 2\beta_t^{1/2}\sigma_{h,t-1}(\mathbf{x}_t) && (35)\\
&\leq \sum_{h\in\mathcal{F}} 2\beta_T^{1/2} \sum_{t=1}^{T_h} \sigma_{h,t-1}(\mathbf{x}_t) && (36)
\end{aligned}
$$

as β_t is non-decreasing. Furthermore, from Lemma 5.4 in Srinivas et al. (2010):

$$
\sum_{t=1}^{T_h} \sigma^2_{h,t-1}(\mathbf{x}_t) \leq C_h \gamma_{h,T_h}/4 \tag{37}
$$

where $C_h \triangleq 8/\log(1+\sigma_h^{-2})$ and γ_{h,T_h} is the maximum information gain from observing T_h noisy evaluations of h. Therefore, applying the Cauchy-Schwarz inequality,

$$
\sum_{t=1}^{T_h} \sigma_{h,t-1}(\mathbf{x}_t) \leq \sqrt{T_h \sum_{t=1}^{T_h} \sigma^2_{h,t-1}(\mathbf{x}_t)} \leq \sqrt{T_h C_h \gamma_{h,T_h}/4}\,. \tag{38}
$$

Hence,

$$
\begin{aligned}
R_T &\leq \sum_{h\in\mathcal{F}} 2\beta_T^{1/2}\sqrt{T_h C_h \gamma_{h,T_h}/4}\\
&\leq \beta_T^{1/2} \sum_{h\in\mathcal{F}} \sqrt{T_h}\sqrt{C_h\gamma_{h,T_h}}\\
&\leq \beta_T^{1/2} \sqrt{\Big(\sum_{h\in\mathcal{F}} T_h\Big)\Big(\sum_{h'\in\mathcal{F}} C_{h'}\gamma_{h',T_{h'}}\Big)} \quad \text{Cauchy-Schwarz inequality}\\
&= \sqrt{T\beta_T \sum_{h'\in\mathcal{F}} C_{h'}\gamma_{h',T_{h'}}}\\
&\leq \sqrt{T\beta_T|\mathcal{F}| \max_{h\in\mathcal{F}} C_h\gamma_{h,T_h}}\\
&\leq \sqrt{|\mathcal{F}|T\beta_T \max_{h\in\mathcal{F}} C_h\gamma_{h,T}}\,.
\end{aligned}
$$

C PROOF OF THEOREM 3.3

From Algorithm 1, we observe that:

Case 1. When $h_t = f$, we show that the instantaneous regrets w.r.t. the objective function and constraints are bounded by $2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$ as follows.

$$r_f(\mathbf{x}_t) \leq 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t) \quad \text{from equation (11)} .$$

Recall that we select $h_t = f$ when $\mathbf{x}_t \in \mathcal{S}_t$, i.e., for all $c \in \mathcal{C},\ \lambda_c - l_{c,t-1}(\mathbf{x}_t) \leq \nu_t = 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$, so

$$\begin{aligned}
\forall c \in \mathcal{C}, r_c(\mathbf{x}_t) &\triangleq \max(0, \lambda_c - c(\mathbf{x}_t)) \\
&\leq \max(0, \lambda_c - l_{c,t-1}(\mathbf{x}_t)) \quad \text{from Lemma 3.1} \\
&\leq \max(0, 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)) \\
&= 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t) .
\end{aligned}$$

Therefore, when $h_t = f$,

$$r(\mathbf{x}_t) \triangleq \max_{h\in\mathcal{F}} r_h(\mathbf{x}_t) \leq 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t) . \tag{39}$$

Case 2. When $h_t = c$, we show that the instantaneous regrets w.r.t. the objective function and constraints are bounded by $2\beta_t^{1/2}\sigma_{c,t-1}(\mathbf{x}_t)$ as follows.

$$\forall c \in \mathcal{C}, r_c(\mathbf{x}_t) \leq 2\beta_t^{1/2}\sigma_{c,t-1}(\mathbf{x}_t) \quad \text{from equation (11)} .$$

Recall that we select $h_t = c$ when $\mathbf{x}_t \in \mathcal{U}_t$, i.e., $\exists c \in \mathcal{C},\ \lambda_c - l_{c,t-1}(\mathbf{x}_t) > \nu_t = 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t)$. This implies that

$$\begin{aligned}
2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t) &< \max_{c\in\mathcal{C}} \lambda_c - l_{c,t-1}(\mathbf{x}_t) \\
&= \lambda_{c_t} - l_{c_t,t-1}(\mathbf{x}_t) \quad \text{where } c_t \text{ is defined in Algorithm 1} .
\end{aligned}$$

$$\begin{aligned}
r_f(\mathbf{x}_t) &\leq 2\beta_t^{1/2}\sigma_{f,t-1}(\mathbf{x}_t) \quad \text{from equation (11)} \\
&< \lambda_{c_t} - l_{c_t,t-1}(\mathbf{x}_t) \\
&\leq u_{c_t,t-1}(\mathbf{x}_t) - l_{c_t,t-1}(\mathbf{x}_t) \quad \text{as } \mathbf{x}_t \in \mathcal{O}_t, \text{ i.e., } u_{c,t-1}(\mathbf{x}_t) \geq \lambda_c\ \forall c \in \mathcal{C} \\
&= 2\beta_t^{1/2}\sigma_{c_t,t-1}(\mathbf{x}_t) .
\end{aligned}$$

Therefore, when $h_t = c$,

$$r(\mathbf{x}_t) \triangleq \max_{h\in\mathcal{F}} r_h(\mathbf{x}_t) \leq 2\beta_t^{1/2}\sigma_{c,t-1}(\mathbf{x}_t) . \tag{40}$$

Combining the above 2 cases in equation (39) and equation (40),

$$r(\mathbf{x}_t) \leq 2\beta_t^{1/2}\sigma_{h_t,t-1}(\mathbf{x}_t) . \tag{41}$$

Hence,

$$R_T \triangleq \sum_{t=1}^{T} r(\mathbf{x}_t) \leq \sum_{t=1}^{T} 2\beta_t^{1/2}\sigma_{h_t,t-1}(\mathbf{x}_t) . \tag{42}$$

Let $T'_h \triangleq \sum_{t=1}^{T} \mathbb{1}_{h=h_t}$, then using the non-decreasing property of β_t, we can rewrite the above inequality as

$$R_T \leq \sum_{h\in\mathcal{F}} 2\beta_T^{1/2} \sum_{t=1}^{T'_h} \sigma_{h,t-1}(\mathbf{x}_t) . \tag{43}$$

The above result is the same as equation (36). Therefore, we can follow the argument in App. B to obtain the same upper confidence bound of R_T:

$$R_T \leq \sqrt{|\mathcal{F}| T \beta_T \max_{h\in\mathcal{F}} C_h \gamma_{h,T}} . \tag{44}$$

D PROOF OF LEMMA 3.5

App. B and C both show that with probability $\geq 1-\delta$,

$$\forall t \geq 1,\ \sum_{t'=1}^{t} \max_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x}_{t'}) \leq \sqrt{|\mathcal{F}| t \beta_t \max_{h\in\mathcal{F}} C_h \gamma_{h,t}} \tag{45}$$

where $u_{r_h,t-1}(\mathbf{x})$ is defined in equation (16) and equation (17). Equivalently,

$$\begin{aligned}
\sqrt{|\mathcal{F}|\beta_t \max_{h\in\mathcal{F}} C_h\gamma_{h,t}/t} &\geq \frac{1}{t}\sum_{t'=1}^{t} \max_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x}_{t'}) && (46)\\
&\geq \frac{1}{t}\sum_{t'=1}^{t} \frac{1}{|\mathcal{F}|}\sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x}_{t'}) && (47)\\
&\geq \frac{1}{t}\sum_{t'=1}^{t} \min_{\mathbf{x}\in\mathcal{X}} \frac{1}{|\mathcal{F}|}\sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x}) && (48)\\
&\geq \min_{t'=1,\ldots,t} \min_{\mathbf{x}\in\mathcal{X}} \frac{1}{|\mathcal{F}|}\sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x})\ . && (49)
\end{aligned}$$

Let

$$\tilde{\mathbf{x}}_{t'} \triangleq \arg\min_{\mathbf{x}\in\mathcal{X}} \sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x}) \tag{50}$$

$$\kappa(t) \triangleq \arg\min_{t'=1,\ldots,t} \sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\tilde{\mathbf{x}}_t)\ . \tag{51}$$

Our estimator is chosen as

$$\tilde{\mathbf{x}}^*_t = \tilde{\mathbf{x}}_{\kappa(t)}\ , \tag{52}$$

then

$$\min_{t'=1,\ldots,t} \min_{\mathbf{x}\in\mathcal{X}} \frac{1}{|\mathcal{F}|}\sum_{h\in\mathcal{F}} u_{r_h,t'-1}(\mathbf{x}) = \frac{1}{|\mathcal{F}|}\sum_{h\in\mathcal{F}} u_{r_h,\kappa(t)-1}(\tilde{\mathbf{x}}^*_t) \geq \frac{1}{|\mathcal{F}|}\sum_{h\in\mathcal{F}} r_h(\tilde{\mathbf{x}}^*_t) = \frac{s(\tilde{\mathbf{x}}^*_t)}{|\mathcal{F}|} \tag{53}$$

where the inequality holds with probability $\geq 1-\delta$. Hence,

$$\Pr\Big\{ s(\tilde{\mathbf{x}}^*_t) \leq |\mathcal{F}| \sqrt{|\mathcal{F}|\beta_t \max_{h\in\mathcal{F}} C_h\gamma_{h,t}/t} \Big\} \geq 1-\delta\ . \tag{54}$$

E ADDITIONAL EXPERIMENT DETAILS

We refrain from conducting a comparison with EPBO and PESC for the following reasons. EPBO is equivalent to our UCB-C method when an appropriate value of ρ is chosen. However, an automated strategy of selecting ρ is left unspecified in the work of Lu and Paulson (2022). On the other hand, the noteworthy advantage of our proposed UCB-C approach is its ability to overcome the need of selecting such a parameter. Regarding PESC in the coupled setting, Takeno et al. (2022) raises the difficulty in implementing and assigning the initial configuration in the PESC package (Spearmint). This difficulty poses obstacles to achieving consistent initial experiment configurations for PESC. Regarding PESC in the decoupled setting, it is not included in the main branch of Spearmint. Therefore, we opt not to use PESC as a baseline in the experiment. On the other hand, the work of Takeno et al. (2022) shows that CMES-IBO outperforms PESC by a large margin in the coupled setting. Hence, we demonstrate the performance of our algorithms by comparing with CMES-IBO.

To ensure that the global optimal solution can be identified and to prevent variations in performance among different methods stemming from being trapped in distinct local optima when using continuous optimization tools, we discretize the input space into $10{,}000$ randomly selected input points in the experiments: [S-A0], [S-A1], [S-A2], [Gas], and [Beam]. For the [CNN] and [QChip] experiments,

the size of the input space matches that of the generated dataset, i.e., 5120 and 393, respectively. In practice, our proposed methods can be implemented in a continuous input domain using any continuous constrained optimization package because only step 3 (selecting $\mathbf{x}_t$) in both Algorithms 1 and 2 involves solving a constrained optimization problem (with a known objective function and constraint).

The GP hyperparameters including the parameters of the SE kernel and the noise variance $\{\sigma_h^2\}_{h\in\mathcal{F}}$ are assumed to be unknown in our experiments. We optimize them after every BO iteration by maximizing the likelihood of the observations using Adam optimizer. We assign prior distributions to the GP hyperparameters to avoid numerical issues when performing the likelihood maximization. The prior distributions of the length-scale, the signal standard deviation, and the noise standard deviation are Gamma$(0.25, 0.5)$, Gamma$(2, 0.15)$, and $\mathcal{N}(0.0, 0.1)$, respectively. Furthermore, the noise standard deviation is constrained to be at least 0.01.

To ensure the initial observations are consistent with the coupled setting, we initialize all experiments with coupled observations, i.e., the evaluations of the objective function and constraints are at the same set of inputs: $\mathcal{D}_{f,0} = \mathcal{D}_{c,0}$ for all $c \in \mathcal{C}$. In the synthetic experiments, the number of initial coupled observations are 3, 5, and 5 for [S-A0], [S-A1], and [S-A2], respectively. In the real-world experiments, the number of initial coupled observations are 7, 30, and 7 for [Gas], [CNN], and [QChip], respectively.

Throughout the remainder of this section, we provide more detailed descriptions of several experiments.

E.1 Synthetic Experiments

Let g_b and g_g denote the Branin-Hoo and the Goldstein-Price functions where the input domain is normalized to range $[0, 1]^2$. They are obtained from `https://www.sfu.ca/~ssurjano`.

Then, the [S-A0] problem is defined as

$$\max_{\mathbf{x}} \; g_b(\mathbf{x}) \text{ s.t. } g_b(\mathbf{x}) \geq 0.6 \; .$$

The [S-A1] problem is defined as

$$\begin{aligned} \max_{\mathbf{x}} \; & g_b(\mathbf{x}) \\ \text{s.t. } & g_b(\mathbf{x}) \geq 0.5 \\ & g_g(\mathbf{x}) \geq 0.7 \; . \end{aligned}$$

The [S-A2] problem is defined as

$$\begin{aligned} \max_{\mathbf{x}} \; & g_b(\mathbf{x}) \\ \text{s.t. } & g_b(\mathbf{x}) \leq 0.6 \\ & g_g(\mathbf{x}) \geq 0.7 \; . \end{aligned}$$

It is noted that although the same function is used in both the objective function and a constraint in the above optimization problems, we treat them as distinct black-box functions and model them with independent GPs.

E.2 Gas Transmission Compressor Design (Kumar et al., 2020)

There are $d = 4$ input dimension: $\mathbf{x} = (x_i)_{i=1}^4$ with the following bounds:

$$\begin{aligned} 20 \leq x_1 \leq 50 \\ 1 \leq x_2 \leq 10 \\ 20 \leq x_3 \leq 50 \\ 0.1 \leq x_4 \leq 60 \; . \end{aligned}$$

The objective function and the constraint are specified as follows.

$$f(\mathbf{x}) = 8.16 \times 10^5 x_1^{1/2} x_2 x_3^{-2/3} x_4^{-1/2} + 3.69 \times 10^4 x_3 + 7.72 \times 10^8 x_1^{-1} x_2^{0.219} - 765.43 \times 10^6 x_1^{-1}$$
$$c(\mathbf{x}) = x_4 x_2^{-2} + x_2^{-2} - 1 \leq 0 \, .$$

We normalize the range of the objective function and the constraint to the range $[-1, 1]$ and the input to to domain $[0, 1]^4$.

E.3 Maximizing Coupling Strength of Synthetic Superconducting Quantum Chip

We address the optimization problem of fine-tuning the critical parameter of coupling strength in a synthetic superconducting quantum chip (Yan et al., 2018). This optimization problem is motivated by the fundamental role that coupling strength plays in influencing the chip's performance (Sete et al., 2021; Wu et al., 2021). Specifically, the coupling strength affects the speed and fidelity of quantum operations, making it a critical factor in quantum chip design (Lu et al., 2012). However, this parameter must be carefully controlled within the constraints of coupling energy to avoid issues such as noise, cross-talk between qubits, and poor gate fidelity (Kwon et al., 2021).

The objective of our optimization problem is to maximize the coupling strength while adhering to the constraint imposed by the coupling energy. We aim to find the optimal configuration of the quantum chip that achieves the highest coupling strength possible within the predefined bounds of the energy constraint. To achieve this, we use a dataset comprising 393 data points, each describing the geometric and electrical features of the quantum chip. These features are utilized to calculate the energy associated with each chip configuration, including the coupling strength. (Li and Jin, 2023).

Generating this dataset is practical because it captures the variability in quantum chip designs, materials, and operational conditions. Furthermore, it allows us to explore and optimize the coupling strength systematically, which is crucial for improving the performance and efficiency of quantum computations (Liu et al., 2022; Miller, 1997). The choice of the energy feature, denoted as Eq_1q_2, as the target for constraint optimization is based on experimental insights and considerations, which may vary for different quantum chips (Li and Jin, 2023). By maintaining precise control over this energy feature within the predefined range, we demonstrate the effectiveness of our proposed optimization method in achieving our objectives in the context of synthetic quantum chip design.

E.4 Dataset Collection

We initiate our experiment by curating a comprehensive dataset comprising 393 data points. Each data point encompasses eleven geometric features describing the physical dimensions and layout of quantum chip components. These are complemented by four electrical features obtained through rigorous simulations utilizing standard electrical simulation software. The energy associated with each quantum chip configuration is derived from these four electrical features. Additionally, we calculate the coupling strength for each specific quantum chip layout using the energy and reference frequency.

E.5 Model Architecture

In quantum chip design, the coupling strength between qubits is a critical parameter. It determines the rate at which qubits can exchange quantum information, and thus influences the speed and fidelity of quantum operations (Lu et al., 2012).

However, the coupling strength must be carefully controlled under the constraints of coupling energy. If the coupling energy is too high (i.e., the qubits are too strongly coupled), it can lead to unwanted effects such as noise or cross-talk between qubits (Kwon et al., 2021). Cross-talk is a phenomenon where a signal transmitted on one qubit influences another qubit, leading to errors in quantum operations (Sarovar et al., 2020).

On the other hand, if the coupling energy is too low (i.e., the qubits are weakly coupled), it can result in poor gate fidelity (Ghosh and Geller, 2010). Gate fidelity is a measure of how accurately quantum gates (the basic operations of a quantum computer) can be implemented. If the gate fidelity is low, the output of a quantum computation may be unreliable (Ghosh and Geller, 2010).

Therefore, in the context of quantum chip design, achieving a higher coupling strength within a certain constraint of coupling energy is often desirable. It allows for fast and accurate quantum operations while avoiding the problems associated with too much or too little coupling.

Precise control over the energy feature within specific bounds is pivotal in optimizing the coupling strength during the quantum chip design process for several reasons:

1. **Optimal Coupling Strength**: The coupling strength between qubits in a superconducting quantum chip is intricately linked to the system's energy feature. Meticulous management of these parameters, closely tied to the energy feature, enables the attainment of optimal coupling strength. This stands as a crucial factor in ensuring the high fidelity of two-qubit gates (Li and Jin, 2023).

2. **Energy Efficiency**: Devices such as superconducting diodes can achieve enhanced energy efficiency when equipped with a series of gates to control the energy flow (Gupta et al., 2023)[3].

3. **Suppression of Energy Loss Channels**: The choice of material and geometric design of the sample plays a pivotal role in minimizing qubit energy loss channels[4] (Lienhard et al., 2019).

The constraints associated with the energy feature in a quantum chip hold significant sway over the performance and viability of quantum computations. These constraints exhibit variability based on the specific quantum chip under consideration, owing to differences in design, materials, and operational conditions (Yang et al., 2020; Hao et al., 2022; Kwon et al., 2021)

For instance, most quantum computers under global development only operate at fractions of a degree above absolute zero, necessitating multi-million-dollar refrigeration (Yang et al., 2020). Nevertheless, researchers have pioneered a proof-of-concept quantum processor unit cell that functions at 1.5 Kelvin, which is 15 times warmer than competing chip-based technologies (Yang et al., 2020). This elevated operating temperature holds promise for more cost-effective and robust quantum computers (Yang et al., 2020).

Furthermore, the incorporation of complex constraints poses a central challenge when applying near-term quantum optimization algorithms to industrially relevant problems (Hao et al., 2022). In general, such constraints cannot be easily encoded in the circuit, and there is no guarantee that quantum circuit measurement outcomes will adhere to these constraints (Hao et al., 2022). Consequently, novel approaches for solving constrained optimization problems with unconstrained, readily implementable quantum ansatze are being proposed (Hao et al., 2022).

However, it's imperative to acknowledge that the appropriate constraints for the energy feature can fluctuate depending on the specific quantum chip under scrutiny. These constraints may necessitate experimental observations for determination, with their range often rooted in prior experimental insights. This is because the energy feature's range can exert a substantial influence on the performance and feasibility of quantum computations (Liu et al., 2022; Miller, 1997).

In our synthetic quantum chip example, we opted to focus on one specific energy feature, denoted as Eq_1q_2, as the target for constraint optimization. In practical applications, the range for this energy feature should be selected based on prior experimental insights and considerations, which may vary for each unique chip. Nevertheless, for our synthetic chip, we delineated the range for this energy feature ourselves, constraining it within the bounds of 30×10^{-23}J to 45×10^{-23}J.

By maintaining precise control over the energy feature within this pre-defined range, we systematically explored and optimized the coupling strength for the synthetic superconducting quantum chip. This experiment serves as a compelling demonstration of the efficacy of our proposed method in achieving our optimization objectives.

[3]This feature has not been integrated into superconducting diodes previously.

[4]Qubits can spontaneously dissipate energy through dielectric defects on the surface and interfaces of the sample or by coupling to unwanted package modes.

E.6 Extraction of Features and Computation

In this study, we extract a set of 11 geometric features that encompass various aspects of the quantum chip, such as parameters related to its length, width, the gap between components (qubit and coupler), and the boundary layer. Following electrical simulations, we obtain four crucial electrical values: the self-capacitance of the two qubits (C_{01} and C_{02}), the self-capacitance of the coupler (C_{0c}), and the mutual capacitance between each qubit (C_{12}) and between each qubit and the coupler (C_{1c} and C_{2c}). For this structural context, we make the assumption of identical and symmetric properties for the two qubits, resulting in $C_{1c} = C_{2c}$ and $C_{01} = C_{02}$.

The energy associated with each configuration of the quantum chip is computed using the following equations (Sete et al., 2021; Li and Jin, 2023):

$$
\begin{aligned}
E_{Cq_1} &= \frac{e^2}{2\left(C_{01} + C_{1c} + C_{12}\right)}, \\
E_{Cc} &= \frac{e^2}{2\left(C_{0c} + C_{1c} + C_{2c}\right)}, \\
E_{Cq_2} &= \frac{e^2}{2\left(C_{02} + C_{2c} + C_{12}\right)}, \\
E_{q_1c} &= e^2\left(-C_{1c}\right)^{-1}, \\
E_{q_2c} &= e^2\left(-C_{2c}\right)^{-1}, \\
E_{q_1q_2} &= e^2\left(-C_{12}\right)^{-1},
\end{aligned}
$$

where e represents the elementary charge in coulombs (C). For convenience in calculation, we adopt $e = 1.602$ (rather than 1.602×10^{-19}). The corresponding coupling strength g is given by (Sete et al., 2021; Li and Jin, 2023):

$$
g = \frac{2\omega_q}{B}\left(A - \frac{\omega_c^2}{\omega_c^2 - \omega_q^2}\right),
$$

where

$$
A = \frac{2E_{12}E_{Cc}}{E_{1c}E_{2c}}, \quad B = \frac{16E_{Cq}E_{Cc}}{E_{1c}E_{2c}},
$$

and ω_c and ω_q represent the frequencies of the coupler and qubits, respectively. These values are contingent on experimental results and may vary across different studies. According to (Li and Jin, 2023), ω_c ranges from 9 GHz to over 16 GHz, while for the Zuchongzhi 2.1 quantum chip, ω_q is 5.099 GHz, and for Sycamore, it is 6.924 GHz. In our synthetic chip, we set $\omega_c = 20$ GHz and $\omega_q = 6.9$ GHz.

These equations encapsulate the energy associated with each component of the quantum chip configuration. The computed energy values furnish valuable insights into the performance and efficiency of the quantum chip.

E.7 Regrets due to Suboptimality and Constraint Violation

We plot the regrets at the estimator $\tilde{\mathbf{x}}_t$ resulting from suboptimality, i.e., $r_f(\tilde{\mathbf{x}}_t)$, and from constraint violation, i.e., $\sum_{c\in\mathcal{C}} r_c(\tilde{\mathbf{x}}_t)$, against the number of input queries in Fig. 4. The use of a log-scale causes the line plot to extend beyond the plotting area when the regret is exactly 0 (a log value of $-\infty$).

Fig. 4a shows that the estimator $\tilde{x}_*$ of the optimal solution is typically a feasible input when there are no active constraints at the optimal solution (such as in the problem [S-A0]). However, in cases where active constraints exist at the optimal solution, accurately estimating the boundary of the feasible region at the optimal solution becomes challenging. Consequently, the estimator often becomes an

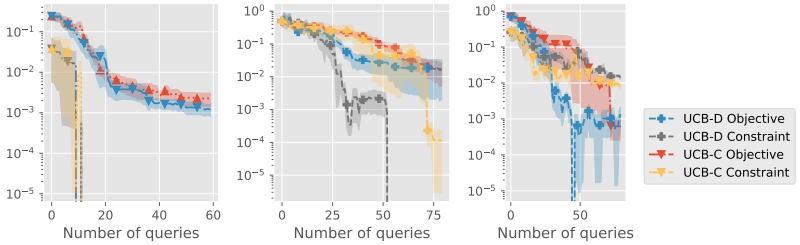

(a) [S-A0] (b) [S-A1] (c) [S-A2]

Figure 4: Plot of the regret $r_f(\tilde{\mathbf{x}}_t)$ at the estimator due to suboptimality (labeled as *UCB-C Objective* and *UCB-D Objective*) and the regret $\sum_{c\in\mathcal{C}} r_c(\tilde{\mathbf{x}}_t)$ due to constraint violation (labeled as *UCB-C Constraint* and *UCB-D Constraint*) against the number of input queries.

infeasible input (Figs. 4b,c). Nevertheless, it is worth noting that the sum of regrets serves as an upper bound for the regrets associated with the objective function and constraint functions. Hence, even if the estimator is infeasible, the no-regret result in Lemma 3.5 shows that we can achieve arbitrarily small constraint violation at the expense of more BO iterations.

Appendix A delves into the rationale behind our selecting an estimator within the optimistic feasible region, even if it may be infeasible, particularly in the context of equality constraints. More generally, if the feasible region has an empty interior (e.g., due to equality constraints), pinpointing a feasible input accurately becomes impossible, regardless of the number of observations gathered. For example, if the feasible region is only a line in the space, then estimating the line without any error (to identify feasible inputs) using only noisy observations is not possible. Our strategy of utilizing the optimistic feasible region circumvents this challenge, albeit with a minor constraint violation if active constraints are present at the optimal solution. It is important to recall that in situations where there are no active constraints at the optimal solution, our estimator is often a feasible input, as illustrated in Fig. 4a. Additionally, as the number of BO iterations increases, the amount of constraint violation approaches zero, allowing for arbitrarily small violations at the expense of a more extended BO procedure.