

A PROOFS

A.1 PROOF OF LEMMA 1

Proof: Similar to lemma 5.1 of Srinivas et al. (2009), with probability at least $1 - 1/2\delta$, $\forall \mathbf{x} \in \tilde{D}, \forall t \geq 1, \forall g \in \{f\} \cup \{\mathcal{C}_k\}_{k \in \mathbf{K}}$,

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

Note that we also take the union bound on $g \in \{f\} \cup \{\mathcal{C}_k\}_{k \in \mathbf{K}}$.

First, by definition $S_{\mathcal{C},t} \triangleq \bigcap_k^{\mathbf{K}} S_{\mathcal{C}_k,t}$, we have $\forall t \leq T, \mathbf{x} \in S_{\mathcal{C},t}, \forall k \in \mathbf{K}$

$$\mathbb{P} \left[\mathcal{C}_k(\mathbf{x}) \geq \text{LCB}_{\mathcal{C}_k,t}(\mathbf{x}) = \mu_{\mathcal{C}_k,t-1}(\mathbf{x}) - \beta_t^{1/2} \sigma_{\mathcal{C}_k,t-1}(\mathbf{x}) > 0 \right] \geq 1 - 1/2\delta$$

meaning with probability at $1 - \delta$, \mathbf{x} lies in the feasible region. At the same time, we have, $\forall t \leq T$

$$\mathbb{P} [\text{UCB}_{f,t}(\mathbf{x}^*) \geq f(\mathbf{x}^*) \geq f(\mathbf{x}) \geq \text{LCB}_{f,t}(\mathbf{x}) \mid \mathcal{C}_k(\mathbf{x}) > 0, \forall k \in \mathbf{K}] \geq 1 - 1/2\delta$$

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

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

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

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. $\mathcal{C}_k(\mathbf{x}^*) > \epsilon_{\mathcal{C}}$ we have $\forall t \leq T, \forall k \in \mathbf{K}$

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

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

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

□

A.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 A.1 *Under the conditions assumed in Theorem 1, let $\alpha_t = \max_{g \in \mathcal{G}} \alpha_{g,t}(\mathbf{x}_{g,t})$ as in Algorithm 1, with $\beta_t = 2 \log(2|\tilde{D}_{\hat{\mathbf{X}}_t}| \pi_t / \delta)$ that is non-increasing, after at most $T \geq \frac{\beta_T \gamma_T C_1}{\epsilon^2}$ iterations, $\alpha_T \leq \epsilon$ Here $C_1 = 8 / \log(1 + \sigma^{-2})$.*

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

$\forall T \geq t \geq 1, \forall g \in \{\mathcal{C}_k\}_{k \in \mathbf{K}}$, when $\tilde{D}_{\hat{\mathbf{X}}_t} \cap U_{g,t} \neq \emptyset$,

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{X}}_t} \cap U_{g,t}} \text{UCB}_{g,t}(\mathbf{x}) - \text{LCB}_{g,t}(\mathbf{x}) = 2\beta_{g,t}^{1/2} \sigma_{g,t-1}(\mathbf{x}_{g,t}) \leq \alpha_t \quad (8)$$

$\forall T \geq t \geq 1, \forall g \in \{\mathcal{C}_k\}_{k \in \mathbf{K}}$, when $\tilde{D}_{\hat{\mathbf{X}}_t} \cap U_{\mathcal{C}_k,t} = \emptyset$, let

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{X}}_t} \cap U_{g,t}} \text{UCB}_{g,t}(\mathbf{x}) - \text{LCB}_{g,t}(\mathbf{x}) = 2\beta_{g,t}^{1/2} \sigma_{g,t-1}(\mathbf{x}_{g,t}) = 0 \leq \alpha_t \quad (9)$$

$\forall T \geq t \geq 1, g = f$

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

$$= 2\beta_{g,t}^{1/2} \sigma_{g,t-1}(\mathbf{x}_{g,t}) \quad (11)$$

$$\leq \alpha_t \quad (12)$$

By lemma 5.4 of Srinivas et al. (2009), with $\beta_t = 2 \log(2(K+1)|\tilde{D}_{\hat{\mathbf{x}}_t}| \pi_t / \delta)$, $\forall g \in \{f\} \cup \mathcal{C}_{k \in \mathbf{K}}$ and $\forall x_t \in \tilde{D}_{\hat{\mathbf{x}}_t}$, we have $\sum_{t=1}^T (2\beta_t^{1/2} \sigma_{g,t-1}(\mathbf{x}_t))^2 \leq C_1 \beta_T \gamma_{g,T}$. By definition of α_t , we have the following

$$\begin{aligned} \sum_{t=1}^T \alpha_t^2 &\leq \sum_{t=1}^T \sum_{g \in \mathcal{G}} (\alpha_{g,t}(\mathbf{x}_{g,t}))^2 \\ &= \sum_{t=1}^T \sum_{g \in \mathcal{G}} (2\beta_{g,t}^{1/2} \sigma_{g,t-1}(\mathbf{x}_{g,t}))^2 \\ &\leq \sum_{g \in \mathcal{G}} C_1 \beta_T \gamma_{g,T} \\ &= C_1 \beta_T \widehat{\gamma}_T \end{aligned}$$

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

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

By the monotonicity assumed in *Assumption 3*, the definition of $U_{g,t}$, $\forall g \in \{\mathcal{C}_k\}_{k \in \mathbf{K}}$, and the definition of $\hat{\mathbf{x}}_t$, for $\forall 1 \leq t_1 < t_2 \leq T$, $\forall g \in \{\mathcal{C}_k\}_{k \in \mathbf{K}}$, we have that $U_{g,t_2} \subseteq U_{g,t_1}$ and $\hat{\mathbf{x}}_{t_2} \subseteq \hat{\mathbf{x}}_{t_1}$. Meaning the search space is shrinking for all constraints and the objective. Together with the monotonicity of UCB and LCB, for $\forall 1 \leq t_1 < t_2 \leq T$, we have $\alpha_{t_2} \leq \alpha_{t_1}$, and therefore

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

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

□

With Lemma A.1, we could first prove that after adequately T rounds of evaluations such that $\epsilon \leq \min_{k \in \mathbf{K}} \epsilon_k$ is sufficiently small, with certain probability, $\mathbf{x}^* \in S_{\mathcal{C},T}$. 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 the high confidence interval of f^* , is bounded by ϵ .

Proof: We first prove that after at most $T \geq \frac{\beta_T \widehat{\gamma}_T C_1}{\epsilon^2}$ iterations, $\mathbb{P}[\mathbf{x}^* \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap S_{\mathcal{C},T}] \geq 1 - 1/2\delta$. Given equation 8 and 9 and Lemma A.1, we have $\forall g \in \mathcal{C}_{k \in \mathbf{K}}, t \geq T$

$$\max_{\mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t}} \text{UCB}_{g,t}(\mathbf{x}) - \text{LCB}_{g,t}(\mathbf{x}) \leq \epsilon \leq \min_{k \in \mathbf{K}} \epsilon_k$$

According to the definition of $U_{g,t}$, $\forall \mathbf{x} \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap U_{g,t}, \forall g \in \mathcal{C}_{k \in \mathbf{K}}$

$$\text{UCB}_{g,t}(\mathbf{x}) \leq \min_{k \in \mathbf{K}} \epsilon_k + \text{LCB}_{g,t}(\mathbf{x}) \leq \min_{k \in \mathbf{K}} \epsilon_k$$

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

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

Hence $\forall t \geq T$

$$\mathbb{P}[\mathbf{x}^* \in \tilde{D}_{\hat{\mathbf{x}}_t} \cap S_{\mathcal{C},t} = \tilde{D}_{\hat{\mathbf{x}}_t} \cap \hat{\mathbf{x}}_{\mathcal{C},t} \setminus \cup_{k \in \mathbf{K}} U_{\mathcal{C}_k,t}] \geq 1 - 1/2\delta$$

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 high-confidence interval of f^* . Given that $\text{LCB}_{f,t,\max} \neq -\infty$, we have

$$\begin{aligned} \max_{\mathbf{x} \in \tilde{D}_{\mathbf{x}_t}} \text{UCB}_{f,T}(\mathbf{x}) - \max_{\mathbf{x} \in \tilde{D}_{\mathbf{x}_t}} \text{LCB}_{f,T}(\mathbf{x}) &\leq \max_{\mathbf{x} \in \tilde{D}_{\mathbf{x}_t}} \text{UCB}_T(\mathbf{x}) - \text{LCB}_{f,T,\max} \\ &\leq 2\beta_{f,T}^{1/2} \sigma_{f,T-1}(\mathbf{x}_{f,T}) \\ &\leq \alpha_T \\ &\leq \epsilon \end{aligned}$$

Combining it with the fact that

$$\mathbb{P} \left[\max_{\mathbf{x} \in \tilde{D}_{\mathbf{x}_t}} \text{LCB}_{f,T}(\mathbf{x}) \leq \max_{\mathbf{x} \in \tilde{D}_{\mathbf{x}_t}} f(\mathbf{x}) = f^* \leq \text{UCB}_{f,T}(\mathbf{x}^*) \leq \max_{\mathbf{x} \in \tilde{D}_{\mathbf{x}_t}} \text{UCB}_{f,T}(\mathbf{x}) \right] \geq 1 - 1/2\delta$$

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

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

□

B DECOUPLED SETTING

In the main paper, we assume both objective f and the constraints $\{\mathcal{C}_k\}_{k \in \mathbf{K}}$ 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 $g_t \in \{f\} \cup \{\mathcal{C}_k\}_{k \in \mathbf{K}}$ for evaluation each time. This typically requires different acquisition from coupled setting (Gelbart et al., 2014). However, we will that our acquisition function and COBALT require minimum adaptation to the decoupled setting while bearing a similar performance guarantee.

B.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 COBALT to the decoupled setting by querying $\mathbf{x}_{g,t}$ on this unknown function $g_t \in \mathcal{G} \subseteq \{f\} \cup \{\mathcal{C}_k\}_{k \in \mathbf{K}}$ at iteration t . The modified algorithm is shown below.

B.2 THOERETICAL GUARANTEE AND PROOF

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

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

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

$$T = \sum_{g \in \{f\} \cup \{\mathcal{C}_k\}_{k \in \mathbf{K}}} T_g$$

Then we have the following guarantee for the performane of COBALT-Decoupled.

Theorem 2 *The width of the resulting confidence interval of the global optimum $f^* = f(\mathbf{x}^*)$ is upper bounded. That is, under the same assumptions in Theorem 1, with $\beta_t = 2\log(2(K+1)|\tilde{D}_{\mathbf{x}_t}| \pi_t / \delta)$ that is constant, and acquisition function in Algorithm 2, $\exists \epsilon \leq \min_{k \in \mathbf{K}} \epsilon_k$, after at most $T \geq \frac{\beta_T \tilde{\gamma}_T C_1}{\epsilon^2}$ iterations, we have $\mathbb{P}[|CI_{f^*,t}| \leq \epsilon, f^* \in CI_{f^*,t} \mid t \geq T] \geq 1 - \delta$ Here $C_1 = 8/\log(1 + \sigma^{-2})$.*

Algorithm 2 CONstrained BO with Adaptive active Learning of *decoupled* unknown constraints (COBALT-Decoupled)

```

1: Input: Search space  $\mathbf{X}$ , initial observation  $\mathbf{D}_0$ , horizon  $T$ ;
2: for  $t = 1$  to  $T$  do
3:   Update the posteriors of  $\mathcal{GP}_{f,t}$  and  $\mathcal{GP}_{C_k,t}$  according to equation 1 and 2
4:   Identify ROIs  $\tilde{\mathbf{X}}_t$ , and undecided sets  $U_{C_k,t}$ 
5:   for  $k \in \mathbf{K}$  do
6:     if  $U_{C_k,t} \neq \emptyset$  then
7:       Candidate for active learning of each constraints:
        $\mathbf{x}_{C_k,t} \leftarrow \arg \max_{\mathbf{x} \in U_{C_k,t}} \alpha_{C_k,t}(\mathbf{x})$  as in equation 6
8:        $\mathcal{G} \leftarrow \mathcal{G} \cup C_{k,t}$ 
9:       Candidate for optimizing the objective:
        $\mathbf{x}_{f,t} \leftarrow \arg \max_{\mathbf{x} \in \tilde{\mathbf{X}}_{f,t}} \alpha_{f,t}(\mathbf{x})$  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:  $\mathbf{x}_t \leftarrow \mathbf{x}_{g,t}$ 
13:      Update the observation set with the candidate and corresponding new observations on  $g_t$ 
        $\mathbf{D}_t \leftarrow \mathbf{D}_{t-1} \cup \{(\mathbf{x}_t, y_{g,t})\}$ 

```

The proof is similar to Appendix A, as the major difference is replacing the upper bound in Lemma A.1 to

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

Proof: We omit the shared part of the proof. Here is the critical difference.

$$\begin{aligned}
\sum_{t=1}^T \alpha_t^2 &= \sum_{t=1}^T \alpha_{g,t}^2(\mathbf{x}_{g,t}) \\
&= \sum_{t=1}^T (2\beta_{g,t}^{1/2} \sigma_{g,t-1}(\mathbf{x}_{g,t}))^2 \\
&\leq \sum_{g \in \mathcal{G}} C_1 \beta_T \gamma_{g,T_g} \\
&= C_1 \beta_T \tilde{\gamma}_T
\end{aligned}$$

□

C REWARD FUNCTION

C.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 $r(\mathbf{x}) = f(\mathbf{x}) \prod_k \mathbb{1}_{C_k(\mathbf{x}) > h_k}$ 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

$$\arg \max_{\mathbf{x} \in \mathbf{X}} r(\mathbf{x}) = \arg \max_{\mathbf{x} \in \mathbf{X}} f(\mathbf{x}) \prod_k \mathbb{P}[Y_{C_k}(\mathbf{x}) > h_k]$$

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

The problem with this product reward, on 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 $\mathbb{P}[Y_{C_k}(\mathbf{x}) > h_k]$. 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, is that 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 if 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 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.

C.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 $Y_{C_k}(\mathbf{x}) = C_k(\mathbf{x})$, we could simply use the indicator function μ_k for each constraint to turn the feasibility function into an indicator function.

$$r(\mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } \mathbb{I}(C_k(\mathbf{x}) > h_k) \quad \forall k \in \mathbf{K} \\ -inf & \text{o.w} \end{cases} \quad (14)$$

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

$$\mathbb{P}[Y_{C_k}(\mathbf{x}) > h_k] \geq \mu_k$$

to explicitly deal with the aleatoric uncertainty. With the percentage point function (PPF), we could transform the probabilistic constraints into a deterministic constraint $\mathbb{I}(C_k(\mathbf{x}) > \hat{h}_k)$ with $\hat{h}_k = PPF(h_k, \sigma, \mu_k)$, meaning \hat{h}_k is the μ_k percent point of a Gaussian distribution with h_k 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 $h_k = 0$.

Throughout the rest of the paper, we want to efficiently locate the global maximizer

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathbf{X}, \forall k \in \mathbf{K}, C_k(\mathbf{x}) > 0} f(\mathbf{x})$$

Equivalently, we seek to achieve the performance guarantee in terms of simple regret at certain time t ,

$$\mathbf{R}_t := r(\mathbf{x}^*) - \max_{\mathbf{x} \in \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t\}} r(\mathbf{x})$$

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

$$P(\max_{\mathbf{x} \in \hat{\mathbf{X}}} \mathbf{R}_T(\mathbf{x}) \geq \epsilon) \leq 1 - \delta$$

D DATASET

Here we offer a more detailed discussion over the construction of the six CBO tasks studied in section 6.

D.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). 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 excludes 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 COBALT. Figure 3 shows the distribution of the objective function and feasible regions.

Ackley-5D-2C The Ackley function is also a popular benchmark for optimization algorithms. Compared with the Rastrigin function, it is highly multimodal similarly, 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: $f(\mathbf{x}) = 20 \exp(-0.2 \sqrt{1/d \sum_{i=1}^d x_i^2}) + \exp(1/d \sum_{i=1}^d \cos(2\pi x_i)) + 20 + \exp(1)$, $d = 5$. 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 $(\|\mathbf{x} - \mathbf{1}\|_2 - 5.5)^2 - 1 > 0$ constructs two feasible regions with one in the center and the other close to the boundary of the search space. The second constraint $-\|\mathbf{x}\|_\infty^2 + 9$ allows one hypercube feasible region in the center.

D.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 at optimizing 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 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.

E DISCUSSIONS

Here we offer additional discussion over the concerns on COBALt.

E.1 EMPTY ROI(S)

It is possible that $\hat{\mathbf{X}}_t$ could be empty at certain t when any intersection results in the empty set. However, according to the assumptions in section 5 and Lemma 1, the properly chosen $\beta_{f,t}$ and $\beta_{c,t}$ that does not result in over-aggressive filtering, the ROI is soundly defined. The algorithm is also robust to empty $U_{C_k,t}$ due to the domain where the acquisition functions defined in equation 6 and equation 5 are maximized.

E.2 COMPARABILITY

Despite both the acquisition function for optimization of the objective and active learning are confidence interval-based, it is possible they are not comparable. In practice, the objective and constraints could be of different scales. With prior knowledge of the scaling difference, one can choose to standardize the values, or equivalently, calibrate the acquisition function accordingly.

E.3 LIMITATIONS

The limitation of COBALt including (1) the insufficiency of identifying the ROIs due to the pointwise comparison in current implementation; (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). We expect the following work could further improve the algorithms efficiency and effectiveness accordingly.