Constrained Bayesian Optimization under Partial Observations: Balanced Improvements and Provable Convergence

Shengbo Wang[1], **Ke Li**[2]

[1]School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu, China
[2]Department of Computer Science, University of Exeter, EX4 4RN, Exeter, UK
shnbo.wang@foxmail.com, k.li@exeter.ac.uk

Abstract

The partially observable constrained optimization problems (POCOPs) impede data-driven optimization techniques since an infeasible solution of POCOPs can provide little information about the objective as well as the constraints. We endeavor to design an efficient and provable method for expensive POCOPs under the framework of constrained Bayesian optimization. Our method consists of two key components. Firstly, we present an improved design of the acquisition functions that introduce balanced exploration during optimization. We rigorously study the convergence properties of this design to demonstrate its effectiveness. Secondly, we propose Gaussian processes embedding different likelihoods as the surrogate model for partially observable constraints. This model leads to a more accurate representation of the feasible regions compared to traditional classification-based models. Our proposed method is empirically studied on both synthetic and real-world problems. The results demonstrate the competitiveness of our method for solving POCOPs.

Introduction

The black-box constrained optimization problem considered in this paper is formulated as:

$$\text{minimize } f(\mathbf{x}) \quad \text{subject to} \quad \vec{g}(\mathbf{x}) \leq 0, \tag{1}$$

where $\mathbf{x} = (x_1, \cdots, x_n)^\top \in \Omega$ denotes the decision variable, $\Omega = [x_i^{\mathrm{L}}, x_i^{\mathrm{U}}]_{i=1}^n \subset \mathbb{R}^n$ denotes the search space, x_i^{L} and x_i^{U} are the lower and upper bounds of x_i respectively. The objective function $f(\mathbf{x})$ and m constraint functions $\vec{g}(\mathbf{x}) = (g_1(\mathbf{x}), \cdots, g_m(\mathbf{x}))^\top$ are: i) *analytically unknown*, i.e., we do not have access to f and $\vec{g}$ directly, but to $\mathbf{x}$ to be determined and their observations $f(\mathbf{x})$ and $\vec{g}(\mathbf{x})$ instead; ii) *computationally expensive*; and iii) *partially observable*, i.e., the values of f and $\vec{g}$ are not observable/measurable when $\mathbf{x}$ is infeasible. We denote the unknown feasible space by $\chi = \{\mathbf{x} \in \Omega | \vec{g}(\mathbf{x}) \leq 0\}$. Partially observable constrained optimization problems (POCOPs) are not uncommon in real-life applications. For example, a robot control task will be suspended when a collision or excessive instantaneous power consumption is detected, where the feedback is merely a failure message rather than any reward (Marco et al. 2021). An AutoML task will be terminated without outputting the performance of a hyperparameter configuration but an error log when there is a memory overflow or computation timeout (Perrone et al. 2019).

Bayesian optimization (BO) is recognized as an effective query-efficient framework for black-box optimization (Garnett 2023). Although there have been dedicated efforts on constraint handling in the context of BO, a.k.a. constrained Bayesian optimization (CBO), most of them are however expected to work with complete observations. Considering missing observations, existing CBO methods may become inefficient due to the following two issues.

- First, existing CBO methods risk overly exploiting known feasible regions in POCOPs. In particular, the update of an acquisition function such as expected improvement (EI) (Jones, Schonlau, and Welch 1998) stagnates when objective values are unobservable outside χ. This stagnation may cause cluttered observations in local feasible regions, resulting in an overconfidence effect in both surrogate modeling and candidate acquisition. Consequently, the efficiency of BO diminishes.
- Second, POCOPs generate mixed data from both feasible and infeasible solutions, which complicate sufficient exploitation using conventional surrogate models. When a probabilistic classifier like Gaussian process classifier (GPC) is employed to distinguish between feasible and infeasible solutions, the available observations lose their utility in refining the GPC model. Brockman et al. leveraged value observations by using regression models and artificially injecting values into infeasible solutions. However, this approach potentially introduces erroneous priors, thereby compromising optimization efficacy.

Bearing these considerations in mind, this paper proposes a novel CBO framework for POCOPs. Our main contributions are outlined as follows.

- To address the risk of overly exploiting evaluated local feasible regions, we propose a novel acquisition function framework. It enhances EI with a balanced constraint handling technique, encapsulated in a general exploration function to enable global search during optimization. We theoretically analyze the convergence properties of this design, and further develop an instance of the exploration function, effectively enhancing the efficiency of our method for solving POCOPs.

- To fully leverage the mixed-observation characteristic of POCOPs, we propose heterogeneous-likelihood Gaussian processes (HLGPs), providing a promising representation of unknown constraints compared to classifier-based models. Further, we employ expectation propagation to manage the non-Gaussian inference, yielding an efficient Gaussian approximation of HLGPs.
- To demonstrate the efficacy of our proposed method, we perform a series of experiments on diverse benchmark problems. These include synthetic problems and real-world applications in reinforcement learning-based control design and machine learning hyperparameter optimization. Experimental results show the competitiveness and better efficiency of our method compared to selected state-of-the-art CBO methods.

Related Works

Constrained Bayesian Optimization

Classic CBO methods usually reshape an unconstrained acquisition function by incorporating feasibility considerations. A prominent approach, the EI with constraints (EIC), first introduced by Schonlau, Welch, and Jones, has been extensively employed to locate feasible solutions with high probability (Gardner et al. 2014; Gelbart, Snoek, and Adams 2014). To further enhance EIC's capability, Letham et al. leveraged a quasi-Monte Carlo approximation regarding observation noises. Furthermore, various strategies such as integration (Gelbart, Snoek, and Adams 2014), rollout (Lam and Willcox 2017), and a two-step lookahead algorithm (Zhang, Zhang, and Frazier 2021), have been proposed to achieve better optimization efficiency. These approaches increase the exploration of unknown regions, albeit at the cost of computational complexity, hindering their scalability for high-dimensional problems.

From the perspective of uncertainty reduction, predictive entropy search (PESC), which selects feasible candidate solutions directly from the search space, offers attractive heuristics to constrained optimization problems (Hernández-Lobato et al. 2015). However, the intractable nature of quadrature calculations during sampling in PESC has been a challenge. To address this, Takeno et al. proposed the min-value entropy search, enabling the sampling process to operate more effectively within the objective space. This concept was subsequently extended to accommodate binary observations (Perrone et al. 2019) and multi-objective scenarios (Belakaria, Deshwal, and Doppa 2019). Besides, the fusion of EI and entropy search showed promise for enhancing exploration (Lindberg and Lee 2015).

To harness the structure inherent in equation (1), researchers such as Gramacy et al. and Picheny et al. proposed the utilization of a Lagrangian method with slack variables, providing the capability to deal with equality constraints. Regarding problems marked by unknown constraints, Ariafar et al. integrated BO with the alternating direction method of multipliers, thereby enabling the exploration of solutions even in the absence of feasible ones. Meanwhile, Eriksson and Poloczek proposed the use of Thompson sampling combined with a trust region approach towards enhancement of the scalability, meanwhile incorporating a thoughtful design aimed at maintaining computational efficiency.

Surrogate Models for Constraints

The aforementioned CBO methods frequently employ Gaussian Process regression (GPR) and GPC to construct surrogate models for unknown constraints. Notably, classifiers such as GPC and support vector machine (SVM) have been found effective in sequential updates when the real value of an infringed constraint remains unobservable (Lindberg and Lee 2015; Perrone et al. 2019; Ariafar et al. 2019; Bachoc, Helbert, and Picheny 2020; Candelieri 2021). Yet, in partially observable scenarios, these classifiers exhibit limitations in utilizing available real-value observations, resulting in a dip in modeling performance. In response to this challenge, Marco et al. enhanced the construction of GPR models by introducing a switched likelihood combined with mixed data observations. As an alternative solution, Pourmohamad and Lee and Zhang, Dai, and Low proposed multivariate GPs (MVGP) with joint distributions of hybrid input to handle mixed observations.

Exploration for Unknown Feasible Regions

In light of the black-box nature of the problem at hand, discerning the feasibility of a solution becomes a significant concern. In this context, Parr et al. interpreted the delicate balance between enhancing the probability of feasibility (POF) and optimizing the objective as a multi-objective optimization issue. Focusing on optimization, Picheny developed a step-wise uncertainty reduction method, which capitalizes on the volume of feasible regions beneath the most promising solution observed up to that point, despite the method's considerable computational complexity. Furthermore, the EIC was adapted in (Lindberg and Lee 2015; Wang and Ierapetritou 2018) to deepen the understanding of global feasible regions. Regarding exploration, level-set or contour estimation techniques have been adapted to locate unknown feasible regions (Ranjan, Bingham, and Michailidis 2008; Bect et al. 2012; Bachoc, Cesari, and Gerchinovitz 2021). Yet, these methods can be excessively aggressive, thereby hindering optimization progress.

Preliminaries of CBO

Conventional BO starts from uniformly sampling a set of solutions according to a space-filling experimental design method. Thereafter, it sequentially updates its next sample until the given computational budget is exhausted. BO consists two main components: i) a surrogate model for approximating the true expensive objective function; and ii) an infill criterion (based on the optimization of an acquisition function) for deciding the next point of merit.

Surrogate Model

Given a set of training data $\mathcal{D} = \left\{\left(\mathbf{x}^i, f(\mathbf{x}^i)\right)\right\}_{i=1}^{N}$, we apply the GPR model to learn a Guassian process $\tilde{f}(\mathbf{x})$ with a prior mean function $m(\mathbf{x})$ and a noise-free likelihood (Rasmussen and Williams 2005). For a candidate solution $\hat{\mathbf{x}}$, the mean and variance of the target $f(\tilde{\mathbf{x}})$ can be predicted as:

$$\begin{aligned}\mu_f(\tilde{\mathbf{x}}) &= m(\tilde{\mathbf{x}}) + \mathbf{k}^{*\top}K^{-1}\mathbf{f},\\ \sigma_f^2(\tilde{\mathbf{x}}) &= k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}) - \mathbf{k}^{*\top}K^{-1}\mathbf{k}^*,\end{aligned} \tag{2}$$

where $\mathbf{k}^*$ is the covariance matrix between X and $\tilde{\mathbf{x}}$, K is the covariance matrix of X, $X = \left(\mathbf{x}^1, \ldots, \mathbf{x}^N\right)^\top$, and $\mathbf{f} = \left(f(\mathbf{x}^1) - m(\mathbf{x}^1), \ldots, f(\mathbf{x}^N) - m(\mathbf{x}^N)\right)^\top$. In this paper, we use the Matérn $5/2$ as the kernel function combined with the constant mean function for all GP models by default. As for the i-th constraint in (1), it will be modeled by an independent GP model $\tilde{g}_i(\tilde{\mathbf{x}})$ whose predictive mean and variance are denoted by $\mu_g^i(\tilde{\mathbf{x}})$ and $\sigma_g^i(\tilde{\mathbf{x}})$ respectively.

Infill Criterion

Instead of directly working on $\tilde{f}(\mathbf{x})$, the actual search process of BO is driven by an acquisition function that naturally strikes a balance between exploitation of the predicted optimum and exploration regarding uncertainty. This paper applies the widely used EI to serve this purpose:

$$\mathrm{EI}(\tilde{\mathbf{x}}|\mathcal{D}) = \sigma_f(\tilde{\mathbf{x}})\big(z\Phi_f(z) + \phi_f(z)\big), \tag{3}$$

where $z = \frac{f^*_{\mathcal{D}} - \mu_f(\tilde{\mathbf{x}})}{\sigma_f(\tilde{\mathbf{x}})}$, $f^*_{\mathcal{D}} = \min\limits_{(\mathbf{x}, f(\mathbf{x}))\in\mathcal{D}} f(\mathbf{x})$, Φ_f and ϕ_f denote the cumulative distribution function and probability density function according to $\tilde{f}$, respectively.

To tackle unknown constraints, EIC was proposed as a product of the EI with POF (Gardner et al. 2014):

$$\mathrm{EIC}(\tilde{\mathbf{x}}|\mathcal{D}) = \mathrm{EI}(\tilde{\mathbf{x}}|\mathcal{D}) \cdot \mathrm{POF}(\tilde{\mathbf{x}}), \tag{4}$$

with

$$\mathrm{POF}(\tilde{\mathbf{x}}) = \mathbb{P}\left[\vec{g}(\tilde{\mathbf{x}}) \leq \lambda\right] = \prod_{i=1}^{m} \Phi_g^i(\lambda), \tag{5}$$

where Φ_g^i denotes the cumulative distribution function of the i-th constraint based on a GPR model $\tilde{g}_i(\tilde{\mathbf{x}}) \sim \mathcal{N}(\mu_g^i(\tilde{\mathbf{x}}), {\sigma_g^i}^2(\tilde{\mathbf{x}}))$, λ is the threshold of a feasible level and is set to a constant 0 in this paper.

Proposed Method

This section delineates the implementation of our method, the CBO with balance (dubbed `CBOB`), for POCOPs. As shown in Algorithm 1, `CBOB` adheres to the conventional CBO procedure while introducing two unique algorithmic components (highlighted by $\triangleright$). The first is a framework for designing an acquisition function, facilitating balanced exploration by effectively harnessing the surrogate models of constraints. The second is a bespoke GP model, specifically formulated to model constraints using partial observations.

Algorithm 1: Pseudo code of `CBOB`

Input: Initial dataset $\mathcal{D} = \left\{\left(\mathbf{x}^i, f(\mathbf{x}^i), \vec{g}(\mathbf{x}^i)\right)\right\}_{i=1}^{N_0}$, budget N, and priors of GPs
Output: The optimal feasible objective $f^*_{\mathcal{D}}$

for $k \leftarrow 1$ **to** N **do**
  Build a GPR model for the black-box objective;
  for $i \leftarrow 1$ **to** m **do**
    $\triangleright$ Update $\mathbf{g}_i$ in $\mathcal{D}$ with modified observations $\tilde{\mu}_g^i$ and $\tilde{\Sigma}_g^i$ using equation (13);
    $\triangleright$ Build an HLGP model based on $\mathbf{g}_i$;
  $\triangleright$ $\mathbf{x}^k \leftarrow \arg\max_{\tilde{\mathbf{x}}\in\Omega} \mathrm{EICB}(\tilde{\mathbf{x}}|\mathcal{D})$;
  if $\mathbf{x}^k$ *is feasible* **then**
    $\mathcal{D} \leftarrow \mathcal{D} \cup \left\{\left(\mathbf{x}^k, f(\mathbf{x}^k), \vec{g}(\mathbf{x}^k)\right)\right\}$;
  else
    $\triangleright$ $g_i^k \leftarrow +1$ for the i-th violated constraints;
    $\vec{g}(\mathbf{x}^k) = (g_1^k, \ldots g_m^k)$;
    $\triangleright$ $\mathcal{D} \leftarrow \mathcal{D} \cup \left\{\left(\mathbf{x}^k, \mathrm{Null}, \vec{g}(\mathbf{x}^k)\right)\right\}$;

A Framework for Acquisition Function Design

Under the condition of partial observations, the EI function only updates upon evaluation of a feasible solution, while the POF predominantly targets known feasible regions. This results in an overemphasis on known feasible regions by the EIC, particularly when tackling POCOPs. Inspired by (Picheny 2014), we posit that prioritizing search towards less explored regions can enhance exploratory capability, thus promoting more global search behaviors in a CBO method. This approach has been empirically substantiated in (Lindberg and Lee 2015; Lam and Willcox 2017; Zhang, Zhang, and Frazier 2021). In this work, we propose a dynamic version of POF (DPOF) that incorporates an additional exploration capability, rather than prioritizing the most uncertain region indiscriminately, as follows:

$$\mathrm{DPOF}(\tilde{\mathbf{x}}) = \prod_{i=1}^{m} \mathrm{Proj}_{[0,1]}\left[(\rho^i(\tilde{\mathbf{x}}) + 1)\Phi_g^i(\lambda)\right], \tag{6}$$

where Proj clips values outside $[0, 1]$ to the boundaries, and ρ^i denotes a general exploration function defined below.

Definition 1 (Exploration function)**.** *A smooth function $\rho^i(\mathbf{x}) : \Omega \to [0, \bar{\rho}]$ is a valid exploration function if it is bounded by $\bar{\rho} > 0$ and $\rho^i(\tilde{\mathbf{x}}) = 0, \forall \sigma_g^i(\mathbf{x}) = 0$.*

With an exploration function ρ^i, DPOF assigns more weights to unknown regions than POF to facilitate a global search. However, to prevent excessive exploration and maintain a high probability of obtaining feasible solutions, we multiply $\Phi_g^i(\lambda)$ with ρ^i in equation (6). Alternatively, this could be viewed as introducing a dynamic constraint threshold $\lambda(\tilde{\mathbf{x}}) = {\Phi_g^i}^{-1}(\mathrm{DPOF}^i(\tilde{\mathbf{x}}))$ that varies across different candidate solutions, where ${\Phi_g^i}^{-1}$ denotes the inverse function of Φ_g^i and $\mathrm{DPOF}^i(\tilde{\mathbf{x}})$ is the i-th factor in equation (6). Note that when $\rho^i \equiv 0$, DPOF simplifies to the traditional POF. Building on this, we introduce a new acquisition function, termed as EI with constraint and balance (EICB), formulated as a product of EI and DPOF:

$$\mathrm{EICB}(\tilde{\mathbf{x}}|\mathcal{D}) = \mathrm{EI}(\tilde{\mathbf{x}}|\mathcal{D}) \cdot \mathrm{DPOF}(\tilde{\mathbf{x}}). \tag{7}$$

Theorem 1. *Assume that the constraint values are fully observable. Assume also that the involved GPs are non-degenerate and satisfy the no-empty-ball property (Vazquez* *and Bect 2010). Let $\mathcal{D}$ be the collected observations with $\left(\mathbf{x}^1, f(\mathbf{x}^1)\right)$ fixed in χ while $\left\{\left(\mathbf{x}^i, f(\mathbf{x}^i)\right)\right\}_{i=2}^{N}$ are sequentially chosen by*

$$\mathbf{x}^i = \arg\max_{\tilde{\mathbf{x}}\in\Omega} \mathrm{EICB}(\tilde{\mathbf{x}}|\mathcal{D}). \tag{8}$$

Then, as $N \to \infty$, almost surely:

1. *the acquisition function $\sup_{\tilde{\mathbf{x}}\in\Omega} \mathrm{EICB}(\tilde{\mathbf{x}}|\mathcal{D}) \to 0$;*
2. *the evaluated best objective $f^*_{\mathcal{D}} \to f^*_{\chi}$;*

where f^*_{χ} represents the global optimum of problem (1).

The proof of Theorem 1 is sketched in Section A of the supplementary document. This theorem suggests that the incorporation of ρ^i as designed in equation (6) does not undermine the asymptotic convergence capability of EI-based acquisition functions, such as EICB. In the following subsection, we propose an instance of the exploration function under Definition 1 that outperforms the EIC in terms of efficiently conducting global optimization for POCOPs.

An instance of the exploration function In the context of EICB framework, exploration during optimization can be facilitated by an apt design of ρ^i. In this paper, we concentrate on identifying promising constraint boundaries, as opposed to aggressively targeting the most uncertain regions, a tactic often employed in level-set estimation and active learning (Ranjan, Bingham, and Michailidis 2008; Bichon et al. 2008; Bect et al. 2012; Bachoc, Cesari, and Gerchinovitz 2021). To this end, we first define a utility function representing the potential of being the boundary (POB) for the i-th constraint at $\tilde{\mathbf{x}} \in \Omega$ as follows:

$$\mathrm{POB}^i(\tilde{\mathbf{x}}) = \begin{cases} 1, & \tilde{g}_i(\tilde{\mathbf{x}}) \in \left[-\varepsilon(\tilde{\mathbf{x}}),\ \varepsilon(\tilde{\mathbf{x}})\right], \\ 0, & \text{otherwise}, \end{cases} \tag{9}$$

where $\varepsilon(\tilde{\mathbf{x}}) = \beta\sigma_g^i(\tilde{\mathbf{x}})$ and $\beta > 0$ represents a confidence level. Taking the expectation of equation (9) over the predicted distribution of $\tilde{g}_i(\tilde{\mathbf{x}})$ and defining $\bar{g}_i(\tilde{\mathbf{x}}) = \mu_g^i(\tilde{\mathbf{x}})/\sigma_g^i(\tilde{\mathbf{x}})$, we obtain a valid exploration function as:

$$\rho^i(\tilde{\mathbf{x}}) = \Phi\left(\beta - \bar{g}_i(\tilde{\mathbf{x}})\right) - \Phi\left(-\beta - \bar{g}_i(\tilde{\mathbf{x}})\right), \tag{10}$$

where Φ is the cumulative distribution function of $\mathcal{N}(0,1)$.

The illustrative example in Figure 1 demonstrates how DPOF, given a surrogate model for a constraint, assigns more weight to the unknown feasible region ($[2.85, 3.45]$ in this case) as β increases. For a previously located feasible region such as $[4.3, 4.7]$, DPOF provides equal weights (approximately 1) to all candidates within this region when $\beta \geq 0.5$. In contrast, POF ($\beta = 0$) assigns differentiated weights based on different $\Phi_g(0)$ values. By refining the boundary, DPOF ensures a more balanced weight distribution within the located feasible region, hence the nomenclature, EICB. We posit that this balanced approach enhances the decision-making capabilities of EI, as compared to the imbalanced weights scenario posed by POF. As a positive consequence, EICB encourages greater exploration towards unknown regions. As ρ^i in equation (10) is bound by 1, DPOF can assign a maximum of $2\Phi_g(0)$ to any given candidate. Empirically, this subtle adjustment leads to enhanced optimization efficiency as evidenced in Figure 1, thanks to the introduction of exploration. Besides, we provide a more aggressive design of ρ^i that bolsters the reduction of uncertainty in global feasible regions in Section B of the supplementary document.

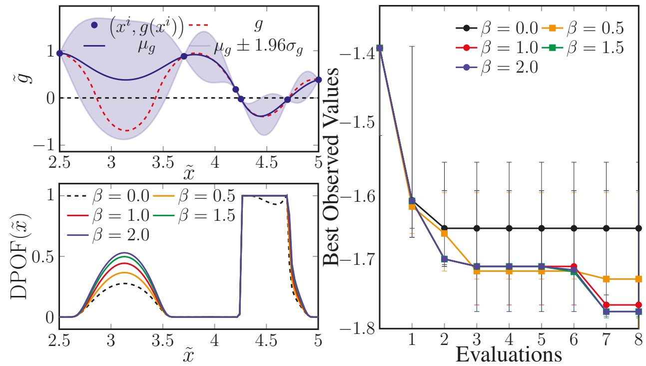

Figure 1: Illustration of EICB with equation (10) by a toy 1-D example (please refer to Section B of the supplementary document for more details). (Left up) The constraint surrogate model with 6 observed data pairs where the red dotted line denotes the true function. (Left down) The DPOF with different β on $[2.5, 5]$, where the dotted line with $\beta = 0$ reduces to POF. (Right) Optimization trajectories on $[0, 10]$ with 5 randomly repeated experiments.

Surrogate Models Under Partial Observations

When dealing with partially observable constraints, observations are composed of two distinct aspects: $i)$ the actual values associated with feasible solutions, and $ii)$ the truncated distribution of possible values for all solutions, such as $\mathbb{P}(\vec{g} > 0) = 1$, as depicted in Figure 2. The simultaneous consideration of these two types of observations can be achieved by attaching individual likelihood distributions to feasible/infeasible solutions, as is done in HLGP.

For the i-th constraint, the posterior of a latent function, $p(\tilde{\mathbf{g}}_i|X, \mathbf{g}_i)$, within an HLGP model is determined via the Bayes rule using the prior distribution $p(\tilde{\mathbf{g}}_i|X) = \mathcal{N}\left(\mathbf{0}, K\right)$ in equation (2), along with the individual likelihood distributions. Specifically, the likelihood can be expressed as:

$$p(g_i^k|\tilde{g}_i^k) = \begin{cases} \mathcal{N}\left(g_i(\mathbf{x}^k), \sigma^2\right), & \text{if } g_i(\mathbf{x}^k) \leq 0, \\ \Phi(\alpha^{-1}g_i(\mathbf{x}^k)), & \text{if } g_i(\mathbf{x}^k) > 0, \end{cases} \tag{11}$$

where $k \in \{1, \ldots, N\}$, $\mathbf{g}_i = (g_i^1, \ldots, g_i^N)^\top$ represents N observations of $g_i(X)$, $\tilde{\mathbf{g}}_i = (\tilde{g}_i^1, \ldots, \tilde{g}_i^N)^\top$ denotes N latent functions, $\sigma \geq 0$ stands for the noise level, and $\alpha > 0$ is a scaling parameter. We set $\sigma = 10^{-6}$ to indicate a noise-free environment and $\alpha = 10^{-6}$ to approximate the truncating step function that implies $\mathbb{P}(g_i(\mathbf{x}^k) > 0) = 1$, $\forall g_i(\mathbf{x}^k) > 0$ (Riihimäki and Vehtari 2010).

HLGP Inference via Expectation Propagation In this paper, we employ expectation propagation (EP) (Minka 2001), a principled and highly efficient approach to handle non-Gaussian likelihoods. This provides Gaussian approximations to both the posterior and predicted distributions of

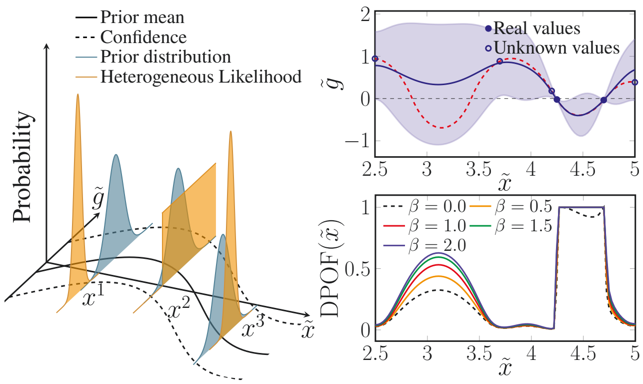

Figure 2: Illustration of HLGPs. (Left) Three observations with heterogeneous likelihood distributions, including a truncated distribution on x^2, and Gaussian distributions on x^1 and x^3. (Right up) An EP-based HLGP model with partial observations. (Right down) The DPOF with different β.

HLGP. First, the posterior is formulated as:

$$p(\tilde{\mathbf{g}}_i|X, \mathbf{g}_i) = \frac{1}{Z} p(\tilde{\mathbf{g}}_i|X) \prod_{k=1}^{N} p(g_i^k|\tilde{g}_i^k), \tag{12}$$

where Z is the normalization factor. For the k-th observation g_i^k, EP assigns it an un-normalized Gaussian distribution $t_i^k \triangleq \tilde{Z}_i^k \mathcal{N}\left(\tilde{\mu}_i^k, \tilde{\sigma}_i^{k\,2}\right)$ to locally approximate its exact likelihood. In this vein, the posterior is approximated by:

$$p(\tilde{\mathbf{g}}_i|X, \mathbf{g}_i) \approx \frac{1}{Z_{\mathrm{EP}}} p(\tilde{\mathbf{g}}_i|X) \prod_{k=1}^{N} t_i^k = \mathcal{N}(\boldsymbol{\mu}_g^i, \Sigma_g^i)$$

$$\text{with } \boldsymbol{\mu}_g^i = \Sigma_g^i \tilde{\Sigma}_g^{i\,-1} \tilde{\boldsymbol{\mu}}_g^i \text{ and } \Sigma_g^i = (K + \tilde{\Sigma}_g^{i\,-1})^{-1}, \tag{13}$$

where $\tilde{\boldsymbol{\mu}}_g^i = \left(\tilde{\mu}_i^1, \cdots, \tilde{\mu}_i^N\right)^\top$, $\tilde{\Sigma}_g^i$ denotes a diagonal matrix with the k-th element being $\tilde{\sigma}_i^{k\,2}$, and Z_{EP} is the marginal likelihood. The site parameters in t_i^k of a Gaussian likelihood in equation (11) are valued by $\tilde{Z}_i^k = 1$, $\tilde{\mu}_i^k = g_i(\mathbf{x}^k)$, $\tilde{\sigma}_i^k = \sigma$. Differently, the site parameters of a non-Gaussian likelihood in equation (11) should be computed by the moment matching (Riihimäki and Vehtari 2010). Detailed formulations of this part are delineated in Section C of the supplementary document. For a candidate solution $\tilde{\mathbf{x}}$, the mean and variance of the HLGP model $\tilde{g}_i(\tilde{\mathbf{x}})$ are predicted as:

$$\begin{aligned} \mu_g^i(\tilde{\mathbf{x}}) &= m(\tilde{\mathbf{x}}) + \mathbf{k}^{*\top}(K + \tilde{\Sigma}_g^i)^{-1}\tilde{\boldsymbol{\mu}}_g^i, \\ \sigma_g^{i\,2}(\tilde{\mathbf{x}}) &= k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}) - \mathbf{k}^{*\top}(K + \tilde{\Sigma}_g^i)^{-1}\mathbf{k}^*. \end{aligned} \tag{14}$$

In principle, the hyperparameters of EP-based GP models should be updated by maximizing the marginal likelihood Z_{EP}. Since (14) resembles (2), from another perspective, EP algorithm serves as a data generator for HLGPs, i.e., assigning *virtual observations* for infeasible solutions with an estimation of noise levels. Accordingly, the hyperparameters of an HLGP model can be optimized by maximizing the marginal likelihood of a vanilla GPR model using these injected observations rather than Z_{EP} for better computation efficiency, as noted in (Rasmussen and Williams 2005).

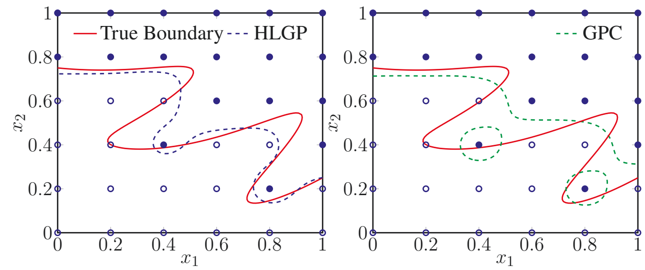

Figure 3: Different curves of feasible boundary predicted by (Left) an HLGP model versus (Right) a GPC model.

Comparison with GPC Finally, we present a brief comparison between the HLGP models in our proposed method and the GPC models frequently used in existing CBO algorithms for unobservable constraints. Notably, executing CBO routines is also feasible by only considering the truncated value distribution in equation (11). We contend that HLGP, by leveraging available observations, constructs a more reliable surrogate model than GPC. As illustrated in Figure 3, 36 equidistant solutions are evaluated, with constraint values being unobservable for half. The true boundary and the predictions made by HLGP and GPC are provided for performance assessment. The GPC model delineates the boundary by maximizing the distance between the nearest distinct solutions, mirroring the approach of an SVM (Rasmussen and Williams 2005). This results in the prediction of three disconnected feasible regions, deviating from the true feasible region. In contrast, HLGP predicts a connected feasible region that covers the majority of the true one. As such, we anticipate that HLGP, with its superior modeling capability, can enhance optimization for POCOPs.

Summary of the `CBOB` Algorithm

To improve the optimization efficiency for POCOPs, `CBOB` fully exploits available observations by re-designing a more balanced acquisition function and constructing principled surrogate models from mixed observations. Our method combines a boundary-based exploration function, as in equation (10), and non-informative likelihoods as in equation (11). Moreover, `CBOB` preserves ample flexibility in designing exploration functions and likelihoods, thus making it adaptable to individual optimization problems and suitable for further investigations.

- Building on the concept of exploring potentially feasible regions as (Lindberg and Lee 2015; Wang and Ierapetritou 2018), `CBOB` introduces extra exploration during optimization but with robust theoretical support underpinning this design. Moreover, our EICB maintains the computational efficiency of the original EIC, unlike the methods proposed in (Gelbart, Snoek, and Adams 2014; Lam and Willcox 2017; Zhang, Zhang, and Frazier 2021).
- Inspired by the level-set estimation methods (Bachoc, Cesari, and Gerchinovitz 2021), we propose an innovative exploration function as equation (10) that emphasizes potential boundaries, integrating with the POF for

efficient optimization. This design avoids aggressively evaluating unknown regions, which is a common tactic in active learning (Ranjan, Bingham, and Michailidis 2008; Bichon et al. 2008; Bect et al. 2012).

- By employing HLGPs, we build a better surrogate model for each partially observable constraint function, outperforming GPC-based methods (Lindberg and Lee 2015; Perrone et al. 2019; Ariafar et al. 2019; Bachoc, Helbert, and Picheny 2020; Candelieri 2021). With the aid of EP and its generated virtual observations, we manage to construct HLGP models with commendable computational efficiency, resulting in an improvement over other models with mixed observations (Pourmohamad and Lee 2016; Zhang, Dai, and Low 2019).

Experiment Setup

In this section, we present the experimental settings used in our empirical study.

Benchmark Suite

Our experiments consider various optimization tasks, including synthetic problems, engineering design cases, hyperparameter optimization (HPO) problems based on scikit-learn (Pedregosa et al. 2011), and reinforcement learning tasks based on Open AI Gym (Brockman et al. 2016), to constitute our benchmark suite. In addition, we consider the following two scenarios of POCOPs.

- S1: The first scenario is that only f is partially observable. Specifically, problems include 10D Keane's bump function (KBF) (Keane 1994), 4D welded beam design (WBD) (Deb 2000), 7D HPO of XGBoost on the California housing dataset with a model size constraint (XGB-H), and 12D Lunar Landing with an energy constraint (Lunar) (Eriksson et al. 2019).
- S2: The second scenario considers both f and $\vec{g}$ are partially observable. Specifically, problems include 10D Ackley function with one constraint (Ackley) (Marco et al. 2021), 4D pressure vessel design (PVD) (Coello and Mezura-Montes 2002), 8D HPO of MLP on the digits dataset with a model size constraint (MLP-D), and 16D Swimmer with an energy constraint (Swimmer) (Wang, Fonseca, and Tian 2020).

Peer Algorithms

We consider three state-of-the-art CBO methods, including `EIC` (Gardner et al. 2014) in the EI-based family, min-value entropy search with constraints (`MESC`) (Takeno et al. 2022) in the information-theoretic family, and Thompson sampling with constraints (`TSC`) (Eriksson and Poloczek 2021) in the stochastic sampling family. For S2, note that both `EIC` and `MESC` can handle binary constraint feedback (Bachoc, Helbert, and Picheny 2020; Perrone et al. 2019). For S1, all algorithms use GPR to build the surrogate models. For S2, we choose either HLGP or GPC for modeling constraints. In particular, we use a dedicated subscript to represent the corresponding surrogate model, e.g., $\texttt{EIC}_c$ and $\texttt{EIC}_h$ denote EIC with GPC and HLGP, respectively.

Algorithm		KBF	WBD	XGB-H	Lunar
`EIC`	BOV	−0.33	2.47	0.281	255
	ROF	99.3%	75.3%	<u>22.4%</u>	82.3%
`MESC`	BOV	−0.29	2.28	**0.271**	255
	ROF	99.5%	<u>18.3%</u>	58.1%	84.5%
`TSC`	BOV	−0.26	2.47	0.283	247
	ROF	99.8%	77.9%	51.2%	87.7%
`CBOB`	BOV	**−0.39**	**2.16**	0.274	**260**
	ROF	<u>99.2%</u>	67.4%	25.8%	<u>70.0%</u>

Table 1: The BOV and ROF of different algorithms in S1.

Algorithm		Ackley	PVD	MLP-D	Swim.
$\texttt{EIC}_c$	BOV	0.70	8155	0.978	322
	ROF	77.9%	5.12%	84.5%	82.3%
$\texttt{MESC}_c$	BOV	1.44	N/A	**0.983**	244
	ROF	78.2%	N/A	83.1%	83.5%
$\texttt{EIC}_h$	BOV	0.66	7598	0.981	245
	ROF	38.4%	<u>3.13%</u>	<u>64.8%</u>	<u>67.2%</u>
$\texttt{MESC}_h$	BOV	0.48	7507	0.98	197
	ROF	54.1%	3.78%	50.1%	79.8%
`CBOB`	BOV	**0.43**	**7198**	0.982	**350**
	ROF	<u>38.0%</u>	3.32%	65.5%	77.1%

Table 2: The BOV and ROF of different algorithms in S2.

General Settings

All algorithms are implemented according to their open-source code (Eriksson and Poloczek 2021; Takeno et al. 2022). In `MESC`, the optimal solutions are sampled 20 times. Both `MESC` and `TSC` sample with a grid size of 1 000. For `CBOB` with equation (10), we fix $\beta = 1.96$ to obtain a 95% confidence level. As equation (10) is a conservative design for exploration, we omit the study on more conservative behaviors with smaller β. Each experiment is independently repeated 20 times with shared random seeds. For all tasks, the Sobol sequence is used to generate $11 \times n$ initial samples, then 100 function evaluations (FEs) are performed in each experiment. Detailed settings of all algorithms and benchmark problems are presented in Section D of the supplemental document. The source code of our project is available[1].

Experiment Results

The optimization trajectories of all experiments are given in Figures 4 and 5. In addition, the median best-evaluated values (BOVs) and average ratios of feasible evaluations (ROFs) of different algorithms are presented in Tables 1 and 2. We empirically study the efficacy of `CBOB` from three aspects: i) the *improvements* on `CBOB` for EI-based CBO methods and GPC-based models; ii) the *competitiveness* of `CBOB` with other peer algorithms; and iii) the *relationship* between efficiency and exploration ability of `CBOB`.

Improvements Although EIC can be more efficient within 10 to 20 FEs, such as in WBD, XGB-H, Ackley, and MLP-D, EICB outperforms EIC in all experiments after 100 FEs,

[1]https://github.com/COLA-Laboratory/CBOB

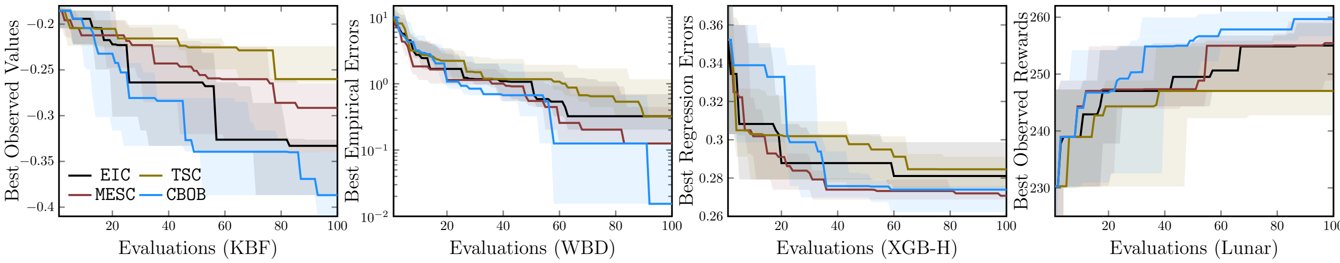

Figure 4: Optimization trajectories of different tasks in S1 (f is partially observable).

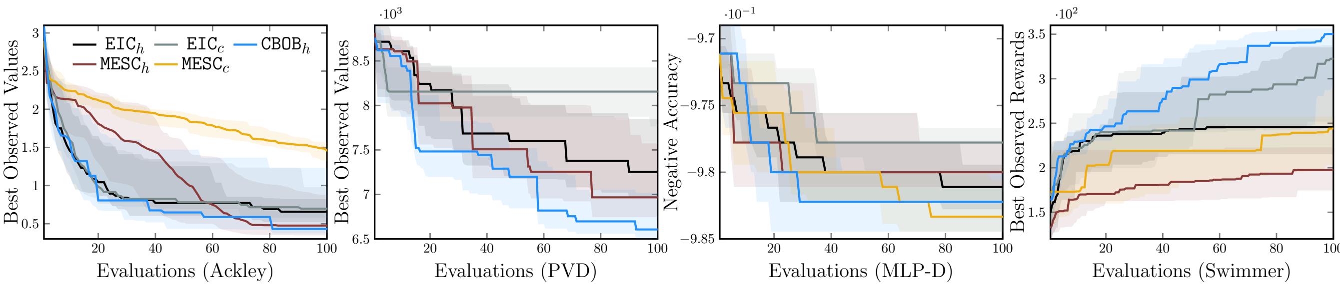

Figure 5: Optimization trajectories of different tasks in S2 (both f and $\vec{g}$ are partially observable).

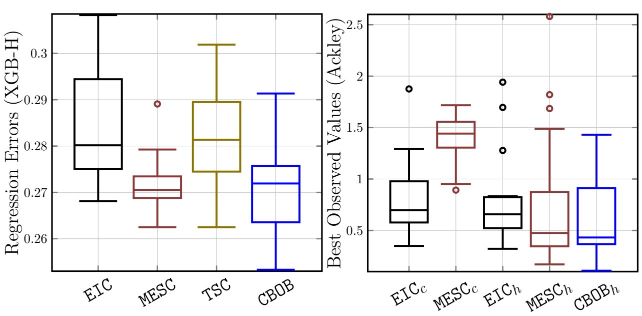

Figure 6: Box plot of the final best observed values of different algorithms on Ackley and XGB-H.

which demonstrates the design of DPOF. As the natural expense of exploration, the value deviation of EICB may be larger during the search. As given in Figure 6, despite $\texttt{EIC}_c$ and $\texttt{EIC}_h$ have smaller deviations of the best evaluated values, EICB obtains a more promising result in a statistical sense. In addition, compared to GPC, HLGP does not always improve EIC and MESC, whereas benefiting CBOB well.

Competitiveness The CBOB shows a strong competitiveness in all experiments against other CBO methods. In addition to the problems that EI-based methods perform well, such as KBF and Swimmer, CBOB remains competitive in problems that EI-based methods struggle, such as XGB-H, PVD and MLP-D. In comparison, MESC, as a promising CBO method in most problems, is inefficient in problems such as KBF and Swimmer. The TSC that showed high efficiency in fully observable environments with trust regions (Eriksson and Poloczek 2021), struggles in solving POCOPs. Moreover, CBOB and other EI-based CBO methods show less efficiency in HPO problems, which agrees with the empirical results in (Watanabe and Hutter 2023).

Relationship between efficiency and exploration In Tables 1 and 2, while CBOB obtains better feasible solutions, its ROFs are relatively low, i.e., more evaluated solutions of CBOB are infeasible. On the one hand, this agrees with the ideas of (Zhang, Zhang, and Frazier 2021) that exploration towards infeasible regions facilitates optimization efficiency of CBO methods. It also explains the larger deviation of CBOB in Figure 6 that infeasible evaluations return little information in POCOPs. Besides, we find this exploration effective as the number of outliers of CBOB reduces, such as MESC_h and CBOB in Ackley of Figure 6. On the other hand, since EI has already considered exploration, EIC can also have fewer ROFs during the search. Differently, in order not to break the balance of EI between exploration and exploitation, CBOB assigns more balanced weights for constraint handling. We highlight that this idea can also be integrated with other acquisition functions, such as the probability of improvement and parzen estimator (Garnett 2023).

Concluding Remarks

This paper designs CBOB that fully exploits the available observations for better exploration and surrogate modeling, by both theoretical and empirical analysis, demonstrating that CBOB has the potential to be a promising CBO method for POCOPs. Further investigations include the in-depth theoretical study of CBOB and the design of the exploration functions that are suitable for individual problems. We will also endeavor to propose a risk-aware improvement of EICB regarding robustness and the reduction of infeasible evaluations, contributing to more practical optimization scenarios.

Acknowledgments

This work was supported in part by the UKRI Future Leaders Fellowship under Grant MR/S017062/1 and MR/X011135/1; in part by NSFC under Grant 62376056 and 62076056; in part by the Royal Society under Grant IES/R2/212077; in part by the EPSRC under Grant 2404317; in part by the Kan Tong Po Fellowship (KTP\R1\231017); and in part by the Amazon Research Award and Alan Turing Fellowship.

References

Ariafar, S.; Coll-Font, J.; Brooks, D. H.; and Dy, J. G. 2019. ADMMBO: Bayesian Optimization with Unknown Constraints using ADMM. *J. Mach. Learn. Res.*, 20: 123:1–123:26.

Bachoc, F.; Cesari, T.; and Gerchinovitz, S. 2021. The Sample Complexity of Level Set Approximation. In *AISTATS'21: Proc. of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130, 424–432. PMLR.

Bachoc, F.; Helbert, C.; and Picheny, V. 2020. Gaussian Process Optimization with Failures: Classification and Convergence Proof. *J. Glob. Optim.*, 78(3): 483–506.

Bect, J.; Ginsbourger, D.; Li, L.; Picheny, V.; and Vázquez, E. 2012. Sequential Design of Computer Experiments for the Estimation of a Probability of Failure. *Stat. Comput.*, 22(3): 773–793.

Belakaria, S.; Deshwal, A.; and Doppa, J. R. 2019. Max-value Entropy Search for Multi-Objective Bayesian Optimization. In *NeurIPS'19: Annual Conference on Neural Information Processing Systems*, 7823–7833.

Bichon, B. J.; Eldred, M. S.; Swiler, L. P.; Mahadevan, S.; and McFarland, J. M. 2008. Efficient Global Reliability Analysis for Nonlinear Implicit Performance Functions. *AIAA J.*, 46(10): 2459–2468.

Brockman, G.; Cheung, V.; Pettersson, L.; Schneider, J.; Schulman, J.; Tang, J.; and Zaremba, W. 2016. OpenAI Gym. arXiv:1606.01540.

Candelieri, A. 2021. Sequential Model based Optimization of Partially Defined Functions Under Unknown Constraints. *J. Glob. Optim.*, 79(2): 281–303.

Coello, C. A. C.; and Mezura-Montes, E. 2002. Constraint-Handling in Genetic Algorithms Through the use of Dominance-Based Tournament Selection. *Adv. Eng. Informatics*, 16(3): 193–203.

Deb, K. 2000. An Efficient Constraint Handling Method for Genetic Algorithms. *Computer Methods in Applied Mechanics and Engineering*, 186(2): 311–338.

Eriksson, D.; Pearce, M.; Gardner, J. R.; Turner, R.; and Poloczek, M. 2019. Scalable Global Optimization via Local Bayesian Optimization. In *NeurIPS'19: Annual Conference on Neural Information Processing Systems*, 5497–5508.

Eriksson, D.; and Poloczek, M. 2021. Scalable Constrained Bayesian Optimization. In *AISTATS'21: Proc. of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130, 730–738. PMLR.

Gardner, J. R.; Kusner, M. J.; Xu, Z. E.; Weinberger, K. Q.; and Cunningham, J. P. 2014. Bayesian Optimization with Inequality Constraints. In *ICML'14: Proc. of the 31th International Conference on Machine Learning*, volume 32, 937–945. JMLR.org.

Garnett, R. 2023. *Bayesian Optimization*. Cambridge University Press.

Gelbart, M. A.; Snoek, J.; and Adams, R. P. 2014. Bayesian Optimization with Unknown Constraints. In *UAI'14: Proc. of the 13th Conference on Uncertainty in Artificial Intelligence*, 250–259. AUAI Press.

Gramacy, R. B.; Gray, G. A.; Digabel, S. L.; Lee, H. K. H.; Ranjan, P.; Wells, G. N.; and Wild, S. M. 2016. Modeling an Augmented Lagrangian for Black-box Constrained Optimization. *Technometrics*, 58(1): 1–11.

Hernández-Lobato, J. M.; Gelbart, M. A.; Hoffman, M. W.; Adams, R. P.; and Ghahramani, Z. 2015. Predictive Entropy Search for Bayesian Optimization with Unknown Constraints. In *ICML'15: Proc. of the 32nd International Conference on Machine Learning 2015*, volume 37, 1699–1707. JMLR.org.

Jones, D. R.; Schonlau, M.; and Welch, W. J. 1998. Efficient Global Optimization of Expensive Black-Box Functions. *J. Glob. Optim.*, 13(4): 455–492.

Keane, A. J. 1994. Experiences with Optimizers in Structural Design. In *Proc. of the conference on adaptive computing in engineering design and control*, volume 94, 14–27.

Lam, R.; and Willcox, K. 2017. Lookahead Bayesian Optimization with Inequality Constraints. In *NeurIPS'17: Annual Conference on Neural Information Processing Systems*, 1890–1900.

Letham, B.; Karrer, B.; Ottoni, G.; and Bakshy, E. 2019. Constrained Bayesian Optimization with Noisy Experiments. *Bayesian Anal.*, 14(2): 495 – 519.

Lindberg, D. V.; and Lee, H. K. 2015. Optimization Under Constraints by Applying an Asymmetric Entropy Measure. *J. Comput. Graph. Stat.*, 24(2): 379–393.

Marco, A.; Baumann, D.; Khadiv, M.; Hennig, P.; Righetti, L.; and Trimpe, S. 2021. Robot Learning With Crash Constraints. *IEEE Robotics Autom. Lett.*, 6(2): 1439–1446.

Minka, T. P. 2001. Expectation Propagation for Approximate Bayesian Inference. In *UAI '01: Proc. of the 17th Conference in Uncertainty in Artificial Intelligence*, 362–369. Morgan Kaufmann.

Parr, J. M.; Keane, A. J.; Forrester, A. I.; and Holden, C. M. 2012. Infill Sampling Criteria for Surrogate-Based Optimization with Constraint Handling. *Eng. Optim.*, 44(10): 1147–1166.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; VanderPlas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.*, 12: 2825–2830.

Perrone, V.; Shcherbatyi, I.; Jenatton, R.; Archambeau, C.; and Seeger, M. W. 2019. Constrained Bayesian

Optimization with Max-Value Entropy Search. *CoRR*, abs/1910.07003.

Picheny, V. 2014. A Stepwise Uncertainty Reduction Approach to Constrained Global Optimization. In *AISTATS'14: Proc. of the 17th International Conference on Artificial Intelligence and Statistics*, volume 33, 787–795. JMLR.org.

Picheny, V.; Gramacy, R. B.; Wild, S. M.; and Digabel, S. L. 2016. Bayesian Optimization Under Mixed Constraints with a Slack-Variable Augmented Lagrangian. In *NeurIPS'16: Annual Conference on Neural Information Processing Systems*, 1435–1443.

Pourmohamad, T.; and Lee, H. K. H. 2016. Multivariate Stochastic Process Models for Correlated Responses of Mixed Type. *Bayesian Anal.*, 11: 797 – 820.

Ranjan, P.; Bingham, D.; and Michailidis, G. 2008. Sequential Experiment Design for Contour Estimation From Complex Computer Codes. *Technometrics*, 50(4): 527–541.

Rasmussen, C. E.; and Williams, C. K. I. 2005. *Gaussian Processes for Machine Learning*. The MIT Press. ISBN 9780262256834.

Riihimäki, J.; and Vehtari, A. 2010. Gaussian Processes with Monotonicity Information. In *AISTATS'10: Proc. of the 13th International Conference on Artificial Intelligence and Statistics*, 645–652. JMLR.org.

Schonlau, M.; Welch, W. J.; and Jones, D. R. 1998. *Global Versus Local Search in Constrained Optimization of Computer Models*, volume 34, 11–25. Institute of Mathematical Statistics. ISBN 0-940600-46-3.

Takeno, S.; Tamura, T.; Shitara, K.; and Karasuyama, M. 2022. Sequential and Parallel Constrained Max-value Entropy Search via Information Lower Bound. In *ICML'22, Proc. of the 39th International Conference on Machine Learning*, volume 162, 20960–20986. PMLR.

Vazquez, E.; and Bect, J. 2010. Convergence Properties of the Expected Improvement Algorithm with Fixed Mean and Covariance Functions. *J. Stat. Plan. Inference*, 140(11): 3088–3095.

Wang, L.; Fonseca, R.; and Tian, Y. 2020. Learning Search Space Partition for Black-box Optimization using Monte Carlo Tree Search. In *NeurIPS'20: Annual Conference on Neural Information Processing Systems*.

Wang, Z.; and Ierapetritou, M. 2018. Constrained Optimization of Black-box Stochastic Systems Using a Novel Feasibility Enhanced Kriging-Based Method. *Comput. Chem. Eng.*, 118: 210–223.

Watanabe, S.; and Hutter, F. 2023. c-TPE: Tree-structured Parzen Estimator with Inequality Constraints for Expensive Hyperparameter Optimization. In *IJCAI'23: Proc. of the 32nd International Joint Conference on Artificial Intelligence*, 4371–4379.

Zhang, Y.; Dai, Z.; and Low, B. K. H. 2019. Bayesian Optimization with Binary Auxiliary Information. In *UAI'19: Proc. of the 35th Conference on Uncertainty in Artificial Intelligence*, volume 115, 1222–1232. AUAI Press.

Zhang, Y.; Zhang, X.; and Frazier, P. I. 2021. Two-step Lookahead Bayesian Optimization with Inequality Constraints. In *NeurIPS'21: Annual Conference on Neural Information Processing Systems*, 12563–12575.