

---

# Sequential- and Parallel- Constrained Max-value Entropy Search via Information Lower Bound

---

Shion Takeno<sup>1,2</sup> Tomoyuki Tamura<sup>3</sup> Kazuki Shitara<sup>4,5</sup> Masayuki Karasuyama<sup>1</sup>

## Abstract

Max-value entropy search (MES) is one of the state-of-the-art approaches in Bayesian optimization (BO). In this paper, we propose a novel variant of MES for constrained problems, called Constrained MES via Information lower BOund (CMES-IBO), that is based on a Monte Carlo (MC) estimator of a lower bound of a mutual information (MI). Unlike existing studies, our MI is defined so that uncertainty with respect to feasibility can be incorporated. We derive a lower bound of the MI that guarantees non-negativity, while a constrained counterpart of conventional MES can be negative. We further provide theoretical analysis that assures the low-variability of our estimator which has never been investigated for any existing information-theoretic BO. Moreover, using the conditional MI, we extend CMES-IBO to the parallel setting while maintaining the desirable properties. We demonstrate the effectiveness of CMES-IBO by several benchmark functions and real-world problems.

## 1. Introduction

Bayesian optimization (BO) has been widely studied as an effective framework for expensive black-box optimization problems. On the other hand, additional unknown *constraints* often exist in real-world problems of a variety of fields such as scientific experiments, industrial product designs, and automatic machine learning (AutoML). For example, materials discovery can be seen as an optimization

of a physical property of materials, such as conductivity, under constraints derived from other physical properties, such as stability. In the case of AutoML, for example, constraints are useful to guarantee the classification accuracy of minority classes for class-imbalanced datasets. However, constrained black-box optimization is often quite difficult because, in many practical problems, a measurement comes at a high cost, and further, functional forms of both the objective and the constraints are unknown. To incorporate unknown constraint functions, *constrained BO* (CBO) has also been studied (Gardner et al., 2014; Gelbart et al., 2014; Schonlau et al., 1998; Snoek, 2013). CBO tries to achieve the sample-efficient constrained optimization based on Bayesian models for objective and constraint functions.

Max-value entropy search (MES) (Wang & Jegelka, 2017) is one of the state-of-the-art approaches in BO. The basic idea is to maximize the mutual information (MI) between a querying point and the optimal objective value. For constrained problems, Perrone et al. (2019) have proposed an extension of MES, but it is restricted to only one constraint (note that although their main focus is on a setting in which only a binary indicator of feasibility is observed for constraints, they also show the ‘real-valued feedback’ case in the appendix). Although it is possible to consider a multiple constraint extension of this approach, which we call constrained MES (CMES), the resulting MI approximation can have a negative value (though the MI should be non-negative) when the number of constraints is more than 5 as we will show in Section 4.1. A more comprehensive review of related studies is shown in Section 5.

In this paper, we propose a novel information-theoretic CBO method called *Constrained Max-value Entropy Search via Information lower BOund* (CMES-IBO), which is based on a Monte Carlo (MC) estimator of a lower bound of an MI. We first define the random ‘max-value’ of a constrained problem so that uncertainty with respect to feasibility (whether the problem has a feasible region or not) can be incorporated. Uncertainty of feasibility has not been considered by existing information-theoretic CBO studies, though for constrained problems in general, to identify the feasibility of the problem is a key issue known as the feasibility problem (Chinneck, 2007). Unlike CMES, our ac-

---

<sup>1</sup>Department of Computer Science, Nagoya Institute of Technology, Aichi, Japan <sup>2</sup>Center for Advanced Intelligence Project, RIKEN, Tokyo, Japan <sup>3</sup>Department of Physical Science and Engineering, Nagoya Institute of Technology, Aichi, Japan <sup>4</sup>Joining and Welding Research Institute, Osaka University, Osaka, Japan <sup>5</sup>Nanostructures Research Laboratory, Japan Fine Ceramics Center, Aichi, Japan. Correspondence to: Masayuki Karasuyama <karasuyama@nitech.ac.jp>.

quisition function guarantees non-negativity. We theoretically and empirically verify that CMES-IBO can provide reasonable acquisition function values even when CMES provides negative values. Further, we provide theoretical analysis that assures the low-variability of our estimator, such as a concentration bound achieving an exponentially fast convergence to the true lower bound in the number of MC samples, which has never been investigated for any existing information-theoretic BO. Moreover, using the conditional MI, we extend CMES-IBO to the parallel setting while maintaining the desirable properties.

Our main contributions are summarized as follows:

1. We develop an MI lower bound based CBO called CMES-IBO. Our MI is well-defined even when a feasible region can be empty. The resulting MI approximation is quite simple and guarantees non-negativity.
2. We reveal advantageous properties of CMES-IBO compared with the direct MES extension (CMES), such as the smaller worst-case error of the MI estimation. We also derive an estimation variability of CMES-IBO that suggests robustness with respect to the number of MC samples.
3. We further develop a parallel extension of CMES-IBO, in which multiple queries can be issued simultaneously. The extension is based on a combination of the conditional MI and the greedy selection.

We demonstrate the effectiveness of CMES-IBO by benchmark and real-world functions.

## 2. Preliminary

We are interested in the maximum value of the objective function  $f : \mathcal{X} \mapsto \mathbb{R}$  under  $C$  constraint functions  $g_c : \mathcal{X} \mapsto \mathbb{R}$  for  $c = 1, \dots, C$ , where  $\mathcal{X} \subset \mathbb{R}^d$  is an input space. Let  $\mathcal{X}_{\text{feasible}} := \{\mathbf{x} \mid g_c(\mathbf{x}) \geq z_c, c = 1, \dots, C\}$  be the feasible region, where  $z_c \in \mathbb{R}$  is a given constant. The optimal solution of this constrained optimization problem is written as  $\mathbf{x}_* := \arg \max_{\mathbf{x} \in \mathcal{X}_{\text{feasible}}} f(\mathbf{x})$ . We consider the setting that the functions  $f$  and  $g_1, \dots, g_C$  are unknown but can be evaluated simultaneously for any given  $\mathbf{x} \in \mathcal{X}$  with high observation cost. We assume that  $f$  and  $g_1, \dots, g_C$  follow independent *Gaussian processes* (GPs) (Rasmussen & Williams, 2005). Suppose that the observations are contaminated with additive Gaussian noises as  $y_{\mathbf{x}}^{(f)} = f(\mathbf{x}) + \varepsilon$  and  $y_{\mathbf{x}}^{(g_c)} = g_c(\mathbf{x}) + \varepsilon$  for  $\forall c$ , where  $\varepsilon \sim \mathcal{N}(0, \sigma_{\text{noise}}^2)$ . For the GP prior, we employ  $\mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$ , which indicates that the prior mean is 0 and covariance is specified by a kernel function  $k : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$ . Then, the predictive distribution for  $f(\mathbf{x})$  given  $\mathcal{D}_n^{(f)} = \{(\mathbf{x}_i, y_{\mathbf{x}_i}^{(f)})\}_{i=1}^n$  is derived

as follows:

$$\begin{aligned} f(\mathbf{x}) | \mathcal{D}_n^{(f)} &\sim \mathcal{N}(\mu^{(f)}(\mathbf{x}), \sigma^{(f)2}(\mathbf{x})), \\ \mu^{(f)}(\mathbf{x}) &= \mathbf{k}^\top (\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I})^{-1} \mathbf{y}^{(f)}, \\ \sigma^{(f)2}(\mathbf{x}) &= k(\mathbf{x}, \mathbf{x}) - \mathbf{k}^\top (\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I})^{-1} \mathbf{k}, \end{aligned}$$

where  $\mathbf{y}^{(f)} = (y_{\mathbf{x}_1}^{(f)}, \dots, y_{\mathbf{x}_n}^{(f)})^\top$ ,  $\mathbf{k} = (k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_n))^\top$ ,  $\mathbf{K} \in \mathbb{R}^{n \times n}$  is a kernel matrix whose  $(i, j)$  element is defined as  $k(\mathbf{x}_i, \mathbf{x}_j)$ , and  $\mathbf{I} \in \mathbb{R}^{n \times n}$  is the identity matrix. The predictive distributions  $\mathcal{N}(\mu^{(g_c)}(\mathbf{x}), \sigma^{(g_c)2}(\mathbf{x}))$  for  $g_c$  can be obtained in the same way using the observations  $y_{\mathbf{x}_1}^{(g_c)}, \dots, y_{\mathbf{x}_n}^{(g_c)}$ . Note that, for simplicity, we here assume the same kernels and noise variances for all the functions, but both the settings can be changed depending on the function.

## 3. Constrained Max-value Entropy Search via Information Lower Bound

### 3.1. Max-value of Constrained Problem

First, we consider the definition of a random variable  $f_*$ , which is the max-value of a constraint optimization problem. The straightforward definition is  $\max_{\mathbf{x} \in \mathcal{X}_{\text{feasible}}} f(\mathbf{x})$ , employed in a prior work on the MES-based CBO (Perrone et al., 2019). However, this definition is not well-defined as a proper random variable when  $\mathcal{X}_{\text{feasible}}$  can be empty, i.e., when  $\Pr(\mathcal{X}_{\text{feasible}} = \emptyset) > 0$ . Note that  $\mathcal{X}_{\text{feasible}}$  is randomly determined through GPs  $g_1, \dots, g_C$ , and thus,  $\Pr(\mathcal{X}_{\text{feasible}} = \emptyset) > 0$  can occur even when the underlying unknown true constraint functions have a non-empty feasible region. Thus, we define  $f_*$  as follows:

$$f_* := \begin{cases} \max_{\mathbf{x} \in \mathcal{X}_{\text{feasible}}} f(\mathbf{x}), & \text{if } \mathcal{X}_{\text{feasible}} \neq \emptyset, \\ -\infty, & \text{otherwise.} \end{cases} \quad (1)$$

Regarding the optimal value as  $-\infty$  when no feasible solution exists is a common convention in classical optimization literature (Boyd & Vandenberghe, 2004). The schematic illustration of  $f_*$  is shown in Fig. 1 (a). The figure indicates that our definition of  $p(f_*)$  can incorporate the uncertainty about whether the problem is feasible or not.

### 3.2. Acquisition Function of CMES-IBO

Let

$$\mathbf{h}_{\mathbf{x}} := (f(\mathbf{x}), g_1(\mathbf{x}), \dots, g_C(\mathbf{x}))^\top \in \mathbb{R}^{C+1},$$

be a vector concatenating the objective and all the constraint function values. Suppose that we already have the dataset  $\mathcal{D}_{t-1}$ , which contains observations for the past  $t-1$  queries. We consider selecting a next query  $\mathbf{x}_t$  by maximiz-

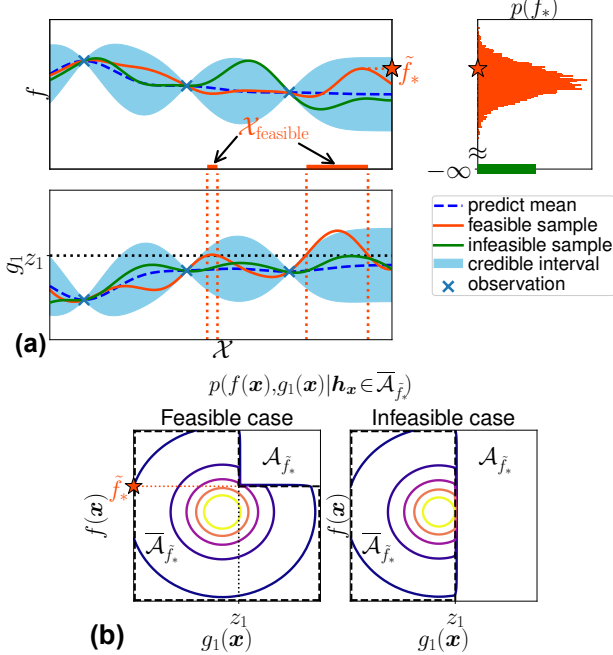


Figure 1: Generative process of the optimal value  $f_*$  and truncation of the predictive distribution. (a) In the left two plots, two sets of sample paths (red and green) are generated from GPs. The red sample path of  $g_1$  has the non-empty feasible region  $\mathcal{X}_{\text{feasible}}$  while the green sample path has an empty feasible region. The red star  $\tilde{f}_*$  represents the optimal value of the red sample path, and for the green sample path, the optimal value is defined as  $-\infty$  as shown in (1). The distribution of the optimal value is plotted in the upper right figure. (b) The two contour plots represent the density function  $p(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{\tilde{f}_*}) = p(f(\mathbf{x}), g_1(\mathbf{x}) | \mathbf{h}_x \in \bar{\mathcal{A}}_{\tilde{f}_*})$  for the feasible and infeasible cases, respectively. For a given  $\tilde{f}_*$ , the density should be 0 in  $\mathcal{A}_{\tilde{f}_*}$ , because  $f(\mathbf{x})$  cannot be larger than the given  $\tilde{f}_*$  for  $\mathbf{x}$  such that  $g_1(\mathbf{x}) \geq z_1$ .

ing the MI<sup>1</sup> between  $\mathbf{h}_x$  and  $f_*$  given  $\mathcal{D}_{t-1}$ :

$$\text{MI}(\mathbf{h}_x; f_* | \mathcal{D}_{t-1}). \quad (2)$$

Hereafter, we omit the conditioning by  $\mathcal{D}_{t-1}$  in MI, entropy, density, and probability when it is obvious from the context. For example,  $p(\mathbf{h}_x | f_*)$  indicates  $p(\mathbf{h}_x | \mathcal{D}_{t-1}, f_*)$ . In MI, we use noiseless  $\mathbf{h}_x$  instead of noisy  $(y_{x_i}^{(f)}, y_{x_i}^{(g_1)}, \dots, y_{x_i}^{(g_C)})$  for computational simplicity.

Since directly evaluating the MI (2) is computationally in-

<sup>1</sup>Note that although  $p(f_*)$  can be seen as a mixture distribution in which one of the components (the green distribution in Fig. 1) consists of a single constant value ( $-\infty$ ), MI still can be defined in this case. Details are shown in Appendix A.

tractable, we consider a lower bound derived as follows:

$$\begin{aligned} \text{MI}(\mathbf{h}_x; f_*) &= \mathbb{E}_{f_*} \left[ \mathbb{E}_{\mathbf{h}_x | f_*} \left[ \log \frac{p(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} \right] \right] \\ &= \mathbb{E}_{f_*} \left[ \mathbb{E}_{\mathbf{h}_x | f_*} \left[ \log \frac{q(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} \right] + D_{\text{KL}}(p(\mathbf{h}_x | f_*) || q(\mathbf{h}_x | f_*)) \right] \\ &\geq \mathbb{E}_{f_*} \left[ \int p(\mathbf{h}_x | f_*) \log \frac{q(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} d\mathbf{h}_x \right], \end{aligned} \quad (3)$$

where  $D_{\text{KL}}(\cdot)$  is Kullback-Leibler (KL) divergence and  $q(\mathbf{h}_x | f_*)$  is an arbitrary probability density function (PDF) that has same support as  $p(\mathbf{h}_x | f_*)$ . See Appendix A for the derivation. The difference between the lower bound and the MI is given as  $\mathbb{E}_{f_*} [D_{\text{KL}}(p(\mathbf{h}_x | f_*) || q(\mathbf{h}_x | f_*))]$ , and thus, the equality holds if  $p(\mathbf{h}_x | f_*)$  is equal to  $q(\mathbf{h}_x | f_*)$  for  $\forall f_*$ .

To use (3), we have to specify the distribution  $q(\mathbf{h}_x | f_*)$ . For this type of KL-based lower bounds, the variational approximation is a well-known approach by which the lower bound is maximized with respect to  $q(\mathbf{h}_x | f_*)$ . However, the expectation and the integral with respect to  $f_*$  and  $\mathbf{h}_x$  in (3) make this optimization computationally too complicated. Instead, we consider setting  $q(\mathbf{h}_x | f_*)$  as a specific distribution that can be easily computed.

The optimal selection of  $q(\mathbf{h}_x | f_*)$  is  $p(\mathbf{h}_x | f_*)$ , i.e., predictive distributions of GPs given the optimal value  $f_*$  which we call a  $\text{GP}|_{f_*}$ . For all existing MES-based approaches, the  $\text{GP}|_{f_*}$  is required in a slightly different context (see Section 4.1 for the case of the existing MES for constrained problems). Since the  $\text{GP}|_{f_*}$  is difficult to evaluate, it is approximated by a distribution defined by truncating the original GP predictive distribution at  $f_*$  in almost all the existing methods (e.g., Wang & Jegelka, 2017; Takeno et al., 2020; 2022; Perrone et al., 2019; Suzuki et al., 2020). For example, in the case of the original MES,  $p(f(\mathbf{x}) | f_*)$  is replaced with the truncated normal distribution  $p(f(\mathbf{x}) | f(\mathbf{x}) \leq f_*)$ , as pointed out in (Takeno et al., 2020). This replacement makes the resulting acquisition function simple and practical performance of MES-based approaches with this replacement has been repeatedly shown. Therefore, we follow this truncation-based approach.

In the case of constrained problems, when the optimal value  $f_*$  is fixed,  $f(\mathbf{x}) > f_*$  does not hold if  $g_c(\mathbf{x}) \geq z_c$  for  $\forall c$  (i.e.,  $\mathbf{x}$  is feasible). This means that  $f(\mathbf{x})$  and  $g_c(\mathbf{x})$  are jointly truncated so that their densities become 0 in  $\mathcal{A}_{f_*} := (f_*, \infty) \times (z_1, \infty) \times \dots \times (z_C, \infty) \subset \mathbb{R}^{C+1}$ . In other words,  $\mathbf{h}_x$  should be in  $\bar{\mathcal{A}}_{f_*} \subset \mathbb{R}^{C+1}$ , which is a complementary set of  $\mathcal{A}_{f_*}$ , as illustrated in Fig. 1 (b). Thus, we set  $q(\mathbf{h}_x | f_*)$  as the following *truncated multi-*

variate normal (TMN) distribution:

$$q(\mathbf{h}_x | f_*) = p(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*}) = \begin{cases} p(\mathbf{h}_x) / \bar{Z}_x(f_*) & \text{for } \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*}, \\ 0 & \text{for } \mathbf{h}_x \in \mathcal{A}_{f_*}, \end{cases} \quad (4)$$

where  $\bar{Z}_x(f_*) = \Pr(\mathbf{h}_x \in \bar{\mathcal{A}}_{f_*})$  is the normalization constant.

By substituting TMN (4) into (3), we obtain a lower bound  $L(x)$  as

$$L(x) := \mathbb{E}_{f_*} \left[ \int_{\bar{\mathcal{A}}_{f_*}} p(\mathbf{h}_x | f_*) \log \frac{p(\mathbf{h}_x) / \bar{Z}_x(f_*)}{p(\mathbf{h}_x)} d\mathbf{h}_x \right] = \mathbb{E}_{f_*} [-\log \bar{Z}_x(f_*)]. \quad (5)$$

See Appendix A for the detailed derivation. In (5),  $C + 1$  dimensional integral with respect to  $\mathbf{h}_x$  is canceled out, which is one of important benefits of our selection of  $q(\mathbf{h}_x | f_*)$ . By applying the MC approximation to  $\mathbb{E}_{f_*}$  in (5), we obtain the CMES-IBO acquisition function

$$\alpha^{\text{IBO}}(x) = -\frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} \log \bar{Z}_x(\tilde{f}_*), \quad (6)$$

where  $\mathcal{F}_*$  is a set of sampled  $f_*$  from the current GPs,  $\tilde{f}_* \in \mathcal{F}_*$  represents each one of sampled values, and  $K = |\mathcal{F}_*|$ .

Let  $Z_x(\tilde{f}_*) := \Pr(\mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) = 1 - \bar{Z}_x(\tilde{f}_*)$ . For a given  $\tilde{f}_* \in \mathcal{F}_*$ , we can compute  $\bar{Z}_x(\tilde{f}_*)$  easily by using

$$\begin{aligned} \bar{Z}_x(\tilde{f}_*) &= 1 - Z_x(\tilde{f}_*) \\ &= 1 - \Pr(f(x) \geq \tilde{f}_*) \prod_{c=1}^C \Pr(g(x) \geq z_c) \\ &= 1 - (1 - \Phi(\gamma_x^{(f)}(\tilde{f}_*))) \prod_{c=1}^C (1 - \Phi(\gamma_x^{(g_c)}(z_c))), \end{aligned}$$

where  $\gamma_x^{(f)}(\tilde{f}_*) := (\tilde{f}_* - \mu^{(f)}(x)) / \sigma^{(f)}(x)$ ,  $\gamma_x^{(g_c)}(z_c) := (z_c - \mu^{(g_c)}(x)) / \sigma^{(g_c)}(x)$ , and  $\Phi$  is the cumulative distribution function (CDF) of the standard normal distribution. This decomposition is possible because  $f$  and  $g_c$  for  $\forall c$  are assumed to be independent, and  $\mathcal{A}_{\tilde{f}_*}$  is a hyperrectangle.

Consequently, the functional form of our acquisition function (6) becomes quite simple. The algorithm and computational complexity for our CMES-IBO is shown in Appendix G. Although we assume that GPs are independent of each other, CMES-IBO can also apply to the correlated setting easily by computing  $Z_x(\tilde{f}_*) = \Pr(\mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  as the CDF of the multivariate (correlated) normal distribution.

### 3.3. Sampling $f_*$

When  $\mathcal{X}$  is a finite set of a moderate size,  $f_*$  can be sampled from the joint predictive distributions of  $\mathbf{h}_x$  for all candidate points  $x \in \mathcal{X}$ . Otherwise (such as a continuous  $\mathcal{X}$ ), we employ an approach using the random Fourier

feature (RFF) (Rahimi & Recht, 2008). We first generate approximate sample paths of objective and constraint functions from Bayesian linear regression models using RFF, and then, we solve a (white-box) constrained optimization problem defined by the generated sample paths by using a general solver (see Appendix G.1 for details). A similar approach has been commonly employed in the entropy-based BO (e.g., Wang & Jegelka, 2017). Particularly when the dataset  $\mathcal{D}_{t-1}$  does not contain any feasible solution (which typically occurs at the beginning of the optimization), sampled  $\mathcal{X}_{\text{feasible}}$  is often empty, as shown in the green sample path in Fig. 1 (b). Therefore, if the solver cannot find any feasible solution,  $\tilde{f}_*$  is set as  $-\infty$  as defined in (1).

### 3.4. Parallelization

We consider a parallel extension, in which  $Q$  queries can be evaluated at every iteration. Due to a space limitation, we here only describe the main idea briefly (see Appendix B for more detail). We employ a greedy selection of the  $Q$  queries using the *conditional mutual information* (CMI). Let  $\mathcal{X}_q := \{x^{(1)}, \dots, x^{(q)}\}$  and  $\mathcal{H}_q := \{\mathbf{h}_{x^{(1)}}, \dots, \mathbf{h}_{x^{(q)}}\}$  be sets of  $q$  inputs and output vectors already selected in the greedy procedure ( $q < Q$ ), respectively. To maximize the MI after adding the  $(q + 1)$ -th query  $x$ , i.e.,  $\text{MI}(\mathcal{H}_q \cup \mathbf{h}_x; f_*)$ , we consider the decomposition

$$\text{MI}(\mathcal{H}_q \cup \mathbf{h}_x; f_*) = \text{MI}(\mathcal{H}_q; f_*) + \text{CMI}(\mathbf{h}_x; f_* | \mathcal{H}_q),$$

where  $\text{CMI}(\mathbf{h}_x; f_* | \mathcal{H}_q) := \mathbb{E}_{\mathcal{H}_q} [\text{MI}(\mathbf{h}_x; f_* | \mathcal{H}_q)]$ . Since the first term does not depend on  $x$ , we only need to maximize  $\text{CMI}(\mathbf{h}_x; f_* | \mathcal{H}_q)$ .

By the same approach as (5), we can obtain the lower bound of  $\text{CMI}(\mathbf{h}_x; f_* | \mathcal{H}_q)$  as

$$L(x | \mathcal{X}_q) := -\mathbb{E}_{\mathcal{H}_q, f_*} [\log \bar{Z}_x(f_* | \mathcal{H}_q)], \quad (7)$$

where  $\bar{Z}_x(f_* | \mathcal{H}_q) := \Pr(\mathbf{h}_x \in \bar{\mathcal{A}}_{f_*} | \mathcal{X}_q, \mathcal{H}_q)$ . See Appendix A for the derivation. By applying the MC estimation, we obtain the acquisition function for the  $(q + 1)$ -th query as

$$\alpha^{\text{IBO}}(x | \mathcal{X}_q) = -\frac{1}{K} \sum_{(\tilde{f}_*, \mathcal{H}_q) \in \mathcal{J}} \log \bar{Z}_x(\tilde{f}_* | \mathcal{H}_q), \quad (8)$$

where  $\mathcal{J}$  is a set of  $K$  sampled  $(\tilde{f}_*, \mathcal{H}_q)$  from the current GPs. Note that jointly sampling  $\tilde{f}_*$  and  $\mathcal{H}_q$  can be easily performed by almost the same procedure as the sequential case, and given  $\tilde{f}_*$  and  $\mathcal{H}_q$ ,  $\bar{Z}_x(\tilde{f}_* | \mathcal{H}_q)$  can be easily computed through the GPs updated by  $\mathcal{H}_q$  (details are in Appendix G.1 and B, respectively). Consequently, our parallel extension is also reduced to a simple acquisition function.



## 4. Analysis

We here consider differences of CMES-IBO from a more direct application of the conventional MES, and we also provide analysis on the estimation variability of our estimator. Although we focus on the sequential case, both of them also hold in the parallel case.

### 4.1. Comparison with Conventional MES

Although we employ the lower bound based approach to derive CMES-IBO, Perrone et al. (2019) have proposed a more direct extension of MES to CBO (Note that Perrone et al. (2019) mainly focus on the binary setting, in which a binary value indicating feasible or not is only observed). However, their derivation is restricted to only one constraint ( $C = 1$ ). To consider the relation with CMES-IBO, we extend this original MES-based approach to the multiple constraints case, which we refer to as CMES.

In CMES, the MI is approximated as

$$\begin{aligned} \text{MI}(\mathbf{h}_x; f_*) &= H(\mathbf{h}_x) - \mathbb{E}_{f_*} [H(\mathbf{h}_x | f_*)] \\ &\approx H(\mathbf{h}_x) - \mathbb{E}_{f_*} [H(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*})] \quad (9) \\ &\approx \frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} \left\{ \frac{Z_x(\tilde{f}_*)}{2\bar{Z}_x(\tilde{f}_*)} R_{\tilde{f}_*} - \log \bar{Z}_x(\tilde{f}_*) \right\} \\ &=: \alpha^{\text{CMES}}(\mathbf{x}), \quad (10) \end{aligned}$$

where

$$R_{\tilde{f}_*} = \frac{\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*) \phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))}{1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))} + \sum_{c=1}^C \frac{\gamma_{\mathbf{x}}^{(g_c)}(z_c) \phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}{1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))},$$

and  $H(\cdot)$  is the differential entropy and  $\phi$  is the PDF of the standard normal distribution. See Appendix E for the derivation. The first approximation (9) is the replacement from  $p(\mathbf{h}_x | f_*)$  to  $p(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*})$  in the entropy of the second term, and the second approximation that derives the next line is the MC estimation of  $\mathbb{E}_{f_*}$  (Note that although the approximation strategy follows (Perrone et al., 2019),  $f_*$  here is based on our definition (1) because  $f_*$  should be a proper random variable to define the MI as we mentioned in Section 3.1).

Although CMES can be seen as a constrained counterpart of the original MES, the following lemma reveals that the MI approximation of CMES (10) can be a negative value when  $C > 5$ .

**Lemma 4.1.** *When  $C > 5$ , for every  $\mathbf{x} \in \mathcal{X}$ , there exist thresholds  $\{z_c\}_{c=1}^C$  that result in  $\Pr(\alpha^{\text{CMES}}(\mathbf{x}) < 0) > 0$ .*

See Appendix E.3 for an interpretation and the proof. In contrast, for CMES-IBO, we have the following remark:

**Remarks 4.1.** *Our acquisition function is bounded from below by an average of probability of improvement (PI)*

from  $\tilde{f}_*$ , i.e.,  $\alpha^{\text{IBO}}(\mathbf{x}) \geq \sum_{\tilde{f}_* \in \mathcal{F}_*} \Pr(\mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})/K \geq 0$ , from which we also see non-negativity.

This remark is immediately derived by applying the well-known inequality  $\log x \leq x - 1$  to (6). Note that  $\Pr(\mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  is PI from  $\tilde{f}_*$  because it is the probability that  $f(\mathbf{x}) \geq \tilde{f}_*$  and  $g_c(\mathbf{x}) \geq z_c$  for  $\forall c$ . The unconstrained MES also has shown a relation with PI (Wang & Jegelka, 2017), while for CMES, it is difficult to see an interpretable relation with PI. Let the worst-case error be the largest error of the MI approximation among all the possible optimal-value samples  $\mathcal{F}_*$ . For problems in which  $\Pr(\alpha^{\text{CMES}}(\mathbf{x}) < 0) > 0$ , the worst-case error of CMES-IBO is smaller than that of CMES. Another important implication of the remark is that although CMES-IBO considers a lower bound,  $\alpha^{\text{IBO}}(\mathbf{x})$  should be substantially larger than 0 if  $\mathbf{x}$  is promising in the sense of the PI, which cannot be guaranteed for CMES as shown in the following toy example.

Figure 2 shows a toy example in which all constraints  $g_1, \dots, g_C$  are the same function just for simplicity (note that the independence assumption on GP models does not mean that observations of different functions cannot take the same value) and change  $C$  from 4 to 7. CMES and CMES-IBO are calculated with  $K = 10$ . In the figure, we also provide a numerical approximation of the original MI naively estimated by using the kernel density estimation (KDE), denoted as KDE-MI. For KDE-MI, since we used a large number of samples 10000 (which cannot be performed in practice), we regard KDE-MI as a pseudo ground truth. To demonstrate negativity of CMES, the threshold  $z_c$  (shown in Fig. 2 (b)) is set so that  $R_{\tilde{f}_*}$  in CMES (10) becomes large negative value around  $x = 0.9$  particularly when  $C$  is large. Then, in fact, CMES takes negative values around the global optimum of KDE-MI in the case of  $C = 6$  and 7. In contrast, CMES-IBO keeps the same global optimum as KDE-MI for all  $C$ . Although CMES has a negative value only when  $C > 5$  as Lemma 4.1 indicates, the acquisition values gradually decrease with the growth of  $C$ . Although values of CMES-IBO are nearly half of KDE-MI, we can clearly see they are highly correlated. Further detail of this example is in Appendix E.4.

Note that, in practice, the problematic behavior of CMES shown in Fig. 2 (c) does not occur frequently depending on a given problem. Although we will see that CMES often reasonably works in later experiments, the potential risk of this negativity issue degrades the reliability of CMES.

### 4.2. Bounds for Estimation Variability

We here analyze variability of our MI estimator (6), that is, how the MI estimator (6) deviates from its expectation (5). The simplicity of our MI estimator (6) enables us to derive the following theorem, while it is difficult to derive a similar guarantee for CMES:




Figure 2: In (a) and (b), the GPs for  $f$  and  $g_c$  are shown, respectively. In (c)-(e), CMES, CMES-IBO, and the KDE-based MI approximation are shown, and each vertical dotted line indicates the maximum.

**Theorem 4.1.** For every  $\mathbf{x} \in \mathcal{X}$ , our MI estimator (6) satisfies  $\nabla_{f_*} [-\log \bar{Z}_{\mathbf{x}}(f_*)] \leq 2$ , where  $\nabla_{f_*}$  represents the variance with respect to  $f_*$ . Moreover, for every  $\xi \geq 0$  and  $\mathbf{x} \in \mathcal{X}$ , following concentration inequality holds

$$\Pr\{|\alpha^{\text{IBO}}(\mathbf{x}) - L(\mathbf{x})| \geq \xi\} \leq \min\{U_1, U_2\}, \quad (11)$$

where  $U_1 = \frac{2}{K\xi^2}$  and  $U_2 = 2 \exp\left[-AK \min\left(\frac{\xi^2}{B^2}, \frac{\xi}{B}\right)\right]$  with constants  $A$  and  $B$ .

See Appendix C for the proof.

The variance of our MI estimator with  $K = 1$  can be bounded by the constant 2, and this fact provides the bound by  $U_1$  as a direct consequence of Chebyshev’s inequality. Most importantly,  $U_2$  yields exponentially fast decay on  $K$  and  $\xi$  in contrast to the linear and quadratic dependence by  $U_1$ . Although the convergence rate of  $U_1$  is worse than  $U_2$ ,  $U_1$  may provide a tighter bound for specific values of  $K$  and  $\xi$ . Note that our bound (11) holds for any MC samples  $K$ , i.e., it is not an asymptotic result such as the central limit theorem. Moreover,  $U_1$  and  $U_2$  do not depend on any parameter of GPs and the problem setting, such as  $\sigma^{(f)^2}(\mathbf{x})$ ,  $\sigma^{(g_c)^2}(\mathbf{x})$ , and  $C$ . Our analysis can be seen as a guarantee of the estimation robustness of CMES-IBO even with a small  $K$ .

## 5. Related Work

Extensions of unconstrained BO methods to CBO have been widely studied. The standard expected improvement (EI) based approaches are called *EI with Constraint* (EIC) (Schonlau et al., 1998; Snoek, 2013; Gelbart et al., 2014; Gardner et al., 2014; Zhang et al., 2021). However, when no feasible solution is obtained, the so-called ‘current best solution’ can not be defined. To avoid this difficulty, heuris-

tic strategies, such as only using the feasible probability (Gelbart et al., 2014) and introducing a threshold hyperparameter (Letham et al., 2019), have been considered, but an appropriate remedy for this issue is still an open problem. Gramacy & Lee (2010) and Picheny (2014) consider the expected reduction of EI and PI, respectively. However, since these approaches require expensive numerical integrations defined on the entire domain  $\mathcal{X}$ , their applicability is limited. A *Thompson Sampling* (TS) based method called *Scalable Constrained BO* (SCBO) (Eriksson & Poloczek, 2021) has been proposed, recently. Since a next query in TS is directly determined by one sampled  $\mathbf{x}_*$ , the decision often has high variability. Further, when sampled  $\mathcal{X}_{\text{feasible}}$  is empty, SCBO selects the input that has the smallest sum of the constraint violations, but the rationale of this selection in a sense of TS is not clarified.

The augmented Lagrangian (AL), which is a classical approach to white-box constrained problems, has also been combined with BO (Gramacy et al., 2016; Picheny et al., 2016). In AL-based methods, initial parameters of a Lagrange multiplier and a penalty coefficient should be specified. In the classical white-box optimization, these parameters can be refined during a large number of iterations. However, for expensive black-box problems, the number of possible iterations is much smaller than the classical optimization, and the effect of the initial parameters can be significant. Ariaifar et al. (2019) proposed *Alternating Direction Method of Multipliers BO* (ADMMBO) based on the famous ADMM algorithm. ADMMBO is only for decoupled setting, in which each one of objective and constraint functions is observed separately. In this paper, we only focus on the setting in which objective and constraint functions are simultaneously observed.

Entropy-based approaches are also studied in CBO literature. Hernández-Lobato et al. (2015; 2016) proposed *Predictive Entropy Search with Constraint* (PESC), which is an extension of a prior work on unconstrained problems (Hernández-Lobato et al., 2014). PESC considers the information gain of the optimum  $\mathbf{x}_*$ , but complicated approximations, for which any relation with the original MI has not been clarified, are required. Another well-known entropy-based BO is MES, for which we already discussed in Section 4.1. The idea of considering  $f_*$  instead of  $\mathbf{x}_*$  was first shown by Hoffman & Ghahramani (2015), and Wang & Jegelka (2017) proposed the current standard form of MES by introducing a simple approximation of the entropy. MES for multi-objective problems with constraints has been recently proposed by two papers (Belakaria et al., 2020; Fernández-Sánchez et al., 2020). Fernández-Sánchez et al. (2020) pointed out that the entropy evaluation of (Belakaria et al., 2020) is incorrect and derived a multi-objective counterpart of CMES (9), but they introduced multiple additional approximations by which the relation with the origi-

nal MI is further unclear. Moreover, it is worth noting that none of the information-theoretic CBO studies have considered the possibility of  $\mathcal{X}_{\text{feasible}} = \emptyset$ . See Appendix D for a detailed discussion.

A simple approach to incorporating multiple constraints is to transform it into a single constraint such as  $g(\mathbf{x}) := \min_c (g_c(\mathbf{x}) - z_c) \geq 0$ . Both of fitting a GP directly to  $g(\mathbf{x})$  and fitting GPs to  $g_c(\mathbf{x})$  individually are possible strategies in this case, but the former discards observations of each  $g_c(\mathbf{x})$  (only use the min value) and the latter makes  $g(\mathbf{x})$  non-Gaussian. The binary setting approach (Perrone et al., 2019) can also be used for this purpose, but then, observations of each  $g_c(\mathbf{x})$  are discarded.

The parallel extension of CBO has not been widely studied. Letham et al. (2019) have shown the parallel CBO based on EIC. Although Eriksson & Poloczek (2021) have not focused on the parallelization, SCBO is applicable to the parallel setting easily because of the nature of TS. For MES, although Takeno et al. (2020; 2022) have proposed a CMI based parallel BO for the multi-fidelity setting, their approach is not based on a lower bound of CMI, and constrained problems have not been studied. Wilson et al. (2018) discussed the MC-based parallel BO in general, but they considered only ‘myopic maximal (MM) acquisition functions.’ MES is not included in the class of MM acquisition functions, and CBO was not considered in (Wilson et al., 2018).

To our knowledge, no existing studies use a lower bound of the MI for BO. Moss et al. (2021) recently proposed a lower bound based extension of MES. However, their bound is a lower bound of the ‘approximate’ MI used in the original MES, which is not a lower bound of the MI, and constrained problems have not been studied. Although we only focus on constrained problems, considering the effectiveness of the lower bound approach to the unconstrained problem is one of our important future works.

Lastly, we mention regret analysis of the entropy-based BO. Wang & Jegelka (2017) show the regret analysis of ‘one sample MES,’ in which only one optimal value is sampled in the MC estimation. To our knowledge, no other regret analysis is known for the entropy-based BO ((Belakaria et al., 2019) show an extension to the multi-objective problem, but their regret can be negative value as pointed out by (Suzuki et al., 2020)). However, the original theorem of MES contains several technical problems. For example, the theorem assumes that  $f$  follows a GP, but the maximum of  $f$  is regarded as a deterministic variable in the proof, which contradicts each other (more details including other issues are in Appendix H). Therefore, convergence guarantee is still an open problem for the entropy-based BO.

## 6. Experiments

We demonstrate the performance of sequential optimization by comparing with CMES, EIC (Gelbart et al., 2014), a TS-based method referred to as TSC, and PESC (Hernández-Lobato et al., 2015) in Spearmint (<https://github.com/HIPS/Spearmint/tree/PESC>). Note that CMES is based on our definition of  $f_*$  (1). TSC is a variant of SCBO (Eriksson & Poloczek, 2021) simplified by omitting the output transformation such as the Gaussian copula-based transformation of the objective and the trust region strategy. Since these two are general strategies applicable to any acquisition functions, to only focus on the difference of acquisition functions, we do not employ them. Although the AL-based methods (Gramacy et al., 2016; Picheny et al., 2016) are not shown as baselines, they are outperformed by PESC and SCBO in prior work (Hernández-Lobato et al., 2015; Eriksson & Poloczek, 2021), respectively.

Performances on GP-derived synthetic functions, benchmark functions, and two real-world problems were evaluated by using the *utility gap* (UG), which is also employed in various prior work (Hernández-Lobato et al., 2015; Picheny et al., 2016; Eriksson & Poloczek, 2021). We set the recommendation at iteration  $t$  as  $\hat{\mathbf{x}}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \mu^{(f)}(\mathbf{x})$ , s.t.  $\forall c, \Pr(g_c(\mathbf{x}) \geq z_c) \geq \sqrt[3]{0.95}$ . Then, UG is defined as  $f_* - f(\hat{\mathbf{x}}_t)$  if the recommendation is feasible, otherwise  $f_* - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ , which indicates that an infeasible recommendation results in the worst UG. The initial inputs are sampled by Latin hypercube sampling (Loh, 1996). The sample size of all MC approximations is set as 10. For the kernel function in GPs, we used a linear combination of the linear kernel  $k_{\text{LIN}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  and RBF kernel  $k_{\text{RBF}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$  defined as  $\sigma_{\text{LIN}}^2 k_{\text{LIN}}(\mathbf{x}, \mathbf{x}') + \sigma_{\text{RBF}}^2 k_{\text{RBF}}(\mathbf{x}, \mathbf{x}')$ , where  $\sigma_{\text{LIN}}^2$  and  $\sigma_{\text{RBF}}^2$  are updated by the marginal likelihood maximization every 5 iteration, except for the GP-derived functions in which all the methods employed the same fixed kernels used to generate the true functions. We used  $k_{\text{LIN}}$  because benchmark functions often contain linear functions. For the GP-derived functions, we generated 10 constrained problems for each of which the experiment is run 10 times, and the mean and standard error of the total 100 trials are reported. For the benchmark and the real-word functions, we report the mean and the standard error of 10 random initializations. Other experimental details such as the number of initial points are shown in Appendix I.

First, we focus on the results for the sequential querying shown in Fig. 3.

**GP-derived synthetic functions (a):** The objective and the  $C = 10$  constraints are sampled from zero-mean GPs with RBF kernels whose length scales are 0.2. The thresholds are set as  $z_c = -0.75$  for  $\forall c$ , and the input do-



main is  $[0, 1]^2$ . We see that EIC and CMES-IBO outperformed other methods, and in particular, CMES is clearly worse than CMES-IBO. Additionally, we run the experiments with the ‘single constraint’ transformation of multiple constraints described in Section 5, in which a GP fits to  $\min_{c=1, \dots, 10} (g_c(\mathbf{x}) - z_c)$  as one constraint. The results shown by dashed lines indicate that this approach deteriorated performance compared with their multiple constraints counterpart.

**Benchmark functions (b)-(e):** We here show results on four benchmark functions called Gardner1 ( $C = 1$  and  $d = 2$ ), G1 ( $C = 9$  and  $d = 13$ ), G7 ( $C = 8$  and  $d = 10$ ), and G10 ( $C = 6$  and  $d = 8$ ). Gardner1 (Gardner et al., 2014) is a simple test problem constructed by sine and cosine functions (see Appendix I for the detailed definition). G1, G7, and G10 are from (Liang et al., 2006). In CBO literature, empirical evaluations are typically performed with  $C \leq 2$ , and thus, the settings  $C = 9, 8$  and  $6$  are large. Moreover, these three functions have large input dimensions  $d$ . We see that CMES-IBO has superior or comparable performance to the other methods in all the four plots. CMES also showed reasonable performance, but we observed substantial differences compared with CMES-IBO in Gardner1, G7, and G10. In G7, the average UG of CMES was larger than 10, while CMES-IBO converged to a much smaller value than 10. In G10, the differences of the average UG between CMES and CMES-IBO may seem small because of the large vertical scale, but after the iteration 30, the differences are larger than 30, and the error bars are not overlapped. For the benchmark functions, the results of PESC are also shown. Note that due to the difficulty of rewriting the Spearmint specification, some settings were not consistent with the other methods (see Appendix I for detail). However, we can see that the performance of PESC is particularly unstable for (c)-(e), which is presumably caused by the difficulty of approximating information gain of the high dimensional  $\mathbf{x}_*$ .

**Reactor network design (f):** In this problem, the product concentration from a sequence of two continuously stirred tank reactors is optimized with capital cost and physical phenomenon constraints (Manousiouthakis & Sourlas, 1992; Kumar et al., 2020). The input dimension is  $d = 6$ . This problem has one inequality constraint and four equality constraints, and we replaced each equality constraint with two inequality constraints, by which we have  $C = 9$  inequality constraints. Since the equality constraints cannot be strictly satisfied when evaluating UG, we set the tolerance error as  $10^{-3}$ . We can see that CMES-IBO also shows faster convergence than other methods in this problem.

**Hyperparameter optimization of neural network (g):** We tested a hyperparameter optimization of convolutional neural network (CNN) for a class-imbalanced setting, in

which, typically, performance for minority classes can deteriorate. We consider optimizing the average accuracy of all classes under the constraints that the recall of each one of classes is larger than 0.5. We fitted the two-layer CNN to the CIFAR10 dataset (Krizhevsky & Hinton, 2009), in which we set the class sizes of labels 0, ..., 4 as 2500, and the others are 5000. To control the imbalance, we gave a weight  $w_m$  to the loss function of class  $m = 1, \dots, 10$  defined by  $w_m = (N/(MN_m))^\rho$ , where  $M = 10$  is the number of classes,  $N$  is the size of the training data,  $N_m$  is the size of class  $m$ , and  $\rho$  is a hyperparameter. Setting  $\rho > 0$  gives larger weights to minority classes (scikit-learn<sup>2</sup> employs  $\rho = 1$ ), while  $\rho = 0$  is reduced to the uniform weighting. In addition to  $\rho \in [0, 2]$ , we also tune other hyperparameters such as learning rate (1 dim), dropout rate (1 dim), and the number of channels (2 dim). As a result, we have  $d = 5$  and  $C = M = 10$ . Other details are shown in Appendix I. CMES-IBO and EIC reduced UG rapidly at the beginning of the iterations, and CMES-IBO achieved a substantially smaller value than the other methods at the end of the iterations.

**Parallel Querying:** For the parallel setting, we parallelize EIC and TSC (see Appendix I for detail) and set the number of parallel queries  $Q = 3$ . The prefix ‘P-’ indicates a parallelized variant of each method (e.g., P-EIC and P-CMES). The results are shown in Fig. 4. Here, G1, G7, and G10 benchmarks were used, and we see that, similarly to the sequential case, CMES-IBO shows efficient convergences compared with the other methods.

Overall, EIC and TSC showed reasonable performance, but in some cases, their performances become much worse compared with the best method in each plot by being trapped at local optima. The performance of CMES was mostly stable compared with EIC and TSC. However, particularly in some large  $C$  experiments such as (a) and (g) in Fig. 3, CMES deteriorates compared to CMES-IBO. This result is consistent with our analysis that CMES can show problematic behavior when  $C$  is large.

Additional evaluations for different sample sizes  $K$ , and the correlated extension are shown in Appendix J.

## 7. Conclusion

We proposed an information-theoretic constrained Bayesian optimization method called CMES-IBO, derived from a lower bound of the MI. We showed its several desired properties, such as the interpretation of the infeasible case, non-negativity, and the bounds of the estimation variability. Moreover, we extended CMES-IBO to the

<sup>2</sup>[https://scikit-learn.org/stable/modules/generated/sklearn.utils.class\\_weight.compute\\_class\\_weight.html](https://scikit-learn.org/stable/modules/generated/sklearn.utils.class_weight.compute_class_weight.html)



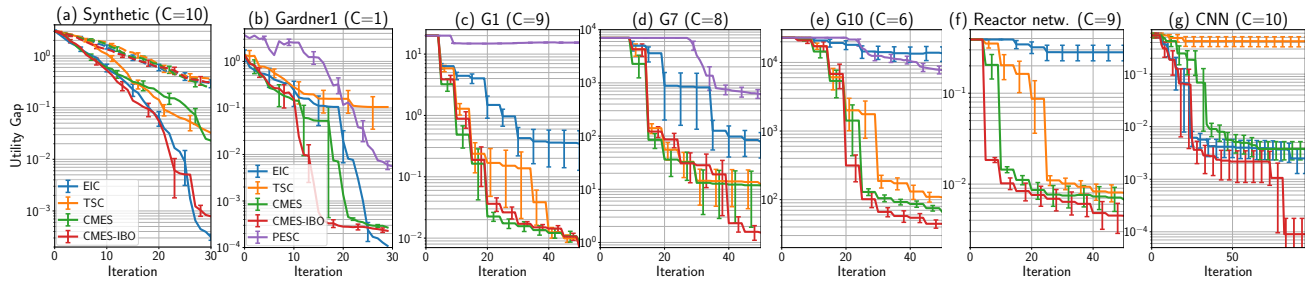


Figure 3: Utility gap in sequential querying (average and standard error). The dashed lines in the synthetic function experiment represent the ‘single constraint’ counterpart of each method that has the same color.

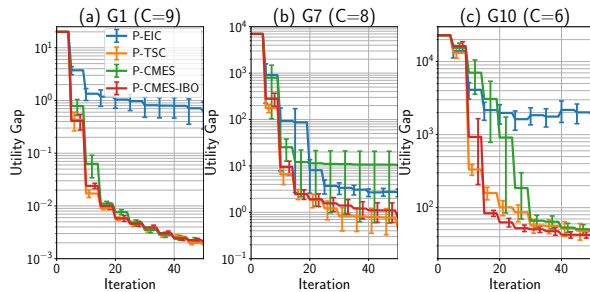


Figure 4: Utility gap in parallel querying (average and standard error).

parallel setting. The effectiveness of our proposed method was shown through various benchmark functions and real-world problems. Although we only focus on MI defined by the noiseless  $h_x$  throughout the paper, the extension for noisy observations is one of our future work.

## Acknowledgments

This work was supported by MEXT KAKENHI 18K04700, 21H03498, 22H00300, JSPS KAKENHI Grant Number JP21J14673, and RIKEN Junior Research Associate Program.

## References

- Abramowitz, M. and Stegun, I. A. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, ninth Dover printing, tenth GPO printing edition, 1964.
- Ariaifar, S., Coll-Font, J., Brooks, D., and Dy, J. AD-MMBO: Bayesian optimization with unknown constraints using ADMM. *Journal of Machine Learning Research*, 20(123):1–26, 2019.
- Beknazaryan, A., Dang, X., and Sang, H. On mutual information estimation for mixed-pair random variables. *Statistics and Probability Letters*, 148:9–16, 2019.

Belakaria, S., Deshwal, A., and Doppa, J. R. Max-value entropy search for multi-objective Bayesian optimization. In *Advances in Neural Information Processing Systems 32*, pp. 7825–7835. Curran Associates, Inc., 2019.

Belakaria, S., Deshwal, A., and Doppa, J. R. Max-value entropy search for multi-objective Bayesian optimization with constraints. *arXiv:2009.01721*, 2020.

Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.

Casella, G. and Berger, R. *Statistical Inference*. Duxbury advanced series in statistics and decision sciences. Thomson Learning, 2002.

Chinneck, J. W. *Feasibility and Infeasibility in Optimization: Algorithms and Computational Methods*. Springer Publishing Company, Incorporated, 2007.

Cover, T. M. and Thomas, J. A. *Elements of Information Theory*. Wiley-Interscience, 2006.

Eriksson, D. and Poloczek, M. Scalable constrained Bayesian optimization. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130, pp. 730–738. PMLR, 2021.

Fernández-Sánchez, D., Garrido-Merchán, E. C., and Hernández-Lobato, D. Max-value entropy search for multi-objective Bayesian optimization with constraints. *arXiv:2011.01150*, 2020.

G, M. B. and Wilhelm, S. Moments calculation for the doubly truncated multivariate normal density. *arXiv:1206.5387*, 2012.

Gardner, J., Kusner, M., Zhixiang, Weinberger, K., and Cunningham, J. Bayesian optimization with inequality constraints. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 937–945. PMLR, 2014.

Gelbart, M. A., Snoek, J., and Adams, R. P. Bayesian optimization with unknown constraints. In *Proceedings of*

- the 30th Conference on Uncertainty in Artificial Intelligence, pp. 250–259. AUAI Press, 2014.
- Genz, A. Numerical computation of multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, 1:141–150, 1992.
- GPy. GPy: A Gaussian process framework in python. <http://github.com/SheffieldML/GPy>, since 2012.
- Gramacy, R. and Lee, H. Optimization under unknown constraints. *Bayesian Statistics 9*, 2010.
- Gramacy, R. B., Gray, G. A., Digabel, S. L., Lee, H. K. H., Ranjan, P., Wells, G., and Wild, S. M. Modeling an augmented Lagrangian for blackbox constrained optimization. *Technometrics*, 58(1):1–11, 2016.
- Hernández-Lobato, J. M., Hoffman, M. W., and Ghahramani, Z. Predictive entropy search for efficient global optimization of black-box functions. In *Advances in Neural Information Processing Systems 27*, pp. 918–926. Curran Associates, Inc., 2014.
- Hernández-Lobato, J. M., Gelbart, M. A., Hoffman, M. W., Adams, R. P., and Ghahramani, Z. Predictive entropy search for Bayesian optimization with unknown constraints. In *Proceedings of the 32th International Conference on Machine Learning*, volume 37, pp. 1699–1707. PMLR, 2015.
- Hernández-Lobato, J. M., Gelbart, M. A., Adams, R. P., Hoffman, M. W., and Ghahramani, Z. A general framework for constrained Bayesian optimization using information-based search. *Journal of Machine Learning Research*, 17(160):1–53, 2016.
- Hoffman, M. W. and Ghahramani, Z. Output-space predictive entropy search for flexible global optimization. In *NIPS Workshop on Bayesian Optimization*, 2015.
- Horrace, W. C. Some results on the multivariate truncated normal distribution. *Journal of Multivariate Analysis*, 94(1):209–221, 2005.
- Johnson, S. G. The NLOpt nonlinear-optimization package. <http://github.com/stevengj/nlopt>, since 2008.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009.
- Kumar, A., Wu, G., Ali, M. Z., Mallipeddi, R., Suganthan, P. N., and Das, S. A test-suite of non-convex constrained optimization problems from the real-world and some baseline results. *Swarm and Evolutionary Computation*, 56:100693, 2020.
- Letham, B., Karrer, B., Ottoni, G., and Bakshy, E. Constrained Bayesian optimization with noisy experiments. *Bayesian Anal.*, 14(2):495–519, 2019.
- Liang, J., Runarsson, T., Mezura-Montes, E., Clerc, M., Suganthan, P., Coello, C., and Deb, K. Problem definitions and evaluation criteria for the CEC 2006 special session on constrained real-parameter optimization. *Nanyang Technological University, Tech. Rep*, 41, 2006.
- Loh, W.-L. On Latin hypercube sampling. *Ann. Statist.*, 24(5):2058–2080, 1996.
- Manousiouthakis, V. and Sourlas, D. A global optimization approach to rationally constrained rational programming. *Chemical Engineering Communications*, 115(1):127–147, 1992.
- Michalowicz, J. V., Nichols, J. M., and Bucholtz, F. *Handbook of Differential Entropy*. Chapman and Hall/CRC, 2014.
- Moss, H. B., Leslie, D. S., Gonzalez, J., and Rayson, P. GIBBON: General-purpose information-based Bayesian optimisation. *arXiv:2102.03324*, 2021.
- Nair, C., Prabhakar, B., and Shah, D. On entropy for mixtures of discrete and continuous variables. *arXiv:cs/0607075*, 2007.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019.
- Perrone, V., Shcherbatyi, I., Jenatton, R., Archambeau, C., and Seeger, M. Constrained Bayesian optimization with max-value entropy search. *arXiv:1910.07003*, 2019.
- Picheny, V. A Stepwise uncertainty reduction approach to constrained global optimization. In *Proceedings of the 17th International Conference on Artificial Intelligence and Statistics*, volume 33, pp. 787–795. PMLR, 2014.
- Picheny, V., Gramacy, R. B., Wild, S., and Le Digabel, S. Bayesian optimization under mixed constraints with a slack-variable augmented Lagrangian. In *Advances in Neural Information Processing Systems 29*, pp. 1435–1443. Curran Associates, Inc., 2016.
- Rahimi, A. and Recht, B. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems 20*, pp. 1177–1184. Curran Associates, Inc., 2008.

- Rasmussen, C. E. and Williams, C. K. I. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- Schonlau, M., Welch, W. J., and Jones, D. R. *Global versus local search in constrained optimization of computer models*, volume 34 of *Lecture Notes–Monograph Series*, pp. 11–25. Institute of Mathematical Statistics, 1998.
- Snoek, J. *Bayesian Optimization and Semiparametric Models with Applications to Assistive Technology*. University of Toronto, 2013.
- Suzuki, S., Takeno, S., Tamura, T., Shitara, K., and Karasuyama, M. Multi-objective Bayesian optimization using Pareto-frontier entropy. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 9279–9288. PMLR, 2020.
- Svanberg, K. A class of globally convergent optimization methods based on conservative convex separable approximations. *SIAM J. on Optimization*, 12(2):555–573, 2002.
- Takeno, S., Fukuoka, H., Tsukada, Y., Koyama, T., Shiga, M., Takeuchi, I., and Karasuyama, M. Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 9334–9345. PMLR, 2020.
- Takeno, S., Fukuoka, H., Tsukada, Y., Koyama, T., Shiga, M., Takeuchi, I., and Karasuyama, M. A generalized framework of multi-fidelity max-value entropy search through joint entropy. *Neural Computation*, 2022. To appear.
- Vershynin, R. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
- Wang, Z. and Jegelka, S. Max-value entropy search for efficient Bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 3627–3635. PMLR, 2017.
- Wilson, J., Hutter, F., and Deisenroth, M. Maximizing acquisition functions for Bayesian optimization. In *Advances in Neural Information Processing Systems 31*, pp. 9906–9917. Curran Associates, Inc., 2018.
- Zhang, Y., Zhang, X., and Frazier, P. Constrained two-step look-ahead Bayesian optimization. In *Advances in Neural Information Processing Systems 34*, 2021.

## A. Derivation of Information Lower Bound

To derive our lower bound, we start from the following representation of the MI:

$$I(\mathbf{h}_x; f_*) = \mathbb{E}_{f_*} \left[ \int p(\mathbf{h}_x | f_*) \log \frac{p(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} d\mathbf{h}_x \right]. \quad (12)$$

Note that since  $p(f_*)$  can be seen as a mixture distribution in which one of the components (the green distribution in Fig. 1) consists of a single constant value ( $-\infty$ ), the expectation regarding  $f_*$  for a given function  $b : \mathbb{R} \cup \{-\infty\} \mapsto \mathbb{R}$  can be defined as

$$\mathbb{E}_{f_*} [b(f_*)] = \Pr(f_* = -\infty)b(-\infty) + \Pr(f_* \neq -\infty) \int_{\mathbb{R}} p(f_* | f_* \neq -\infty)b(f_*)df_*. \quad (13)$$

The MI (12) defined through this expectation can be seen as a special case of an MI for a mixture distribution shown in (Nair et al., 2007; Beknazaryan et al., 2019).

Then, our information lower bound can be derived as follows:

$$\begin{aligned} I(\mathbf{h}_x; f_*) &= \mathbb{E}_{f_*} \left[ \int p(\mathbf{h}_x | f_*) \log \frac{p(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} d\mathbf{h}_x \right] \\ &= \mathbb{E}_{f_*} \left[ \int p(\mathbf{h}_x | f_*) \left( \log \frac{q(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} + \log \frac{p(\mathbf{h}_x | f_*)}{q(\mathbf{h}_x | f_*)} \right) d\mathbf{h}_x \right] \\ &= \mathbb{E}_{f_*} \left[ \int p(\mathbf{h}_x | f_*) \log \frac{q(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} d\mathbf{h}_x + D_{\text{KL}}(p(\mathbf{h}_x | f_*) || q(\mathbf{h}_x | f_*)) \right] \\ &\geq \mathbb{E}_{f_*} \left[ \int p(\mathbf{h}_x | f_*) \log \frac{q(\mathbf{h}_x | f_*)}{p(\mathbf{h}_x)} d\mathbf{h}_x \right] \\ &= \mathbb{E}_{f_*} \left[ \int p(\mathbf{h}_x | f_*) \log \frac{p(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*})}{p(\mathbf{h}_x)} d\mathbf{h}_x \right]. \end{aligned} \quad (14)$$

In the last line, we replace  $q(\mathbf{h}_x | f_*)$  with  $p(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*})$  as defined in (4). Here, from the definition of  $f_*$ , we see  $p(\mathbf{h}_x | f_*) = p(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*}) = 0$  for  $\mathbf{h}_x \in \mathcal{A}_{f_*}$ . Additionally, in information theory,  $0 \log 0$  is treated as zero based on that  $\lim_{x \rightarrow 0} x \log x = 0$  (Cover & Thomas, 2006). Thus, we obtain

$$\begin{aligned} (14) &= \mathbb{E}_{f_*} \left[ \int_{\bar{\mathcal{A}}_{f_*}} p(\mathbf{h}_x | f_*) \log \frac{p(\mathbf{h}_x)}{\bar{Z}_x(f_*)p(\mathbf{h}_x)} d\mathbf{h}_x \right] \\ &= -\mathbb{E}_{f_*} \left[ \int_{\bar{\mathcal{A}}_{f_*}} p(\mathbf{h}_x | f_*) \log \bar{Z}_x(f_*) d\mathbf{h}_x \right] \\ &= -\mathbb{E}_{f_*} \left[ \log \bar{Z}_x(f_*) \int_{\bar{\mathcal{A}}_{f_*}} p(\mathbf{h}_x | f_*) d\mathbf{h}_x \right] \\ &= -\mathbb{E}_{f_*} [\log \bar{Z}_x(f_*)]. \end{aligned}$$

Note that although the expectation  $\mathbb{E}_{f_*}$  is written as the mixture (13), the MC approximation can be directly performed because a sample from the mixture distribution can be obtained through the procedure described in Section G.1, by which our acquisition function (6) is derived.

## B. Details of Parallelization

In this section, we provide the detailed derivation for the parallel setting omitted in the main paper. We assume that  $Q$  workers are available by which we can select  $Q$  queries as a batch at every iteration. Let  $\mathcal{X}_q := \{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(q)}\}$  be a set of input queries and  $\mathcal{H}_q := \{\mathbf{h}_{\mathbf{x}^{(1)}}, \dots, \mathbf{h}_{\mathbf{x}^{(q)}}\}$  be a set of output vectors for an integer  $1 \leq q \leq Q$ . A naïve extension is to maximize MI between  $\mathcal{H}_Q$  and  $f_*$ , i.e.,  $\text{MI}(\mathcal{H}_Q; f_*)$ , with respect to the queries  $\mathcal{X}_Q$ . However, this approach results in a  $Qd$  dimensional acquisition function maximization, which is often extremely hard. Instead, we propose a greedy approximation of  $\max_{\mathcal{X}_Q} \text{MI}(\mathcal{H}_Q; f_*)$  using the *conditional mutual information* (CMI).



The  $q$ -th step of our greedy selection is defined as  $\arg \max_{\mathbf{x}^{(q)}} \text{MI}(\mathcal{H}_q; f_*)$ , where  $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(q-1)}$  are already fixed by the previous steps. This procedure maximizes the additional information gain produced by  $\mathbf{h}_{\mathbf{x}^{(q)}}$ , which can be seen through the following expansion:

$$\text{MI}(\mathcal{H}_q; f_*) = \text{MI}(\mathcal{H}_{q-1}; f_*) + \text{CMI}(\mathbf{h}_{\mathbf{x}^{(q)}}; f_* | \mathcal{H}_{q-1}),$$

where the second term is  $\text{CMI}(\mathbf{h}_{\mathbf{x}^{(q)}}; f_* | \mathcal{H}_{q-1}) := \mathbb{E}_{\mathcal{H}_{q-1}} [\text{MI}(\mathbf{h}_{\mathbf{x}^{(q)}}; f_* | \mathcal{H}_{q-1})]$ . Importantly, the first term does not depend on  $\mathbf{x}^{(q)}$  anymore. Therefore, we obtain  $\arg \max_{\mathbf{x}^{(q)}} \text{MI}(\mathcal{H}_q; f_*) = \arg \max_{\mathbf{x}^{(q)}} \text{CMI}(\mathbf{h}_{\mathbf{x}^{(q)}}; f_* | \mathcal{H}_{q-1})$ , which means that the greedy selection can be performed by the maximization of CMI.

Our lower bound approach used in (5) can be applied to CMI as follows:

$$\text{CMI}(\mathbf{h}_{\mathbf{x}^{(q)}}; f_* | \mathcal{H}_{q-1}) \geq -\mathbb{E}_{\mathcal{H}_{q-1}, f_*} [\log(1 - Z_{\mathbf{x}^{(q)}}(f_* | \mathcal{X}_{q-1}))] =: L_{\text{par}}(\mathbf{x}^{(q)}), \quad (7)$$

where  $Z_{\mathbf{x}^{(q)}}(f_* | \mathcal{X}_{q-1}) := \Pr(\mathbf{h}_{\mathbf{x}^{(q)}} \in \mathcal{A}_{f_*} | \mathcal{X}_{q-1}, \mathcal{H}_{q-1})$ . The derivation of this lower bound is omitted because it is almost identical with the sequential case (see Appendix A). Suppose that  $m_{\mathbf{x}^{(q)}}^{(f)}$  and  $s_{\mathbf{x}^{(q)}}^{(f)2}$  are the predictive mean and variance of  $f(\mathbf{x})$  after conditioning by  $\mathcal{H}_{q-1}$ , respectively, and that  $m_{\mathbf{x}^{(q)}}^{(g_c)}$  and  $s_{\mathbf{x}^{(q)}}^{(g_c)2}$  are those for  $g_c(\mathbf{x})$ . Then, from the independence assumption,  $Z_{\mathbf{x}^{(q)}}(f_* | \mathcal{X}_{q-1}) = (1 - \Phi(\eta_{\mathbf{x}^{(q)}}^{(f)}(f_*))) \prod_{c=1}^C (1 - \Phi(\eta_{\mathbf{x}^{(q)}}^{(g_c)}(z_c)))$ , where  $\eta_{\mathbf{x}^{(q)}}^{(f)}(f_*) = (f_* - m_{\mathbf{x}^{(q)}}^{(f)})/s_{\mathbf{x}^{(q)}}^{(f)}$  and  $\eta_{\mathbf{x}^{(q)}}^{(g_c)}(z_c) = (z_c - m_{\mathbf{x}^{(q)}}^{(g_c)})/s_{\mathbf{x}^{(q)}}^{(g_c)}$ ,  $c = 1, \dots, C$ . By applying the Monte Carlo estimation to the expectation in (7), we obtain the acquisition function for the  $q$ -th query as

$$\alpha^{\text{IBO}}(\mathbf{x}^{(q)} | \mathcal{X}_{q-1}) = -\frac{1}{K} \sum_{(\tilde{f}_*, \mathcal{H}_{q-1}) \in \mathcal{J}} \log(1 - Z_{\mathbf{x}^{(q)}}(\tilde{f}_* | \mathcal{X}_{q-1})), \quad (8)$$

where  $\mathcal{J}$  is a set of  $K$  sampled  $(\tilde{f}_*, \mathcal{H}_{q-1})$  from the current GPs.

### C. Proof of Theorem 4.1

Let  $f_{*(1)}, \dots, f_{*(K)}$  be i.i.d. random variables sampled from the distribution of  $f_*$ , and let  $D_k := -\log(1 - \Pr(\mathbf{h}_{\mathbf{x}} \in \mathcal{A}_{f_{*(k)}}))$ , where  $\mathcal{A}_{f_{*(k)}} := (f_{*(k)}, \infty) \times (z_1, \infty) \times \dots \times (z_C, \infty)$  for  $\forall k \in \{1, \dots, K\}$ . Then, our acquisition function can be written as

$$\alpha^{\text{IBO}}(\mathbf{x}) = \frac{1}{K} \sum_{k=1}^K D_k.$$

For any  $b \in \mathbb{R} \cup \{-\infty\}$ , the following inequality holds:

$$\Pr(\mathbf{h}_{\mathbf{x}} \in \mathcal{B}) \leq \Pr(f_* > b),$$

where  $\mathcal{B} := (b, \infty) \times (z_1, \infty) \times \dots \times (z_C, \infty)$ . This is because  $\mathbf{h}_{\mathbf{x}} \in \mathcal{B}$  implies that (i)  $\mathcal{X}_{\text{feasible}}$  is not empty since at least  $\mathbf{x}$  is feasible and (ii)  $f_* = \max_{\mathbf{x}' \in \mathcal{X}_{\text{feasible}}} f(\mathbf{x}') \geq f(\mathbf{x}) > b$ . By substituting the sampled  $f_{*(1)}, \dots, f_{*(K)}$  into  $b$ , we see

$$\Pr(\mathbf{h}_{\mathbf{x}} \in \mathcal{A}_{f_{*(k)}}) \leq \Pr(f_* > f_{*(k)}),$$

and

$$1 - \Pr(\mathbf{h}_{\mathbf{x}} \in \mathcal{A}_{f_{*(k)}}) \geq 1 - \Pr(f_* > f_{*(k)}) = F_*(f_{*(k)}), \quad (15)$$

where  $F_*(b) := \Pr(f_* \leq b)$  is CDF of  $f_*$ .

Let us consider the distribution of  $F_*(f_{*(k)})$  induced by  $f_{*(k)}$ . Define  $\tau := \Pr(f_* = -\infty)$ . Then,  $f_{*(k)} = -\infty$  with probability  $\tau$  because  $f_{*(k)}$  is a random sample of  $f_*$ . This implies  $F_*(f_{*(k)}) = \Pr(f_* = -\infty) = \tau$  with probability  $\tau$ . Moreover, the CDF of  $F_*(f_{*(k)})$ , defined as  $\Pr(F_*(f_{*(k)}) \leq \alpha)$ , is continuous for  $\alpha \in (\tau, 1)$ . Thus, as with the probability integral transform (e.g., see (Casella & Berger, 2002)),  $\Pr(F_*(f_{*(k)}) \leq \alpha)$  can be obtained as

$$\begin{aligned} \Pr(F_*(f_{*(k)}) \leq \alpha) &= \Pr(F_*^{-1}(F_*(f_{*(k)})) \leq F_*^{-1}(\alpha)) \\ &= \Pr(f_{*(k)} \leq F_*^{-1}(\alpha)) \\ &= F_*(F_*^{-1}(\alpha)) \\ &= \alpha, \end{aligned}$$

where  $F_*^{-1}(\cdot)$  is the quantile function of  $f_*$ . Consequently, CDF of  $F_*(f_{*(k)})$  is obtained as,

$$\Pr(F_*(f_{*(k)}) \leq \alpha) = \begin{cases} 0 & \text{if } \alpha < \tau, \\ \alpha & \text{if } \alpha \geq \tau. \end{cases}$$

From above,  $F_*(f_{*(k)}) = \tau$  with probability  $\tau$ , and  $F_*(f_{*(k)}) \sim \text{Unif}(\tau, 1)$  with probability  $1 - \tau$ . Then, we define the random variable  $U_k$  as follows:

$$U_k = \begin{cases} \tilde{U}_k & \text{if } f_{*(k)} = -\infty, \\ F_*(f_{*(k)}) & \text{if } f_{*(k)} \neq -\infty, \end{cases}$$

where  $\tilde{U}_k \sim \text{Unif}(0, \tau)$ . In this definition, we have  $U_k \sim \text{Unif}(0, \tau)$  with probability  $\tau$ , and  $U_k \sim \text{Unif}(\tau, 1)$  with probability  $1 - \tau$ . Therefore,  $U_k \sim \text{Unif}(0, 1)$ . Since  $F_*(f_{*(k)}) \geq U_k$  from the definition, by combining (15), we obtain

$$1 - \Pr(\mathbf{h}_x \in \mathcal{A}_{f_{*(k)}}) \geq F_*(f_{*(k)}) \geq U_k. \quad (16)$$

From this inequality and the monotonicity of logarithm, we can transform

$$D_k = -\log(1 - \Pr(\mathbf{h}_x \in \mathcal{A}_{f_{*(k)}})) \leq -\log(F_*(f_{*(k)})) \leq -\log(U_k) =: M_k.$$

$D_k$  and  $M_k$  are nonnegative random variables, and  $\Pr(D_k \leq M_k) = 1$ . Hence, we see that  $0 \leq \mathbb{E}_{\mathcal{F}_*}[D_k] \leq \mathbb{E}_{\mathcal{F}_*}[M_k]$ ,  $\mathbb{E}_{\mathcal{F}_*}[D_k^2] \leq \mathbb{E}_{\mathcal{F}_*}[M_k^2]$ , and  $\Pr(D_k \geq b) \leq \Pr(M_k \geq b)$  for  $\forall b \in \mathbb{R}$ .

We will show that variance and concentration bounds for  $D_k$  can be derived from that of  $M_k$ . Using the inverse probability integral transform (Casella & Berger, 2002),  $-\log(U_k)$  follows an exponential distribution with the rate parameter  $\lambda = 1$ :

$$M_k = -\log(U_k) \sim \exp(\lambda = 1).$$

Therefore,  $\mathbb{E}_{f_{*(k)}}[M_k^2] = 2$ , and we can derive the inequality below:

$$\begin{aligned} \mathbb{V}_{f_{*(k)}}[D_k] &= \mathbb{E}_{f_{*(k)}}[D_k^2] - \mathbb{E}_{f_{*(k)}}[D_k]^2 \\ &\leq \mathbb{E}_{f_{*(k)}}[D_k^2] \\ &\leq \mathbb{E}_{f_{*(k)}}[M_k^2] \\ &= 2. \end{aligned}$$

Moreover, we see that  $D_k$  satisfies the condition being a *sub-exponential random variable* (Proposition 2.7.1 (a) in Vershynin, 2018), that is,

$$\Pr(|D_k| \geq \xi) = \Pr(D_k \geq \xi) \leq \Pr(M_k \geq \xi) = \exp(-\xi) \text{ for all } \xi \geq 0.$$

Consequently,  $D_k$  is a sub-exponential random variable, whose variance is bounded from above  $\mathbb{V}_{f_{*(k)}}[D_k] \leq 2$ .

We can directly apply Chebyshev's and Bernstein's inequalities from the bound of variance and the fact that  $D_k$  is a sub-exponential random variable. First, applying Chebyshev's inequality, we can see immediately

$$\Pr(|\alpha^{\text{IBO}}(\mathbf{x}) - L(\mathbf{x})| \geq \xi) \leq \frac{2}{K\xi^2}, \quad (17)$$

for any  $\xi \geq 0$ . Second, applying Bernstein's inequality (Corollary 2.8.3 in Vershynin, 2018), we can derive

$$\Pr(|\alpha^{\text{IBO}}(\mathbf{x}) - L(\mathbf{x})| \geq \xi) \leq 2 \exp\left[-AK \min\left(\frac{\xi^2}{B^2}, \frac{\xi}{B}\right)\right], \quad (18)$$

where  $A$  is a constant and  $B$  can be provided as a *sub-exponential norm*<sup>3</sup> (Definition 2.7.5 in Vershynin, 2018) of  $D_k - \mathbb{E}_{f_{*(k)}}[D_k]$ . For all  $t \geq 0$  and  $\mathbf{x} \in \mathcal{X}$ , the sub-exponential norm  $\|D_k - \mathbb{E}_{f_{*(k)}}[D_k]\|_{\psi_1}$  can be bounded as

$$\begin{aligned} \|D_k - \mathbb{E}_{f_{*(k)}}[D_k]\|_{\psi_1} &\leq \|D_k\|_{\psi_1} + \|\mathbb{E}_{f_{*(k)}}[D_k]\|_{\psi_1} && \text{(triangle inequality)} \\ &\leq \|M_k\|_{\psi_1} + \|\mathbb{E}_{f_{*(k)}}[M_k]\|_{\psi_1} && (\Pr(D_k \leq M_k) = 1) \\ &= \|M_k\|_{\psi_1} + \|1\|_{\psi_1} && (\mathbb{E}_{f_{*(k)}}[M_k] = 1) \\ &= 2 + 1/\log 2. && (M_k \sim \exp(\lambda = 1)) \end{aligned}$$

<sup>3</sup>The sub-exponential norm for a random variable  $X$  is defined as follows:  $\|X\|_{\psi_1} := \inf\{\xi > 0 \mid \mathbb{E}[\exp(|X|/\xi)] \leq 2\}$ .

Therefore,  $B$  can be set as  $2 + 1/\log 2$ . Consequently, we can derive Theorem 4.1 combining these two inequalities (17) and (18).

The above derivations can be directly extended to the parallel setting. Let  $D_k^{(q)} := -\log(1 - \Pr(\mathbf{h}_x \in \mathcal{A}_{f_*(k)} | \mathcal{H}_{q-1}))$  and  $\alpha^{\text{IBO}}(\mathbf{x} | \mathcal{X}_{q-1}) = 1/K \sum_{k=1}^K D_k^{(q)}$ . All the inequalities shown in the sequential case hold for an arbitrary dataset, which can be the observed dataset  $\mathcal{D}_t$  augmented by  $\mathcal{H}_q$ . Thus, for an arbitrary fixed  $\mathcal{H}_{q(k)}$ , we obtain  $\mathbb{E}_{f_*(k) | \mathcal{H}_{q(k)}} [D_k^{(q)^2}] \leq 2$  and  $\Pr(|D_k^{(q)}| > \xi | \mathcal{H}_{q(k)}) \leq \exp(-\xi)$  for all  $\xi \geq 0$ , where the expectation and probability are taken by  $f_*(k) | \mathcal{H}_{q(k)}, \mathcal{D}_t$ . Therefore, we obtain the bounds of the variance and the tail probability as follows:

$$\begin{aligned} \mathbb{V}_{f_*(k), \mathcal{H}_{q(k)}} [D_k^{(q)}] &= \mathbb{E}_{f_*(k), \mathcal{H}_{q(k)}} [D_k^{(q)^2}] - \mathbb{E}_{f_*(k), \mathcal{H}_{q(k)}} [D_k^{(q)}]^2 \leq \mathbb{E}_{\mathcal{H}_{q(k)}} [\mathbb{E}_{f_*(k) | \mathcal{H}_{q(k)}} [D_k^{(q)^2}]] \leq \mathbb{E}_{\mathcal{H}_{q(k)}} [2] = 2, \\ \Pr(|D_k^{(q)}| > \xi) &= \mathbb{E}_{\mathcal{H}_{q(k)}} [\Pr(|D_k^{(q)}| > \xi | \mathcal{H}_{q(k)})] \leq \mathbb{E}_{\mathcal{H}_{q(k)}} [\exp(-\xi)] = \exp(-\xi) \text{ for all } \xi \geq 0. \end{aligned}$$

Consequently,  $D_k^{(q)}$  is a sub-exponential random variable, whose variance is bounded from above by 2. Hence, we can apply Chebyshev's and Bernstein's inequalities as with the sequential case.

## D. Discussion on Infeasibility

Here, we discuss the case that sample paths generated for sampling  $f_*$  do not have any feasible solution. Perrone et al. (2019) defined  $f_*$  as  $\max f(\mathbf{x})$ , s.t.  $g_c(\mathbf{x}) \geq z_c$  for  $\forall c = 1, \dots, C$ , by which the infeasible case is not considered. However, particularly when the dataset  $\mathcal{D}_{t-1}$  does not contain any feasible solution (which typically occurs at the beginning of the optimization), it often occurs that a sample path generated for sampling  $f_*$  does not have any feasible solution. An illustrative example is shown in Figure 1 (a), in which the feasible region of the green sample path is empty. In fact, GPs can generate an infeasible sample path, i.e.,  $\Pr(\mathcal{X}_{\text{feasible}} = \emptyset) > 0$  unless we observe a noiseless feasible solution. Then, there is a possibility that  $f_*$  is not defined, by which MI even cannot be theoretically defined.

Other than using our definition of  $f_*$  (1), another possible approach to avoiding this problem is to assume the existence of feasible solutions in the GPs, but this approach has at least the following three disadvantages. First, this approach ignores uncertainty about the existence of a feasible solution. Identifying the existence of a feasible solution for a given problem is a key issue in constrained problems (called the feasibility problem (Chinneck, 2007)). However, information gain about feasibility cannot be incorporated in this approach. Second, we need to generate sample paths that have feasible solutions for sampling  $f_*$ . A naïve approach is the rejection sampling, in which generated constraint functions are rejected if no feasible solution exists, but this may require a huge number of samplings. Third, the predictive distribution of  $\mathbf{h}_x$  under this condition  $p(\mathbf{h}_x | \text{at least one feasible } \mathbf{x} \text{ exists})$  is not a Gaussian distribution anymore, and is analytically intractable. This makes the entire computational procedures of both CMES-IBO and CMES much more complicated. Perrone et al. (2019) did not mention the above issues at all and used the usual predictive distributions of GPs in the entire acquisition function evaluation without any justification. Note that the same problem exists in the PES-based CBO (Hernández-Lobato et al., 2015; 2016). On the other hand, our definition of  $f_*$  (1) resolves the above issues.

We further discuss how CMES-IBO balances the effect of the objective and constraint functions when infeasible sample paths are generated. For simplicity, we first consider the case of  $K = 1$ . The acquisition function (6) can be written as  $-\log(1 - P_{\text{imp}} \times P_{\text{fea}})$ , where  $P_{\text{imp}} := \Pr(f(\mathbf{x}) \geq \tilde{f}_*)$  and  $P_{\text{fea}} := \prod_{c=1}^C \Pr(g_c(\mathbf{x}) \geq z_c)$ . In this criterion,  $P_{\text{imp}}$  and  $P_{\text{fea}}$  represent benefits for obtaining a larger value of  $f(\mathbf{x})$  and for obtaining a feasible solution, respectively. If the sample path does not have a feasible region,  $P_{\text{imp}} = \Pr(f(\mathbf{x}) \geq -\infty) = 1$ . Thus, only the probability of being feasible  $P_{\text{fea}}$  is maximized. On the other hand, when  $K > 1$ , the balance of the effect of the objective and the constraint functions are balanced through the frequency that sampled constraint functions have a feasible region among  $K$  samplings. If the frequency is low, the probability of being feasible has a dominant effect, while if the frequency is high, the effect of the objective function becomes strong. Note that sampled constraint functions can have feasible solutions even when the dataset  $\mathcal{D}_{t-1}$  does not contain any feasible solution (as illustrated in Figure 1 (a)).

## E. Constrained Extension of Conventional Max-value Entropy Search

We here extend cMES (Perrone et al., 2019) to the multiple constraints setting and the parallel setting, and further, we show an example that its approximate MI can be negative.

### E.1. Acquisition Function of CMES

Following the conventional MES and cMES (Wang & Jegelka, 2017; Perrone et al., 2019), we approximate MI as below:

$$\begin{aligned} (2) &= H(\mathbf{h}_x) - \mathbb{E}_{f_*} [H(\mathbf{h}_x | f_*)] \\ &\approx H(\mathbf{h}_x) - \mathbb{E}_{f_*} [H(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*})]. \end{aligned} \quad (9)$$

For the approximation, the conditioning on  $f_*$  in the second term is replaced with  $\mathbf{h}_x \in \bar{\mathcal{A}}_{f_*}$  following the existing approach (Perrone et al., 2019). Further, the expectation over  $p(f_*)$  in the second term of (9) is also analytically intractable. This expectation can be approximated by the MC estimation, for which we employ the same RFF-based approach described in Appendix G.1. Thus, we obtain the following acquisition function as an approximation of (9):

$$\alpha^{\text{CMES}}(\mathbf{x}) := H(\mathbf{h}_x) - \frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} H(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{\tilde{f}_*}). \quad (19)$$

The first term is easy to calculate because it is the entropy of the multivariate normal distribution. Hence, we need to calculate the entropy in the second term  $H(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{\tilde{f}_*})$ .

Although Perrone et al. (2019) derived a closed-form of the entropy of this TMN distribution specific for  $C = 1$ , for  $C > 1$ , directly evaluating it is not trivial as pointed out by (Fernández-Sánchez et al., 2020). We show that another analytical representation of this entropy actually can be derived for general  $C$  by transforming the domain of the integration as follows:

$$\begin{aligned} H(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{\tilde{f}_*}) &= \int_{\bar{\mathcal{A}}_{\tilde{f}_*}} -\frac{p(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} \log \frac{p(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} d\mathbf{h}_x \\ &= \int -\frac{p(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} \log \frac{p(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} d\mathbf{h}_x + \int_{\mathcal{A}_{\tilde{f}_*}} \frac{p(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} \log \frac{p(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} d\mathbf{h}_x \\ &= \frac{1}{\bar{Z}_x(\tilde{f}_*)} \int -p(\mathbf{h}_x) \log p(\mathbf{h}_x) d\mathbf{h}_x + \frac{1}{\bar{Z}_x(\tilde{f}_*)} \int_{\mathcal{A}_{\tilde{f}_*}} p(\mathbf{h}_x) \log p(\mathbf{h}_x) d\mathbf{h}_x + \log(\bar{Z}_x(\tilde{f}_*)) \\ &= \frac{1}{\bar{Z}_x(\tilde{f}_*)} \int -p(\mathbf{h}_x) \log p(\mathbf{h}_x) d\mathbf{h}_x + \frac{Z_x(\tilde{f}_*)}{\bar{Z}_x(\tilde{f}_*)} \int_{\mathcal{A}_{\tilde{f}_*}} \frac{p(\mathbf{h}_x)}{Z_x(\tilde{f}_*)} \log \frac{p(\mathbf{h}_x)}{Z_x(\tilde{f}_*)} d\mathbf{h}_x + \log(\bar{Z}_x(\tilde{f}_*)) + \frac{Z_x(\tilde{f}_*) \log(Z_x(\tilde{f}_*))}{\bar{Z}_x(\tilde{f}_*)} \\ &= \frac{H(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} - \frac{Z_x(\tilde{f}_*)}{\bar{Z}_x(\tilde{f}_*)} \int_{\mathcal{A}_{\tilde{f}_*}} -p(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) \log p(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) d\mathbf{h}_x + \log(\bar{Z}_x(\tilde{f}_*)) + \frac{Z_x(\tilde{f}_*) \log(Z_x(\tilde{f}_*))}{\bar{Z}_x(\tilde{f}_*)} \\ &= \frac{H(\mathbf{h}_x)}{\bar{Z}_x(\tilde{f}_*)} - \frac{Z_x(\tilde{f}_*) H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})}{\bar{Z}_x(\tilde{f}_*)} + \log(\bar{Z}_x(\tilde{f}_*)) + \frac{Z_x(\tilde{f}_*) \log Z_x(\tilde{f}_*)}{\bar{Z}_x(\tilde{f}_*)}, \end{aligned} \quad (20)$$

where TMN  $p(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  is defined as

$$p(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) = \begin{cases} p(\mathbf{h}_x) / Z_x(\tilde{f}_*) & \text{if } \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*} \\ 0 & \text{otherwise} \end{cases}.$$

In the last equation,  $Z_x(\tilde{f}_*)$ ,  $\bar{Z}_x(\tilde{f}_*)$  and  $H(\mathbf{h}_x)$  are easy to compute. Although  $H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  in the second term is still the entropy of TMN, we show the calculation of this entropy in Appendix F for both independent and correlated cases.

Consequently, in the independent case, by substituting (20) into (19) and replacing  $H(\mathbf{h}_x)$  and  $H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  with their closed-forms (shown in Appendix F), we obtain

$$\alpha^{\text{CMES}}(\mathbf{x}) = \frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} \frac{Z_x(\tilde{f}_*)}{2(1 - Z_x(\tilde{f}_*))} R_{\tilde{f}_*} - \log(1 - Z_x(\tilde{f}_*)),$$

where

$$R_{\tilde{f}_*} := \frac{\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*) \phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))}{1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))} + \sum_{c=1}^C \frac{\gamma_{\mathbf{x}}^{(g_c)}(z_c) \phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}{1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}.$$



In the correlated case,  $H(\mathbf{h}_x)$  and  $H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  are also replaced with their closed-forms (shown in Appendix F), and we obtain

$$\alpha^{\text{CMES}}(\mathbf{x}) = \frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} \frac{Z_{\mathbf{x}}(\tilde{f}_*)}{2(1 - Z_{\mathbf{x}}(\tilde{f}_*))} \left( \text{Tr} \left( \Sigma^{-1}(\mathbf{x}) (\Sigma^{\text{TN}}(\mathbf{x}) + \mathbf{d}\mathbf{d}^\top) \right) - C - 1 \right) - \log(1 - Z_{\mathbf{x}}(\tilde{f}_*)), \quad (21)$$

where  $\text{Tr}(\cdot)$  is the trace of a matrix,  $\boldsymbol{\mu}(\mathbf{x})$  and  $\Sigma(\mathbf{x})$  are the expectation and covariance matrix of  $\mathbf{h}_x$ , respectively,  $\boldsymbol{\mu}^{\text{TN}}(\mathbf{x})$  and  $\Sigma^{\text{TN}}(\mathbf{x})$  are those for  $\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}$ , and  $\mathbf{d} = \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}) - \boldsymbol{\mu}(\mathbf{x})$ . See Appendix F for details.

## E.2. Parallelization of CMES

We can consider the CMI maximization in the same way as CMES-IBO. Suppose that we already select the  $q - 1$  queries, and we need to select the next  $q$ -th query. The CMI can be approximated as

$$\begin{aligned} \mathbb{E}_{\mathcal{H}_{q-1}} [I(\mathbf{h}_{\mathbf{x}^{(q)}}; f_* | \mathcal{H}_{q-1})] &= \mathbb{E}_{\mathcal{H}_{q-1}} [H(\mathbf{h}_{\mathbf{x}^{(q)}} | \mathcal{H}_{q-1})] - \mathbb{E}_{\mathcal{H}_{q-1}, f_*} [H(\mathbf{h}_{\mathbf{x}^{(q)}} | f_*, \mathcal{H}_{q-1})] \\ &\approx \mathbb{E}_{\mathcal{H}_{q-1}} [H(\mathbf{h}_{\mathbf{x}^{(q)}} | \mathcal{H}_{q-1})] - \mathbb{E}_{\mathcal{H}_{q-1}, f_*} [H(\mathbf{h}_{\mathbf{x}^{(q)}} | \mathbf{h}_{\mathbf{x}^{(q)}} \in \bar{\mathcal{A}}_{f_*}, \mathcal{H}_{q-1})] \end{aligned}$$

As described in Appendix B, we define  $f(\mathbf{x}) | \mathcal{D}_{t-1}, \mathcal{X}_{q-1}, \mathcal{H}_{q-1} \sim \mathcal{N}(m_{\mathbf{x}^{(q)}}^{(f)}, s_{\mathbf{x}^{(q)}}^{(f)2})$  and  $g_c(\mathbf{x}) | \mathcal{D}_{t-1}, \mathcal{X}_{q-1}, \mathcal{H}_{q-1} \sim \mathcal{N}(m_{\mathbf{x}^{(q)}}^{(g_c)}, s_{\mathbf{x}^{(q)}}^{(g_c)2})$   $c = 1, \dots, C$ . Importantly,  $s_{\mathbf{x}^{(q)}}^{(f)2}, s_{\mathbf{x}^{(q)}}^{(g_c)2}, \dots, s_{\mathbf{x}^{(q)}}^{(g_C)2}$  only depend on  $\mathcal{X}_{q-1}$ , not  $\mathcal{H}_{q-1}$ . Thus, given  $\mathcal{X}_{q-1}$ , the first term  $\mathbb{E}_{\mathcal{H}_{q-1}} [H(\mathbf{h}_{\mathbf{x}^{(q)}} | \mathcal{H}_{q-1})]$  can be calculated analytically. The second term is also calculated in the same manner as CMES with given  $\tilde{f}_*$  and  $\mathcal{H}_{q-1}$ :

$$H(\mathbf{h}_{\mathbf{x}^{(q)}} | \mathbf{h}_{\mathbf{x}^{(q)}} \in \bar{\mathcal{A}}_{\tilde{f}_*}, \mathcal{H}_{q-1}) = H(\mathbf{h}_{\mathbf{x}^{(q)}} | \mathcal{H}_{q-1}) - \frac{Z_{\mathbf{x}^{(q)}}(\tilde{f}_* | \mathcal{X}_{q-1})}{2(1 - Z_{\mathbf{x}^{(q)}}(\tilde{f}_* | \mathcal{X}_{q-1}))} R_{\tilde{f}_*}^{(q)} + \log(1 - Z_{\mathbf{x}^{(q)}}(\tilde{f}_* | \mathcal{X}_{q-1})),$$

where

$$R_{\tilde{f}_*}^{(q)} := \frac{\eta_{\mathbf{x}^{(q)}}^{(f)}(\tilde{f}_*) \phi(\eta_{\mathbf{x}^{(q)}}^{(f)}(\tilde{f}_*))}{1 - \Phi(\eta_{\mathbf{x}^{(q)}}^{(f)}(\tilde{f}_*))} + \sum_{c=1}^C \frac{\eta_{\mathbf{x}^{(q)}}^{(g_c)}(z_c) \phi(\eta_{\mathbf{x}^{(q)}}^{(g_c)}(z_c))}{1 - \Phi(\eta_{\mathbf{x}^{(q)}}^{(g_c)}(z_c))}.$$

Finally, by applying the MC approximation, we obtain the acquisition function below

$$\alpha^{\text{CMES}}(\mathbf{x}^{(q)} | \mathcal{X}_{q-1}) = \frac{1}{|\mathcal{J}|} \sum_{(\tilde{f}_*, \mathcal{H}_{q-1}) \in \mathcal{J}} \frac{Z_{\mathbf{x}^{(q)}}(\tilde{f}_* | \mathcal{X}_{q-1})}{2(1 - Z_{\mathbf{x}^{(q)}}(\tilde{f}_* | \mathcal{X}_{q-1}))} R_{\tilde{f}_*}^{(q)} - \log(1 - Z_{\mathbf{x}^{(q)}}(\tilde{f}_* | \mathcal{X}_{q-1})). \quad (22)$$

We can also derive the acquisition function in the correlated setting, but we omit the detailed derivation.

## E.3. Interpretation and Proof of Lemma 4.1

**Interpretation of negativity** At first, we provide an intuition that the acquisition function of CMES becomes negative when  $C > 5$ . In (9), CMES replaces  $\mathbb{E}_{f_*} [H(\mathbf{h}_x | f_*)]$  with  $\mathbb{E}_{f_*} [H(\mathbf{h}_x | \mathbf{h}_x \in \bar{\mathcal{A}}_{f_*})]$ , in which the variable of the expectation  $f_*$  and the condition of the entropy  $\mathbf{h}_x \in \bar{\mathcal{A}}_{f_*}$  are not identical unlike  $\mathbb{E}_{f_*} [H(\mathbf{h}_x | f_*)]$ . Thus, (9) cannot be seen as an MI anymore, and the non-negative guarantee is not maintained. The negativity is derived by minimizing  $\alpha^{\text{CMES}}(\mathbf{x})$  with respect to  $\{z_c\}_{c=1}^C$  in Lemma 4.1. The ‘ $C > 5$ ’ is the boundary that an upper bound of this minimum becomes less than 0.

**Proof** Before the proof of Lemma 4.1, we prove the following auxiliary lemma:

**Lemma E.1.** Let  $a(\gamma) := \frac{\gamma \phi(\gamma)}{1 - \Phi(\gamma)}$  for  $\gamma \in \mathbb{R}$ . Then,  $\min_{\gamma \in \mathbb{R}} a(\gamma) < -0.29$ .

*Proof.* Using the approximation of CDF (26.2.17 in Abramowitz & Stegun, 1964), we obtain

$$\Phi(\gamma) = 1 - \phi(\gamma)(b_1 r + b_2 r^2 + b_3 r^3 + b_4 r^4 + b_5 r^5) + \epsilon(\gamma), \quad (23)$$

where

$$\begin{aligned} |\epsilon(\gamma)| &< 7.5 \times 10^{-8}, r = 1/(1 + p\gamma), p = 0.2316419, \\ b_1 &= 0.319381530, b_2 = -1.821255978, b_3 = -0.356563782, b_4 = 1.330274429, b_5 = 1.781477937. \end{aligned}$$

By substituting (23) into  $a(\gamma)$ , we obtain

$$a(\gamma) = \frac{\gamma}{b_1 r + b_2 r^2 + b_3 r^3 + b_4 r^4 + b_5 r^5 + \epsilon(\gamma)/\phi(\gamma)} \text{ for } \forall \gamma \in \mathbb{R}.$$

Since  $\epsilon(\gamma)/\phi(\gamma) = \sqrt{2\pi}e^{\gamma^2}\epsilon(\gamma) < 9\epsilon(\gamma) < 9 \times 7.5 \times 10^{-8} < 10^{-6}$  for  $\gamma \in [-1, 0]$ , we obtain an upper bound

$$a(\gamma) < \frac{\gamma}{b_1 r + b_2 r^2 + b_3 r^3 + b_4 r^4 + b_5 r^5 + 10^{-6}} \text{ for } \forall \gamma \in [-1, 0]. \quad (24)$$

This upper bound consists of elementary arithmetic that can be computed without approximations such as the numerical integration. We evaluated this upper bound (24) with 28 digits of precision by *decimal* package (<https://docs.python.org/3/library/decimal.html>) in Python, and we confirmed  $\min_{\gamma \in \mathbb{R}} a(\gamma) < a(-0.84) < -0.29$ .  $\square$

Next, we prove Lemma 4.1. By using  $\log(1 - Z) \geq (1 - \frac{1}{1-Z}) = \frac{-Z}{1-Z}$  for  $Z \in [0, 1)$ , an upper bound of the CMES acquisition function (10) can be derived as

$$\begin{aligned} & \frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} \left\{ \frac{Z_{\mathbf{x}}(\tilde{f}_*)}{2(1 - Z_{\mathbf{x}}(\tilde{f}_*))} \left( \frac{\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)\phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))}{1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))} + \sum_{c=1}^C \frac{\gamma_{\mathbf{x}}^{(g_c)}(z_c)\phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}{1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))} \right) - \log(1 - Z_{\mathbf{x}}(\tilde{f}_*)) \right\} \\ & \leq \frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} \left\{ \frac{Z_{\mathbf{x}}(\tilde{f}_*)}{2(1 - Z_{\mathbf{x}}(\tilde{f}_*))} \left( \frac{\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)\phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))}{1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))} + \sum_{c=1}^C \frac{\gamma_{\mathbf{x}}^{(g_c)}(z_c)\phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}{1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))} \right) - \left( \frac{-Z_{\mathbf{x}}(\tilde{f}_*)}{1 - Z_{\mathbf{x}}(\tilde{f}_*)} \right) \right\} \\ & = \frac{1}{K} \sum_{\tilde{f}_* \in \mathcal{F}_*} \frac{Z_{\mathbf{x}}(\tilde{f}_*)}{(1 - Z_{\mathbf{x}}(\tilde{f}_*))} \left\{ \frac{1}{2} \left( \frac{\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)\phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))}{1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))} + \sum_{c=1}^C \frac{\gamma_{\mathbf{x}}^{(g_c)}(z_c)\phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}{1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))} \right) + 1 \right\}. \end{aligned}$$

Since  $\frac{Z_{\mathbf{x}}(\tilde{f}_*)}{1 - Z_{\mathbf{x}}(\tilde{f}_*)} \geq 0$ , we see that the sign of the above upper bound is determined by

$$s(z_1, \dots, z_C) := \frac{1}{2} a(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)) + \frac{1}{2} \sum_{c=1}^C a(\gamma_{\mathbf{x}}^{(g_c)}(z_c)) + 1. \quad (25)$$

Note that  $s$  is defined as a function of  $z_c$ .

From Lemma E.1, we see

$$\min_{z_1, \dots, z_C} s(z_1, \dots, z_C) < \frac{1}{2} a(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)) - \frac{0.29C}{2} + 1.$$

Moreover, since  $\tilde{f}_*$  is sampled from  $p(f_*)$  whose support is  $\mathbb{R}$ ,  $a(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))$  can be smaller than  $-0.29$ . Hence, if  $0.29(C + 1)/2 > 1$ , (25) can be negative. Thus, (10) can be negative when  $C > 5$ , since  $0.29(C + 1)/2 = 0.29(6 + 1)/2 = 2.03/2 > 1$  when  $C = 6$ .

#### E.4. Illustrative Example that CMES Takes a Negative Value

First, we describe the detailed setting of Figure 2 omitted in the main text. In the plots, we have 200 equally spaced grid points as  $\mathcal{X}$ . We sampled  $f_*$  through the sampling from the predictive multivariate normal distribution on these grid points. We assume the posterior of  $g_c$  for  $\forall c$  are the same for brevity, and this may seem to contradict the assumption that  $f$  and  $g_c$  for  $\forall c$  are independent. However, since the independence assumption does not mean that observations of different functions cannot take the same value, Fig. 2 does not contradict the assumption. In fact, this setting (same observations for all constraints) is not an essential requirement for negativity. Typically, CMES (10) is negatively biased when  $(z_c - \mu^{(g_c)}(\mathbf{x}))/\sigma^{(g_c)}(\mathbf{x}) \approx -0.84$  for  $\forall c$  that can happen even if observations are different for each constraint. Moreover, since  $(z_c - \mu_{t-1}^{(g_c)}(\mathbf{x}))/\sigma_{t-1}^{(g_c)}(\mathbf{x}) \approx -0.84$  suggests a high probability satisfying the  $c$ -th constraint  $\Pr(g_c(\mathbf{x}) \geq z_c) \approx 0.8$ , this negative bias can be a significant problem.

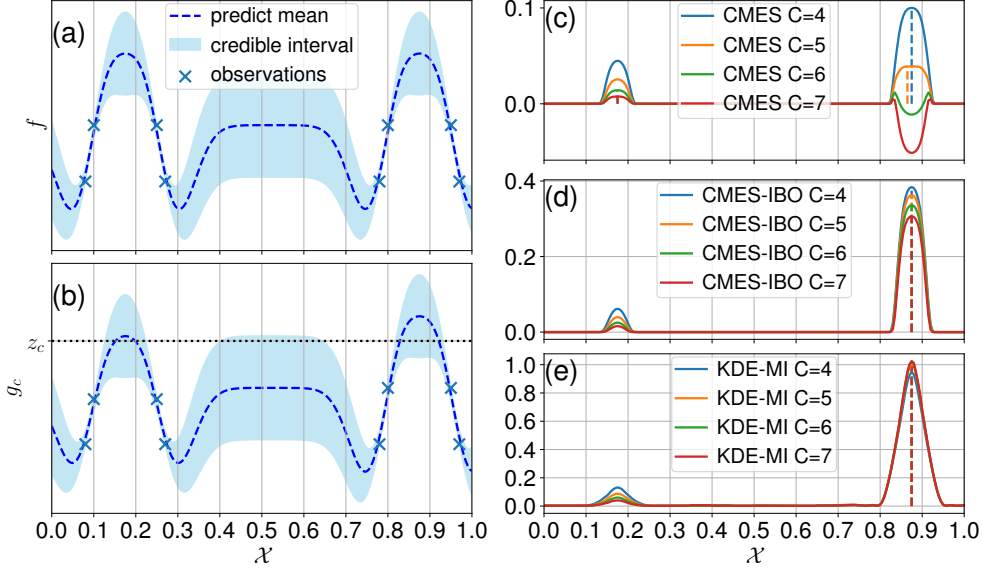


Figure 5: In (a) and (b), the predictions of  $f$  and  $g_c$  are shown, respectively. The solid line and shaded area represent the predictive mean and the credible interval, and the cross mark is the training data. The horizontal dotted line expresses  $z_c$ . In (c)-(e), CMES, CMES-IBO, and the KDE-based MI approximation for  $C = 4, 5, 6,$  and  $7$  are shown, and each vertical dotted line indicates the maximum.

The detailed procedure of KDE-MI is as follows. We used the following representation of the MI (Nair et al., 2007) because, in this form, KDE is required only for one dimensional densities:

$$\begin{aligned} \text{MI}(\mathbf{h}_x; f_*) &= \mathbb{E}_{\mathbf{h}_x} \left[ \Pr(f_* = -\infty | \mathbf{h}_x) \log \frac{\Pr(f_* = -\infty | \mathbf{h}_x)}{\Pr(f_* = -\infty)} + \Pr(f_* \neq -\infty | \mathbf{h}_x) \log \frac{\Pr(f_* \neq -\infty | \mathbf{h}_x)}{\Pr(f_* \neq -\infty)} \right. \\ &\quad \left. + \Pr(f_* \neq -\infty | \mathbf{h}_x) \int p(f_* | \mathbf{h}_x, f_* \neq -\infty) \log \frac{p(f_* | \mathbf{h}_x, f_* \neq -\infty)}{p(f_* | f_* \neq -\infty)} df_* \right], \end{aligned}$$

where the expectation over  $\mathbf{h}_x$  is approximated by the MC estimation with 10000 samples. The probabilities  $\Pr(f_* = -\infty | \mathbf{h}_x)$  and  $\Pr(f_* = -\infty)$  are estimated by the percentage of infeasible samples in 10000 samples ( $\Pr(f_* \neq -\infty | \mathbf{h}_x)$  and  $\Pr(f_* \neq -\infty)$  are estimated in the same way). The one-dimensional density functions  $p(f_* | f_* \neq -\infty)$  and  $p(f_* | \mathbf{h}_x, f_* \neq -\infty)$  are estimated by KDE fitted to feasible samples among 10000 samples of  $f_*$  and  $f_* | \mathbf{h}_x$ , respectively. Note that  $p(f_* | \mathbf{h}_x)$  is approximated by generating 10000  $f_*$  for each MC sample of  $\mathbf{h}_x$ . We approximate the one-dimensional integral with respect to  $f_*$  in the outer expectation by Gauss-Legendre quadrature. We empirically confirmed that the estimation variance of this procedure was low enough.

Here, we show an additional example in Figure 5, in which the number of the MC samplings in CMES and CMES-IBO is set as  $K = 10000$ . The other settings are the same as Figure 2. Since we set  $K$  as 10000 in Figure 5, we expect that the MC estimation are highly accurate. For CMES, we see that the behavior is slightly improved compared with the case of  $K = 10$  in Figure 2 (In particular, the acquisition function of  $C = 5$  in Figure 5 (c) is more reasonable than that of Figure 2 (c)). However, in the case of  $C = 6$  and  $7$ , CMES still takes negative values. In this example, querying to the inputs around 0.2 and 0.9 are obviously beneficial because both predictions of  $f(x)$  and  $g_c(x)$  are large. In particular, the inputs around 0.9 can be a better selection than around 0.2 because the probability of being feasible is the highest, while the uncertainty is almost the same as around 0.2. The  $\arg \max$  of KDE-MI and CMES-IBO are the same point around 0.9 regardless of the number of constraints  $C$ . On the other hand, the acquisition function values of CMES for a point around 0.9 rapidly decrease with the number of constraints  $C$ . Although CMES selects a point around 0.9 for  $C = 4$  and  $5$ , but for  $C = 6$  and  $7$ , CMES selects a point around 0.2.

## F. Derivation of Entropy of TMN

In this section, we derive the entropy of TMN truncated by a hyperrectangle  $H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  for both independent and correlated cases.

### F.1. Independent Case

In this case,  $H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  can be decomposed to the sum of the entropy of each element (Horrace, 2005):

$$H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) = H(f(\mathbf{x}) | f(\mathbf{x}) \geq \tilde{f}_*) + \sum_{c=1}^C H(g_c(\mathbf{x}) | g_c(\mathbf{x}) \geq z_c).$$

The analytical form of the entropy of the one-dimensional truncated normal distribution is known (Michalowicz et al., 2014), by which we obtain

$$\begin{aligned} H(f(\mathbf{x}) | f(\mathbf{x}) \geq \tilde{f}_*) &= \log(\sqrt{2\pi}e\sigma^{(f)}(\mathbf{x})(1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)))) + \frac{\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)\phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))}{2(1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)))}, \\ H(g_c(\mathbf{x}) | g_c(\mathbf{x}) \geq z_c) &= \log(\sqrt{2\pi}e\sigma^{(g_c)}(\mathbf{x})(1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c)))) + \frac{\gamma_{\mathbf{x}}^{(g_c)}(z_c)\phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}{2(1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c)))}. \end{aligned}$$

Moreover, from the independence of each elements of  $\mathbf{h}_x$ , we see  $H(\mathbf{h}_x) = \log(\sqrt{2\pi}e\sigma^{(f)}(\mathbf{x})) + \sum_{c=1}^C \log(\sqrt{2\pi}e\sigma^{(g_c)}(\mathbf{x}))$ . Then, the entropy of TMN is obtained as

$$H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) = H(\mathbf{h}_x) + \log(Z_{\mathbf{x}}(\tilde{f}_*)) + \frac{R_{\tilde{f}_*}}{2},$$

where

$$R_{\tilde{f}_*} = \frac{\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*)\phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))}{1 - \Phi(\gamma_{\mathbf{x}}^{(f)}(\tilde{f}_*))} + \sum_{c=1}^C \frac{\gamma_{\mathbf{x}}^{(g_c)}(z_c)\phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}{1 - \Phi(\gamma_{\mathbf{x}}^{(g_c)}(z_c))}.$$

### F.2. Correlated Case

Supposed that  $\mathbf{h}_x$  has a correlated distribution defined as

$$\mathbf{h}_x | \mathcal{D}_{t-1} \sim \mathcal{N}(\boldsymbol{\mu}(\mathbf{x}), \boldsymbol{\Sigma}(\mathbf{x})),$$

where  $\boldsymbol{\mu}(\mathbf{x})$  and  $\boldsymbol{\Sigma}(\mathbf{x})$  can be obtained by an arbitrary multi-output GP model (Rasmussen & Williams, 2005). Let  $\boldsymbol{\mu}^{\text{TN}}(\mathbf{x})$  and  $\boldsymbol{\Sigma}^{\text{TN}}(\mathbf{x})$  be the expectation and covariance matrix of  $\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}$ , respectively, whose analytical expressions are known (G & Wilhelm, 2012). We denote  $\mathbb{E}_{\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}}$ , which is the expectation with respect to  $\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}$ , as  $\mathbb{E}_{\text{TN}}$  for brevity. Then, we can transform the entropy  $H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*})$  into

$$\begin{aligned} H(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) &= \mathbb{E}_{\text{TN}} \left[ -\log p(\mathbf{h}_x | \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) \right] \\ &= \mathbb{E}_{\text{TN}} \left[ -\log \frac{p(\mathbf{h}_x)}{Z} \right] \\ &= \mathbb{E}_{\text{TN}} [-\log p(\mathbf{h}_x)] + \log Z \\ &= \frac{1}{2} \mathbb{E}_{\text{TN}} [(\mathbf{h}_x - \boldsymbol{\mu}(\mathbf{x}))^\top \boldsymbol{\Sigma}^{-1}(\mathbf{x})(\mathbf{h}_x - \boldsymbol{\mu}(\mathbf{x}))] + \frac{1}{2} |2\pi \boldsymbol{\Sigma}(\mathbf{x})| + \log Z \\ &= \frac{1}{2} \text{Tr} \left( \underbrace{\boldsymbol{\Sigma}^{-1}(\mathbf{x}) \mathbb{E}_{\text{TN}} [(\mathbf{h}_x - \boldsymbol{\mu}(\mathbf{x}))(\mathbf{h}_x - \boldsymbol{\mu}(\mathbf{x}))^\top]}_{=: \mathbf{V}} \right) + \frac{1}{2} |2\pi \boldsymbol{\Sigma}(\mathbf{x})| + \log Z, \end{aligned}$$



**Algorithm 1** Sequential- and parallel- CMES-IBO.

---

```

1: function CMES-IBO( $\mathcal{D}_0, \mathcal{X}, Q, K$ )
2:   for  $t = 0, \dots, T$  do
3:     for  $k = 1, \dots, K$  do
4:       Sample  $\tilde{f}_{(k)}, \tilde{g}_{1(k)}, \dots, \tilde{g}_{C(k)}$  from current GPs
5:        $\tilde{f}_{*(k)} \leftarrow \begin{cases} \max_{\mathbf{x} \in \tilde{\mathcal{X}}_{\text{feasible}}} \tilde{f}_{(k)}(\mathbf{x}), & \text{if } \tilde{\mathcal{X}}_{\text{feasible}} \neq \emptyset, \\ -\infty, & \text{otherwise,} \end{cases}$ 
           where  $\tilde{\mathcal{X}}_{\text{feasible}} := \{\mathbf{x} \mid \tilde{g}_c(\mathbf{x}) \geq z_c, c = 1, \dots, C\}$ 
6:     end for
7:      $\mathcal{F}_* = \{\tilde{f}_{*(k)}\}_{k=1}^K$ 
8:      $\mathbf{x}_t^{(1)} \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha^{\text{IBO}}(\mathbf{x})$  (6)
9:     if  $Q > 1$  then
10:      for  $q = 2, \dots, Q$  do
11:         $\mathcal{J} = \left\{ \left( \tilde{f}_{*(k)}, \left\{ \left( \tilde{f}_{(k)}(\mathbf{x}_t^{(q')}), \tilde{g}_{1(k)}(\mathbf{x}_t^{(q')}), \dots, \tilde{g}_{C(k)}(\mathbf{x}_t^{(q')}) \right)^\top \right\}_{q'=1}^{q-1} \right) \right\}_{k=1}^K$ 
12:         $\mathbf{x}_t^{(q)} \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \alpha_t^{\text{IBO}}(\mathbf{x} \mid \mathcal{X}_{q-1})$  (8)
13:      end for
14:    end if
15:    Evaluate  $f, g_1, \dots, g_C$  at  $\mathcal{X}_q$ 
16:    Update  $\mathcal{D}_t^{(f)}, \mathcal{D}_t^{(g_1)}, \dots, \mathcal{D}_t^{(g_C)}$  adding new observations
17:  end for
18: end function

```

---

where  $|\cdot|$  is the determinant of a matrix. Further, by using  $\mathbf{d} = \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}) - \boldsymbol{\mu}(\mathbf{x})$ , we obtain

$$\begin{aligned}
 \mathbf{V} &= \mathbb{E}_{\text{TN}}[(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}) + \mathbf{d})(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}) + \mathbf{d})^\top] \\
 &= \mathbb{E}_{\text{TN}}[(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}))(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}))^\top + \mathbf{d}(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}))^\top + (\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}))\mathbf{d}^\top + \mathbf{d}\mathbf{d}^\top] \\
 &= \mathbb{E}_{\text{TN}}[(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}))(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}))^\top] + \mathbf{d}\mathbf{d}^\top \\
 &= \boldsymbol{\Sigma}^{\text{TN}}(\mathbf{x}) + \mathbf{d}\mathbf{d}^\top,
 \end{aligned}$$

where we use  $\mathbb{E}_{\text{TN}}[(\mathbf{h}_x - \boldsymbol{\mu}^{\text{TN}}(\mathbf{x}))] = 0$ . Consequently, we can derive

$$H(\mathbf{h}_x \mid \mathbf{h}_x \in \mathcal{A}_{\tilde{f}_*}) = \frac{1}{2} \text{Tr}(\boldsymbol{\Sigma}^{-1}(\mathbf{x})(\boldsymbol{\Sigma}^{\text{TN}}(\mathbf{x}) + \mathbf{d}\mathbf{d}^\top)) + \frac{1}{2} |2\pi\boldsymbol{\Sigma}(\mathbf{x})| + \log Z.$$

## G. Computation of CMES-IBO and CMES

In this section, we describe several computational details. First, we present the general algorithm for sequential- and parallel- CMES-IBO in Algo. 1, where sequential CMES-IBO corresponds to the case that  $Q = 1$ . By replacing (6) and (8) in Algo. 1 with those of CMES (10) and (22), this algorithm can also be seen as the algorithm for sequential- and parallel- CMES.

### G.1. Sampling from Posterior

Our proposed method needs to sample the maximum value  $f_*$  defined by (1). We employ an approach using random Fourier features (RFF) (Rahimi & Recht, 2008), which has been used in various entropy-based BO and CBO methods (Hernández-Lobato et al., 2014; 2015; Wang & Jegelka, 2017). In the RFF-based sampling, RFF  $\phi(\mathbf{x}) \in \mathbb{R}^D$  are first generated, and a Bayesian linear regression model  $f(\mathbf{x}) \approx \boldsymbol{\omega}^\top \phi(\mathbf{x})$  with the coefficients  $\boldsymbol{\omega} \in \mathbb{R}^D$  is constructed. By sampling the coefficients  $\boldsymbol{\omega}$  from the Gaussian posterior, a continuous sample path can be derived. Let  $\tilde{f}_{(k)}$  and  $\tilde{g}_{c(k)}$  ( $c \in \{1, \dots, C\}$ ) be the  $k$ -th set of sample paths ( $k \in \{1, \dots, K\}$ ) for the objective and constraint functions. The  $k$ -th sample of  $f_*$ , can be obtained by solving a constrained optimization problem  $\max_{\mathbf{x} \in \mathcal{X}} \tilde{f}_{(k)}(\mathbf{x})$ , s.t.  $\tilde{g}_{c(k)}(\mathbf{x}) \geq z_c$ ,  $c = 1, \dots, C$ . We can apply any constrained optimization modules such as the method of moving asymptotes (Svanberg, 2002). Another approach

would be the discretization-based sampling through multivariate normal distribution (Perrone et al., 2019), though the number of discretization points strongly depends on the input dimension  $d$  for sufficiently accurate sampling.

For parallel querying, we need to sample the  $f_*$  and  $\mathcal{H}_{q-1}$  simultaneously. By using RFF-based sampling, these random variables are sampled from a joint distribution easily, i.e.,  $\mathcal{H}_{q-1}$  can be sampled as a value of the sample paths  $\tilde{f}_{(k)}(\mathbf{x})$  and  $\tilde{g}_{c(k)}(\mathbf{x})$  at  $\mathbf{x} \in \mathcal{X}_{q-1}$ . Moreover, in the MC approximation for  $q > 1$ , we can reuse the sample paths  $\tilde{f}_{(k)}, \tilde{g}_{1(k)}, \dots, \tilde{g}_{C(k)}$  that are sampled at  $q = 1$  as described in Algo. 1. Because of this reuse, the computational cost of parallel CMES and CMES-IBO are much smaller than  $Q$  times that of sequential CMES and CMES-IBO, respectively.

## G.2. Computational Complexity of CMES-IBO and CMES

**Independent Case** To generate sample paths, the Bayesian linear regression requires  $O(CD^3 + KD^2)$  for the given RFF, but  $C, K$ , and  $D$  are usually not large values. The number of constraints  $C$  is often less than 10. In our empirical observations, the performance of CMES-IBO were stable with a small number of MC samplings  $K$  such as 10. The dimension of the RFF features  $D$  is typically set as less than 1000. The complexity for solving the generated constraint problems depends on the solver. Many gradient-based methods have been known for the white-box constrained optimization such as the sequential quadratic programming and the interior point method. The gradient of a sample path can be obtained by  $O(dD)$ , and thus, those gradient-based solvers can be applied efficiently.

Once we obtain the samples of  $f_*$ , the acquisition functions CMES-IBO (6) and CMES (10) can be easily calculated for  $\forall \mathbf{x}$  by using the GP posterior of  $f$  and  $g_c$ . Since we employ a standard GP model, the posterior can be computed with  $O(n^3)$  for a given kernel matrix, where  $n$  is the number of the observed points. Note that the  $O(n^3)$  computation is required only once through the acquisition function maximization, because by storing  $(\mathbf{K} + \sigma_{\text{noise}}^2 \mathbf{I})^{-1}$ , the predictive mean and variance can be evaluated by  $O(n^2)$  for a given  $\mathbf{x}$ . In our setting, the function evaluation is assumed to be expensive, by which we usually only have a moderate size of a training set (typically, at most several hundreds of points).

In the case of the parallel setting, we need  $m_{\mathbf{x}^{(q)}}^{(f)}$  and  $s_{\mathbf{x}^{(q)}}^{(f)2}$ , which are the predictive mean and variance of  $f(\mathbf{x})$  after conditioning  $\mathcal{H}_{q-1}$ . This conditional density is written as  $p(f(\mathbf{x}) \mid \tilde{f}_{(k)}(\mathbf{x}^{(1)}), \dots, \tilde{f}_{(k)}(\mathbf{x}^{(q-1)}))$ , which can be easily determined through the standard conditional Gaussian formula. Given the  $q \times q$  predictive covariance matrix of the current GP for  $\mathbf{x}^{(q)}$  and  $\mathcal{X}_{q-1}$ , we can obtain  $m_{\mathbf{x}^{(q)}}^{(f)}$  and  $s_{\mathbf{x}^{(q)}}^{(f)2}$  with  $O(q^3)$ . For  $m_{\mathbf{x}^{(q)}}^{(g_c)}$  and  $s_{\mathbf{x}^{(q)}}^{(g_c)2}$ , which are for constraint functions, the same calculation can be applied.

**Correlated Case** For multi-output GP model, the posterior can be computed with  $O(C^3 n^3)$ . For CMES-IBO and CMES, the computations of CDF of the multivariate normal distribution are needed in the correlated case, as shown in Appendix F.2. CMES-IBO performs  $K$  times  $C + 1$  dimensional CDF for the acquisition function (6), in which  $Z_{\mathbf{x}}(f_*) = \Pr(\mathbf{h} \in \mathcal{A}_{\tilde{f}_*})$  is now CDF of a correlated multivariate normal distribution. On the other hand, in addition to these CDF computations, CMES performs  $K$  times computations of the expectation and covariance matrix of TMN in the acquisition function (21), each of which needs computing  $C$  dimensional CDF  $C + 1$  times and  $C - 1$  dimensional CDF  $C(C + 1)/2$  times, respectively (G & Wilhelm, 2012). Using a well-known MC-based algorithm (Genz, 1992), the computational complexity of CDF with respect to the dimension  $C$  is known to be  $O(C^2)$ . Therefore, the computational complexity of CMES-IBO with respect to  $K$  and  $C$  is  $O(KC^2)$ , but that of CMES is  $O(KC^4)$ , which severely limits the applicability of CMES to the large number of constraints.

## H. Considerations on Regret Bound of MES

In this section, we describe flaws in the theoretical analysis of MES (Wang & Jegelka, 2017). We use the same notation as (Wang & Jegelka, 2017) throughout this section. Wang & Jegelka (2017) showed the bound of simple regret  $r_T := \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - \max_{t \in [1, T]} f(\mathbf{x}_t)$  as follows:

**Theorem H.1** (Theorem 3.2 in (Wang & Jegelka, 2017)). *Let  $F$  be the cumulative probability distribution for the maximum of any function  $f$  sampled from  $GP(\mu, k)$  over the compact search space  $\mathcal{X} \subset \mathbb{R}^d$ , where  $k(\mathbf{x}, \mathbf{x}') \leq 1, \forall \mathbf{x}, \mathbf{x}' \in \mathcal{X}$ . Let  $f_* = \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$  and  $w = F(f_*) \in (0, 1)$ , and assume the observation noise is iid  $\mathcal{N}(0, \sigma^2)$ . If in each iteration  $t$ , the query point is chosen as  $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \gamma_{y_*^t}(\mathbf{x}) \frac{\psi(\gamma_{y_*^t}(\mathbf{x}))}{2\Psi(\gamma_{y_*^t}(\mathbf{x}))} - \log(\Psi(\gamma_{y_*^t}(\mathbf{x})))$ , where  $\gamma_{y_*^t}(\mathbf{x}) = \frac{y_*^t - \mu_t(\mathbf{x})}{\sigma_t(\mathbf{x})}$  and  $y_*^t$  is*

drawn from  $F$ , then with probability at least  $1 - \delta$ , in  $T' = \sum_{i=1}^T \log_w \frac{\delta}{2\pi_i}$  number of iterations, the simple regret satisfies

$$r_{T'} \leq \sqrt{\frac{C\rho_T}{T}}(\nu_{t^*} + \zeta_T), \quad (26)$$

where  $C = 2/\log(1 + \sigma^{-2})$  and  $\zeta_T = (2\log(\frac{\pi_T}{\delta}))^{\frac{1}{2}}$ ,  $\pi_i$  satisfies  $\sum_{i=1}^T \pi_i^{-1} \leq 1$  and  $\pi_t > 0$ , and  $t^* = \arg \max_t \nu_t$  with  $\nu_t \triangleq \min_{\mathbf{x} \in \mathcal{X}, y_*^t > f_*} \gamma_{y_*^t}(\mathbf{x})$ , and  $\rho_T$  is the maximum information gain of at most  $T$  selected points.

In this theorem, the simple regret for ‘one sample MES’ is analyzed, in which only one max-value is sampled at every iteration in the algorithm. We summarize five main flaws in this theorem as follows (note that each of them is related to each other):

1. From the assumption  $f \sim \mathcal{GP}(\mu, k)$ , the max-value  $f_*$  is a random variable, but it is not treated as a random variable in their analysis. For example, Lemma C.1 in (Wang & Jegelka, 2017) obviously assumes that  $f(\mathbf{x}_t)$  follows the Gaussian distribution, and therefore, treating  $f_*$  as a deterministic variable contradicts the assumption of this lemma. This problem closely related to other flaws, and thus, it is difficult to correct.
2. In BO literature, usually, the number of iterations is represented by  $T$ , and Wang & Jegelka (2017) use this notation except for the theorem. On the other hand, in the theorem, the number of iteration is represented by  $T'$  instead of  $T$ , and the variable  $T$  in the bound (26) is defined by the number of ‘partitioning’ of the entire iterations in their proof. Each partition is defined so that it satisfies a probabilistic condition related to max-values, which we omit details here. An important issue is that, for this replaced  $T$ , dependency on the actual number of iterations is not clarified. Therefore, (26) does not reveal the convergence rate with respect to the number of iterations.
3. Since  $\nu_{t^*}$  depends on  $y_*^t$  in each iteration,  $\nu_{t^*}$  is a random variable. This variable remains in the final upper bound, but the convergence rate of  $\nu_{t^*}$  has not been shown.
4. In the theorem, Wang & Jegelka (2017) assumed that  $y_*^t$  is drawn from the prior  $F$ . This assumption clearly does not match the basic idea of MES that selects the next query based on the mutual information estimated through samples of the max-value from the ‘posterior’ distribution.
5. Wang & Jegelka (2017) claimed that the above simple regret bound can adapt to the setting that  $y_*^t$  is drawn from the posterior. However, this claim is not proven and not obvious. In the proof, Wang & Jegelka (2017) implicitly assumed  $\Pr(y_*^t < f_*, y_*^{t+1} < f_*) = \Pr(y_*^t < f_*) \Pr(y_*^{t+1} < f_*)$ . Through the basic formula of the joint probability of independent variables, this decomposition is allowed if  $f_*$  is deterministic, and  $y_*^t$  and  $y_*^{t+1}$  are independent of each other. However, both of these two conditions are not satisfied in the case of sampling from the posterior. For the first condition, as we already mention,  $f_*$  should be a random variable (actually, this first condition is not satisfied even for the case of sampling from the prior). For the second condition, when  $y_*^t$  and  $y_*^{t+1}$  are sampled from posterior,  $y_*^t$  depends on  $\mathcal{D}_{t-1}$  and  $y_*^{t+1}$  depends on  $\mathcal{D}_t$ . However, since  $\mathcal{D}_{t-1} \subset \mathcal{D}_t$ , dependency obviously exists between  $y_*^t$  and  $y_*^{t+1}$  (note that now  $\mathcal{D}_{t-1}$  and  $\mathcal{D}_t$  are also random variables not only because of the random noise of observations, but also the randomness in the acquisition function). In the proof, the decomposition is not justified when the two conditions are not satisfied. This decomposition plays a key role in the proof, and we consider that the above issue is not easy to avoid.

We first tried to extend this theorem to our constrained problem, but we consider that correcting the above issues is at least not trivial. Therefore, theoretical analysis of the MES-based methods (even for the simple unconstrained case) is still an open problem, to the best of our knowledge.

## I. Details of Experimental Settings

In this section, we describe additional details of experimental settings.

### I.1. Other Experimental Settings

Observations for objective and constraint functions are standardized so that they have zero mean and unit variance, respectively, before the regression at every iteration. We used the open-source library called GPpy (GPpy, since 2012) for

the GP regression. The candidate intervals of hyperparameters  $\sigma_{\text{LIN}}^2$  and  $\sigma_{\text{RBF}}^2$  are set to  $[0, 1]$ , and the variance of the noise  $\sigma^2$  is set to  $10^{-6}$ . As we mentioned in the main paper, other hyperparameters are chosen by marginal likelihood maximization. We employed an Automatic Relevance Determination (ARD) for the RBF kernel. Thus, the RBF kernel has a hyperparameter  $\ell := (\ell_1, \dots, \ell_d)^\top \in \mathbb{R}^d$ , and it is written as

$$k_{\text{RBF}}(\mathbf{x}, \mathbf{x}') := \exp\left(-\frac{1}{2} \sum_{i=1}^d \frac{(x_i - x'_i)^2}{\ell_i^2}\right),$$

where  $x_i$  and  $x'_i$  are the  $i$ -th elements of  $\mathbf{x}$  and  $\mathbf{x}'$ , respectively. For each one of  $\ell_i$  for  $i = 1, \dots, d$ , we set the same interval of the candidate value. For the benchmark functions and the reactor network design problem, we set  $\ell_i \in [10^{-1}s_i, 10s_i]$ , where  $s_i$  is the length of the interval for the possible value of  $x_i$  in the input domain, and the thresholds  $z_c = 0$  for  $c = 1, \dots, C$ . For the CNN hyperparameter optimization problem, we set  $\ell_i \in [10^{-1}, 2]$  and the thresholds  $z_c = 0.5$  for  $c = 1, \dots, C$ . We used the method of moving asymptotes (Svanberg, 2002) in NLOpt (Johnson, since 2008) for the constrained optimization of the sample path in CMES-IBO, CMES, TSC, and their parallel extensions. The numbers of initial points are set as 3 for the GP-derived synthetic function, 5 for the two-dimensional benchmark functions,  $5d$  for the Hartmann6 function, and 25 for others. Finally, the hyperparameter  $M$  of P-EIC is selected as the minimum of the current observations.

## 1.2. Settings for PESC

Due to the difficulty of rewriting the Spearmint specification, the default settings such as for kernels and inner optimization were used. Although PESC has the same number of initial points, their locations are not identical to other methods because of the same reason.

## 1.3. Details of Benchmark Functions

Here, we provide detailed information on benchmark functions. Note that we change the sign from the original functions if it is required to formalize as the maximization problem.

**Gardner1** Gardner1 is a simple test problem in which functions are constructed by sine and cosine functions (Gardner et al., 2014). The input dimension  $d = 2$ , the input domain  $\mathcal{X} = [0, 6]^2$ , and the number of constraints  $C = 1$ . The detailed forms of functions are

$$\begin{aligned} f(\mathbf{x}) &= -\cos(2x_1) \cos(x_2) - \sin(x_1), \\ g_1(\mathbf{x}) &= -\cos(x_1) \cos(x_2) + \sin(x_1) \sin(x_2) + 0.5. \end{aligned}$$

**Gardner2** Gardner2 also has a simple form, but its feasible region is very small (Gardner et al., 2014). The input dimension  $d = 2$ , the input domain  $\mathcal{X} = [0, 6]^2$ , and the number of constraints  $C = 1$ . The detailed forms of functions are

$$\begin{aligned} f(\mathbf{x}) &= -\sin(x_1) - x_2, \\ g_1(\mathbf{x}) &= -\sin(x_1) \sin(x_2) - 0.95. \end{aligned}$$

**Gramacy** The Gramacy function is from (Gramacy et al., 2016). The input dimension  $d = 2$ , the input domain  $\mathcal{X} = [0, 1]^2$ , and the number of constraints  $C = 2$ . The detailed forms of functions are

$$\begin{aligned} f(\mathbf{x}) &= -x_1 - x_2, \\ g_1(\mathbf{x}) &= \frac{1}{2} \sin(2\pi(x_1^2 - 2x_2)) + x_1 + 2x_2 - 1.5, \\ g_2(\mathbf{x}) &= -x_1^2 - x_2^2 + 1.5. \end{aligned}$$

**Hartmann6** The Hartmann6 function is used in (Letham et al., 2019). The input dimension  $d = 6$ , the input domain



$\mathcal{X} = [0, 1]^6$ , and the number of constraints  $C = 1$ . The detailed forms of functions are

$$f(\mathbf{x}) = \sum_{i=1}^4 \alpha_i \exp\left(-\sum_{j=1}^6 A_{ij}(x_j - P_{ij})^2\right),$$

$$g_1(\mathbf{x}) = -\|\mathbf{x} - \mathbf{x}'\| + 1,$$

where

$$\boldsymbol{\alpha} = (1.0, 1.2, 3.0, 3.2)^\top,$$

$$\mathbf{A} = \begin{pmatrix} 10 & 3 & 17 & 3.5 & 1.7 & 8 \\ 0.05 & 10 & 17 & 0.1 & 8 & 14 \\ 3 & 3.5 & 1.7 & 10 & 17 & 8 \\ 17 & 8 & 0.05 & 10 & 0.1 & 14 \end{pmatrix},$$

$$\mathbf{P} = 10^{-4} \begin{pmatrix} 1312 & 1696 & 5569 & 124 & 8283 & 5886 \\ 2329 & 4135 & 8307 & 3736 & 1004 & 9991 \\ 2348 & 1451 & 3522 & 2883 & 3047 & 6650 \\ 4047 & 8828 & 8732 & 5743 & 1091 & 381 \end{pmatrix}.$$

#### I.4. Details of CNN hyperparameter optimization problem

We used PyTorch (Paszke et al., 2019) to construct the CNN model. The search domain for this problem is set as follows: the learning rate in  $\{10^{-3}, 10^{-2}, 10^{-1}, 10^0\}$ , the dropout rate in  $\{2^{-4}, 2^{-3}, 2^{-2}, 2^{-1}\}$ , the number of channels in  $\{2^3, 2^4, 2^5, 2^6\}$  for two layers, and the coefficient  $\rho$  in  $\{0, 0.1, \dots, 1.9\}$ . Thus, the size of input domain  $|\mathcal{X}| = 5120$ . To define  $\mathbf{x}$ , we apply  $\log_{10}$  and  $\log_2$  transformation to the learning rate and the other inputs except for  $\rho$ , respectively. Moreover, we apply the logit transformation to accuracy and recall for transforming the output domain from  $[0, 1]$  to  $\mathbb{R}$ . Under the setting  $z_c = 0.5$  for  $c = 1, \dots, 10$ , the number of feasible points is 28, which is only about 0.5% ( $\approx 28/5120$ ) of the entire candidates  $\mathcal{X}$ . The unconstrained optimal point of the accuracy  $\arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$  is different from the constrained optimal point  $\mathbf{x}_*$ , meaning that the unconstrained optimization of the average accuracy sacrifices recalls of some classes. For the training of CNN, we used the default weighted cross-entropy loss function in PyTorch, which is defined as follows:

$$\text{LOSS}(\mathbf{l}, \hat{\mathbf{l}}) = \sum_{i=1}^N -\frac{w_{l_i}}{\sum_{m=1}^M w_m} \log \hat{l}_{i,l_i},$$

where  $\mathbf{l} = (l_1, \dots, l_N) \in \mathbb{R}^N$ ,  $l_i$  is a correct label of  $i$ -th data,  $\hat{l}_{i,m}$  is a predicted probability that  $i$ -th data has label  $m$ , and a  $(i, m)$  element of  $\hat{\mathbf{l}} \in \mathbb{R}^{N \times M}$  is  $\hat{l}_{i,m}$ .

## J. Additional Experiments

### J.1. Results with the Different Number of MC samples

By using benchmark functions from CBO literature, we here show the results with three different numbers of MC samples  $K = 1, 10$ , and 50 of CMES and CMES-IBO. Figure 6 shows that even with small  $K = 1$ , these two methods maintain reasonable performance compared with other methods. In particular, for CMES-IBO, this result can be explained by the fact that  $\alpha^{\text{IBO}}(\mathbf{x})$  can be seen as the PI from  $\tilde{f}_*$  even with  $K = 1$  (Remark 4.1).

### J.2. Experimental Results of the Utility Gap Defined by Best Observed Value

Figure 7 shows the utility gap defined by the best observed value, which is  $\min\{f_* - f(\mathbf{x}_i) \mid (\mathbf{x}_i, y_i) \in \mathcal{D}_{t-1}, g_c(\mathbf{x}_i) \geq z_c, \forall c\}$  when feasible solution is already obtained, otherwise  $f_* - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$ . We see that all the plots show the similar behavior as Fig. 3. This indicates that the recommendation  $\hat{\mathbf{x}}_t$  is mostly selected around the training data since  $\hat{\mathbf{x}}_t$  is chosen under constraints  $\Pr(g_c(\mathbf{x}) \geq z_c) \geq \sqrt[4]{0.95}$  for  $\forall c$ . Note that PESC is not shown in Fig. 7 since Spearmint does not provide the best observed values.


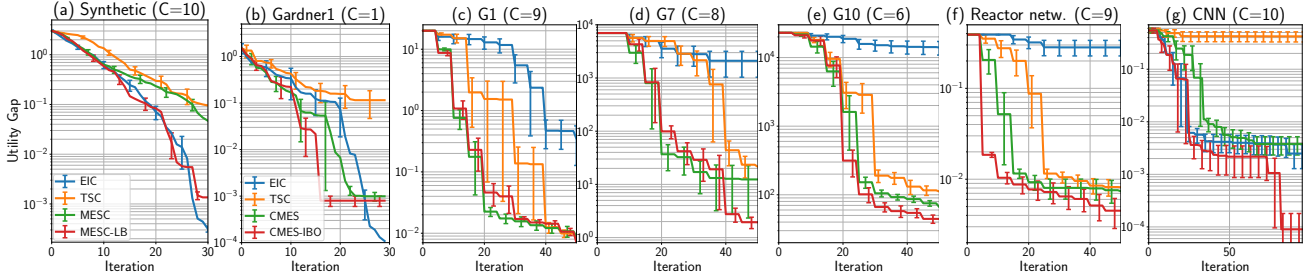

 Figure 6: Utility gap of benchmark functions with different  $K$  (average and standard error).


Figure 7: Utility gap defined by the best observed value in sequential querying (average and standard error).

### J.3. Experiment for Correlated Synthetic Function

We evaluate the performance of CMES-IBO with correlated GPs on synthetic functions ( $d = 3$  and  $C = 3$ ). Because of its high computational complexity, we did not employ CMES here (see Appendix G.2). The objective and  $C = 3$  constraints were sampled from GPs with the RBF kernel, and the length scales of the kernels were 0.1. All the correlation coefficients among the objective and constraint functions were set as 0.5. The same settings of GPs for the function generation were also used in each BO method. The number of initial points is set as 5. The threshold for constraints were set as  $z_c = 0$  for  $\forall c$ , and the input domain was  $[0, 1]^3$ . We sampled 5 sets of functions, and the experiment of each set ran 10 times. We report the mean and standard error of these 50 trials. Figure 8 shows the result. C-CMES-IBO denotes the correlated extension of CMES-IBO. We see that C-CMES-IBO improves CMES-IBO in this case and is superior to the other methods. Although the correlated model can improve performance when the true function has a strong correlation, we mainly focus on the independent setting because the efficiency is often comparable to the correlated model, and the computations of independent CMES-IBO is much simpler than correlated case.

### J.4. Experiment for Computational Time Evaluation

In Table 1, we evaluate computational time of the acquisition function maximization at  $t = 30$ . EIC and TSC were relatively faster than the others. CMES-IBO and CMES took longer times because of sampling of  $f_*$  for which we need to solve constraint optimization problems  $K$  times, while the acquisition function maximization itself was fast enough because of their simple closed forms of the acquisition functions. Although it is not shown here, PESC has the same computational requirement for sampling  $\mathbf{x}_*$  because the same constraint optimization problems need to be solved. It is worth noting that the  $K$  times sampling of  $f_*$  can be accelerated by computing in parallel because each one of the samplings is independent.

 Table 1: Computational time (sec.) for the Gramacy function at  $t = 30$  (mean  $\pm$  standard error)

	sampling of $f_*$	optimization of $\alpha(\mathbf{x})$
EIC	NA	$1.148 \pm 0.287$
TSC	NA	$2.670 \pm 0.464$
CMES	$20.728 \pm 1.498$	$0.780 \pm 1.317$
CMES-IBO	$21.836 \pm 0.933$	$1.163 \pm 2.130$




Figure 8: Utility gap of correlated GP-derived synthetic functions (average and standard error).