
Bayesian Optimization for Distributionally Robust Chance-constrained Problem

Yu Inatsu¹ Shion Takeno¹ Masayuki Karasuyama¹ Ichiro Takeuchi^{1,2}

Abstract

In black-box function optimization, we need to consider not only controllable design variables but also uncontrollable stochastic environment variables. In such cases, it is necessary to solve the optimization problem by taking into account the uncertainty of the environmental variables. Chance-constrained (CC) problem, the problem of maximizing the expected value under a certain level of constraint satisfaction probability, is one of the practically important problems in the presence of environmental variables. In this study, we consider distributionally robust CC (DRCC) problem and propose a novel DRCC Bayesian optimization method for the case where the distribution of the environmental variables cannot be precisely specified. We show that the proposed method can find an arbitrary accurate solution with high probability in a finite number of trials, and confirm the usefulness of the proposed method through numerical experiments.

1. Introduction

In this study, we consider a black-box function optimization problem with two types of variables called *design variables* which are fully controllable and *environmental variables* which change randomly depending on the uncertainty of the environment. Under the presence of these two types of variables, the goal is to identify the design variables that optimize the black-box function by taking into account the uncertainty of environmental variables. In the past few years, Bayesian Optimization (BO) framework that takes the uncertain environmental variables into considerations have been studied in various setups (see §1.1). In this paper,

¹Department of Computer Science, Nagoya Institute of Technology, Aichi, Japan ²RIKEN Center for Advanced Intelligence Project, Tokyo, Japan. Correspondence to: Yu Inatsu <inatsu.yu@nitech.ac.jp>.

we study one of such problems called *distributionally robust chance-constrained (DRCC)* problem. The DRCC problem is an instance of constrained optimization problems in an uncertain environment, which is important in a variety of practical problems in science and engineering.

The goal of a CC problem is to identify the design variables that maximize the expectation of the objective function under the constraint that the probability of the constraint function exceeding a given threshold is greater than a certain level. Let $f(\mathbf{x}, \mathbf{w})$ and $g(\mathbf{x}, \mathbf{w})$ be the unknown objective and constraint functions, respectively, both of which depend on the design variables $\mathbf{x} \in \mathcal{X}$ and the environmental variables $\mathbf{w} \in \Omega$. For a given threshold $h \in \mathbb{R}$ and a level $\alpha \in (0, 1)$, the CC problem is formulated as

$$\operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \int_{\Omega} f(\mathbf{x}, \mathbf{w}) p^{\dagger}(\mathbf{w}) d\mathbf{w} \quad (1a)$$

$$\text{subject to } \int_{\Omega} \mathbb{I}[g(\mathbf{x}, \mathbf{w}) > h] p^{\dagger}(\mathbf{w}) d\mathbf{w} > \alpha, \quad (1b)$$

where $\mathbb{I}[\cdot]$ is the indicator function and $p^{\dagger}(\mathbf{w})$ is the probability density function of the environmental variables \mathbf{w} . When $p^{\dagger}(\mathbf{w})$ is known, there is a method to solve the CC problem (see §1.1).

In this study, we consider the case where $p^{\dagger}(\mathbf{w})$ is unknown as is commonly encountered in practice. Here, we formulate the uncertainty of $p^{\dagger}(\mathbf{w})$ using a measure called *distributionally robustness*. Let \mathcal{A} be the user-specified candidate distribution family of \mathbf{w} . Then, the DRCC problem is defined as

$$\operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} F(\mathbf{x}) \text{ subject to } G(\mathbf{x}) > \alpha, \quad (2)$$

where $F(\mathbf{x})$ and $G(\mathbf{x})$ are defined as

$$F(\mathbf{x}) = \inf_{p(\mathbf{w}) \in \mathcal{A}} \int_{\Omega} f(\mathbf{x}, \mathbf{w}) p(\mathbf{w}) d\mathbf{w}, \quad (3a)$$

$$G(\mathbf{x}) = \inf_{p(\mathbf{w}) \in \mathcal{A}} \int_{\Omega} \mathbb{I}[g(\mathbf{x}, \mathbf{w}) > h] p(\mathbf{w}) d\mathbf{w}. \quad (3b)$$

Many real-world engineering problems can be formulated as DRCC problems, including controlling the position of a robot, controlling the switching of wind turbines, and

optimizing an infectious disease simulator. The solution of this problem is robust with respect to the misspecification of the distributions because $F(\mathbf{x})$ and $G(\mathbf{x})$ are defined by considering the worst case scenario among the candidate distribution families.

For the surrogate models for unknown objective function $f(\mathbf{x}, \mathbf{w})$ and the constraint function $g(\mathbf{x}, \mathbf{w})$, we employ Gaussian Process (GP) models and study the above DRCC problem in the context of BO framework. The main technical challenges in this problem is in the characterization of the posterior distributions of $F(\mathbf{x})$ and $G(\mathbf{x})$. In this study, we derive credible intervals of $F(\mathbf{x})$ and $G(\mathbf{x})$ which can be effectively used for solving the DRCC problem in BO framework. We call the proposed method *Distributionally Robust Chance Constrained Bayesian Optimization (DRCC-BO)* method.

1.1. Related Work

Black-box function optimization problems using a GP surrogate model (Williams & Rasmussen, 2006) have been extensively studied in the context of BO (see, e.g., (Settles, 2009; Shahriari et al., 2016)). The constraint in the form of (2) is closely related with level set estimation (LSE) problem in which a GP surrogate model is often employed (Bryan et al., 2006; Gotovos et al., 2013; Zanette et al., 2018; Inatsu et al., 2020a; Sui et al., 2015; Turchetta et al., 2016; Sui et al., 2018; Wachi et al., 2018). The problem (2) is also closely related to constrained BO, which has also been studied extensively in the literature (Gardner et al., 2014; Hernández-Lobato et al., 2016). In the past few years, various problem settings concerning uncertain environmental variables have been considered in BO and LSE problems. The most standard approach to deal with the uncertainty in environmental variables is to consider the expected value of $f(\mathbf{x}, \mathbf{w})$ or/and $g(\mathbf{x}, \mathbf{w})$. Fortunately, when the GP is employed as the surrogate model for $f(\mathbf{x}, \mathbf{w})$ and $g(\mathbf{x}, \mathbf{w})$, its expected value is also represented as a GP, so the acquisition functions (AFs) of BO and LSE problems can be easily constructed.

On the other hand, in many practical problems, the expected value is often not enough, and other risk measures that can no longer be expressed as GPs, such as variance and tail probability, need to be considered. The objective function $F(\mathbf{x})$ and the constraint function $G(\mathbf{x})$ of the DRCC problem (2) are also difficult to handle because they are also not represented as GPs even when $f(\mathbf{x}, \mathbf{w})$ and $g(\mathbf{x}, \mathbf{w})$ are GPs. In the past few years, there have been several studies that consider various risk measures of the uncertainty in the environmental variables \mathbf{w} and robustness with respect to $p^\dagger(\mathbf{w})$ (Iwazaki et al., 2020; 2021a; Inatsu et al., 2020b; Bogunovic et al., 2018; Nguyen et al., 2021b;a; Iwazaki et al., 2021b; Inatsu et al., 2021).

In (Amri et al., 2021) proposed a BO method for the CC problem, but they assumed that the distribution $p^\dagger(\mathbf{w})$ is known.

Distributionally robust optimization (DRO) problems have long been studied in robust optimization community for ordinary optimization problems in which the objective function and the constraint functions are explicitly formulated (in contrast to expensive black-box functions as we consider in this study) (Scarf, 1958; Rahimian & Mehrotra, 2019). DRCC problem with explicitly formulated objective and constraint functions were also studied in (Xie, 2021; Ho-Nguyen et al., 2021), and they were applied to practical problems called power flow optimization (Xie & Ahmed, 2017; Fang et al., 2019). On the other hand, the study of DRO problems for black-box functions with high evaluation cost has only recently started. In (Kirschner et al., 2020; Nguyen et al., 2020), BO methods to find the design variable that maximizes $F(\mathbf{x})$ in (3a) was studied in DR setting. Furthermore, in (Inatsu et al., 2021), a BO method to efficiently identify the design variables which satisfy $G(\mathbf{x}) > \alpha$ for $G(\mathbf{x})$ in (3b) was studied in DR setting. However, to the best of our knowledge, there is no existing studies that can be used directly in the DRCC-BO problem.

1.2. Contribution

The contributions of this paper are as follows:

- A BO method for DRCC problem called DRCC-BO method is proposed. Specifically, we propose a novel AF for DRCC problem based on the credible intervals of $F(\mathbf{x})$ in (3a) and $G(\mathbf{x})$ in (3b) when GPs are employed as the surrogate models for $f(\mathbf{x}, \mathbf{w})$ and $g(\mathbf{x}, \mathbf{w})$.
- Under mild conditions, we showed that the proposed method can find an arbitrarily accurate solution to the DRCC problem with high probability in a finite number of trials.
- We also showed that by designing the DRCC problem with an appropriate choice of candidate distribution families, without knowing the true distribution $p^\dagger(\mathbf{w})$ the proposed method can find an arbitrary accurate solution even for the CC problem (Theorem 4.6).
- The performance of the proposed method is confirmed through numerical experiments with synthetic as well as simulator-based functions.

2. Preliminary

Let $f : \mathcal{X} \times \Omega \rightarrow \mathbb{R}$ and $g : \mathcal{X} \times \Omega \rightarrow \mathbb{R}$ be the expensive-to-evaluate black-box functions. We assume that \mathcal{X} and Ω are finite sets. For each $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$, the values of $f(\mathbf{x}, \mathbf{w})$ and $g(\mathbf{x}, \mathbf{w})$ are observed as $y^{(f)} = f(\mathbf{x}, \mathbf{w}) + \varepsilon_f$

and $y^{(g)} = g(\mathbf{x}, \mathbf{w}) + \varepsilon_g$, where ε_f and ε_g are independent Gaussian distributions following $\varepsilon_f \sim \mathcal{N}(0, \sigma_{f,\text{noise}}^2)$, $\varepsilon_g \sim \mathcal{N}(0, \sigma_{g,\text{noise}}^2)$. In this study, we consider the following two cases for the observation of \mathbf{w} :

Uncontrollable: For each trial t , \mathbf{w} cannot be controlled, and its realization is generated from the unknown distribution P^\dagger .

Simulator: For each trial t , \mathbf{w} can be chosen arbitrarily.

For example, in the robot position control and wind turbine switching problems described in §1, these settings can be formulated as the uncontrollable setting if actual values of random environment \mathbf{w} (and f and g) can be obtained sequentially while actually performing the work. On the other hand, if the simulation can be performed while arbitrarily changing the value of a random environmental variable \mathbf{w} (e.g., recovery rate), as in the simulation of an infectious disease, it can be interpreted as the simulator setting. Moreover, we consider the following \mathcal{A}_t as a family of candidate distributions for P^\dagger :

$$\mathcal{A}_t = \{\text{probability function } p(\mathbf{w}) \mid d(p(\mathbf{w}), p_t^*(\mathbf{w})) \leq \epsilon_t\},$$

where $p_t^*(\mathbf{w})$ is a user-specified reference distribution, $d(\cdot, \cdot)$ is a given distance function between distributions, and $\epsilon_t > 0$. Then, under a given threshold h , the DR expectation function $F_t(\mathbf{x})$ and DR probability function $G_t(\mathbf{x})$ are defined for each $\mathbf{x} \in \mathcal{X}$ and $t \geq 1$ as

$$F_t(\mathbf{x}) = \inf_{p(\mathbf{w}) \in \mathcal{A}_t} \sum_{\mathbf{w} \in \Omega} f(\mathbf{x}, \mathbf{w}) p(\mathbf{w}),$$

$$G_t(\mathbf{x}) = \inf_{p(\mathbf{w}) \in \mathcal{A}_t} \sum_{\mathbf{w} \in \Omega} \mathbb{1}[g(\mathbf{x}, \mathbf{w}) > h] p(\mathbf{w}).$$

The objective of this study is to efficiently find the optimal design variable \mathbf{x}_t^* that maximizes $F_t(\mathbf{x})$ such that $G_t(\mathbf{x})$ exceeds a given level $\alpha \in (0, 1)$. In other words, \mathbf{x}_t^* satisfies that

$$\mathbf{x}_t^* = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} F_t(\mathbf{x}) \quad \text{subject to} \quad G_t(\mathbf{x}) > \alpha.$$

If the optimal solution \mathbf{x}_t^* does not exist, it is formally defined as $F(\mathbf{x}_t^*) = \min_{\mathbf{x} \in \mathcal{X}} F_t(\mathbf{x})$.

2.1. Gaussian Process

In this study, we use GPs as the surrogate model for the black-box functions f and g . First, we assume that the GPs, $\mathcal{GP}(0, k^{(f)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')))$, and $\mathcal{GP}(0, k^{(g)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')))$ are prior distributions of f and g , respectively. Here, $k^{(f)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}'))$, and $k^{(g)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}'))$ are positive-definite kernels. Then, under the given dataset $\{(\mathbf{x}_i, \mathbf{w}_i, y_i^{(f)})\}_{i=1}^t$, the posterior

distribution of f also follows the GP, and its posterior mean $\mu_t^{(f)}(\mathbf{x}, \mathbf{w})$ and posterior variance $\sigma_t^{(f)2}(\mathbf{x}, \mathbf{w})$ are given by

$$\mu_t^{(f)}(\mathbf{x}, \mathbf{w}) = \mathbf{k}_t^{(f)}(\mathbf{x}, \mathbf{w})^\top (\mathbf{K}_t^{(f)} + \sigma_{f,\text{noise}}^2 \mathbf{I}_t)^{-1} \mathbf{y}_t^{(f)},$$

$$\sigma_t^{(f)2}(\mathbf{x}, \mathbf{w}) = k^{(f)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}, \mathbf{w})) - \mathbf{k}_t^{(f)}(\mathbf{x}, \mathbf{w})^\top (\mathbf{K}_t^{(f)} + \sigma_{f,\text{noise}}^2 \mathbf{I}_t)^{-1} \mathbf{k}_t^{(f)}(\mathbf{x}, \mathbf{w}),$$
(4)

where $\mathbf{k}_t^{(f)}(\mathbf{x}, \mathbf{w})$ is a t -dimensional vector with i th element $k^{(f)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}_i, \mathbf{w}_i))$, $\mathbf{y}_t^{(f)} = (y_1^{(f)}, \dots, y_t^{(f)})^\top$, \mathbf{I}_t , being the $t \times t$ identity matrix, and $\mathbf{K}_t^{(f)}$ is the $t \times t$ matrix with (j, k) element $k^{(f)}((\mathbf{x}_j, \mathbf{w}_j), (\mathbf{x}_k, \mathbf{w}_k))$. As in the case of f , under the dataset $\{(\mathbf{x}_i, \mathbf{w}_i, y_i^{(g)})\}_{i=1}^t$, the posterior distribution of g is also a GP, and its posterior mean $\mu_t^{(g)}(\mathbf{x}, \mathbf{w})$ and posterior variance $\sigma_t^{(g)2}(\mathbf{x}, \mathbf{w})$ can be obtained by using the same formula as (4).

3. Proposed Method

In this section, we propose a BO method for efficiently solving the DRCC problem. In our setting, because $f(\mathbf{x}, \mathbf{w})$ and $g(\mathbf{x}, \mathbf{w})$ are random functions, $F_t(\mathbf{x})$ and $G_t(\mathbf{x})$ are also random functions. Thus, one of the natural BO methods is to use credible intervals of $F_t(\mathbf{x})$ and $G_t(\mathbf{x})$. Unfortunately, although $f(\mathbf{x}, \mathbf{w})$ and $g(\mathbf{x}, \mathbf{w})$ follow GPs, $F_t(\mathbf{x})$ and $G_t(\mathbf{x})$ do not follow GP. Hence, credible intervals of $F_t(\mathbf{x})$ and $G_t(\mathbf{x})$ cannot be constructed based on the property of Normal distribution. In the next section, we describe how to construct credible intervals based on (Kirschner et al., 2020) and (Inatsu et al., 2021).

3.1. Credible Interval

For each input $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$ and trial t , we define a credible interval of $f(\mathbf{x}, \mathbf{w})$ as $Q_t^{(f)}(\mathbf{x}, \mathbf{w}) = [l_t^{(f)}(\mathbf{x}, \mathbf{w}), u_t^{(f)}(\mathbf{x}, \mathbf{w})]$. Here, the lower bound $l_t^{(f)}(\mathbf{x}, \mathbf{w})$ and upper bound $u_t^{(f)}(\mathbf{x}, \mathbf{w})$ are given by

$$l_t^{(f)}(\mathbf{x}, \mathbf{w}) = \mu_t^{(f)}(\mathbf{x}, \mathbf{w}) - \beta_{f,t}^{1/2} \sigma_t^{(f)}(\mathbf{x}, \mathbf{w}),$$

$$u_t^{(f)}(\mathbf{x}, \mathbf{w}) = \mu_t^{(f)}(\mathbf{x}, \mathbf{w}) + \beta_{f,t}^{1/2} \sigma_t^{(f)}(\mathbf{x}, \mathbf{w}),$$

where $\beta_{f,t}^{1/2} \geq 0$. Similar to the same definition of $Q_t^{(f)}(\mathbf{x}, \mathbf{w})$, we define a credible interval of $g(\mathbf{x}, \mathbf{w})$ as $Q_t^{(g)}(\mathbf{x}, \mathbf{w}) = [l_t^{(g)}(\mathbf{x}, \mathbf{w}), u_t^{(g)}(\mathbf{x}, \mathbf{w})]$. Furthermore, we construct a credible interval of $\mathbb{1}[g(\mathbf{x}, \mathbf{w}) > h]$ by using $Q_t^{(g)}(\mathbf{x}, \mathbf{w})$. Let $\eta > 0$ be a user-specified overestimation parameter.¹ Then, we define the credible interval of

¹The parameter η is necessary to theoretical guarantees. Details are given in Section 4.

$\mathbb{1}[g(\mathbf{x}, \mathbf{w}) > h]$ as

$$Q_t^{(\mathbb{1})}(\mathbf{x}, \mathbf{w}; \eta) \equiv [l_t^{(\mathbb{1})}(\mathbf{x}, \mathbf{w}; \eta), u_t^{(\mathbb{1})}(\mathbf{x}, \mathbf{w}; \eta)] \\ = \begin{cases} [1, 1] & \text{if } l_t^{(g)}(\mathbf{x}, \mathbf{w}) > h - \eta, \\ [0, 1] & \text{if } l_t^{(g)}(\mathbf{x}, \mathbf{w}) \leq h - \eta \text{ and } u_t^{(g)}(\mathbf{x}, \mathbf{w}) > h, \\ [0, 0] & \text{if } l_t^{(g)}(\mathbf{x}, \mathbf{w}) \leq h - \eta \text{ and } u_t^{(g)}(\mathbf{x}, \mathbf{w}) \leq h. \end{cases}$$

Next, using $Q_t^{(f)}(\mathbf{x}, \mathbf{w})$, we define a credible interval of $F_t(\mathbf{x})$ as $Q_t^{(F_t)}(\mathbf{x}) \equiv [l_t^{(F_t)}(\mathbf{x}), u_t^{(F_t)}(\mathbf{x})]$, where $l_t^{(F_t)}(\mathbf{x})$ and $u_t^{(F_t)}(\mathbf{x})$ are calculated as

$$l_t^{(F_t)}(\mathbf{x}) = \inf_{p(\mathbf{w}) \in \mathcal{A}_t} \sum_{\mathbf{w} \in \Omega} l_t^{(f)}(\mathbf{x}, \mathbf{w}) p(\mathbf{w}), \\ u_t^{(F_t)}(\mathbf{x}) = \inf_{p(\mathbf{w}) \in \mathcal{A}_t} \sum_{\mathbf{w} \in \Omega} u_t^{(f)}(\mathbf{x}, \mathbf{w}) p(\mathbf{w}). \quad (5)$$

Note that if the distance function $d(\cdot, \cdot)$ is the $L1$ (or $L2$)-norm, (5) can be formulated as a linear (or second-order cone) programming problem. In both cases, optimization solvers exist to easily calculate (5). Similarly, we define a credible interval of $G_t(\mathbf{x})$ as $Q_t^{(G_t)}(\mathbf{x}; \eta) \equiv [l_t^{(G_t)}(\mathbf{x}; \eta), u_t^{(G_t)}(\mathbf{x}; \eta)]$, and its lower and upper bounds are given by

$$l_t^{(G_t)}(\mathbf{x}; \eta) = \inf_{p(\mathbf{w}) \in \mathcal{A}_t} \sum_{\mathbf{w} \in \Omega} l_t^{(\mathbb{1})}(\mathbf{x}, \mathbf{w}; \eta) p(\mathbf{w}), \\ u_t^{(G_t)}(\mathbf{x}; \eta) = \inf_{p(\mathbf{w}) \in \mathcal{A}_t} \sum_{\mathbf{w} \in \Omega} u_t^{(\mathbb{1})}(\mathbf{x}, \mathbf{w}; \eta) p(\mathbf{w}). \quad (6)$$

Moreover, using $Q_t^{(G_t)}(\mathbf{x}; \eta)$, we define an estimated upper (resp. lower) set H_t (resp. L_t) and a potential upper set M_t . Let $\xi > 0$ be a user-specified accuracy parameter. Then, we define H_t , L_t and M_t as

$$H_t = \{\mathbf{x} \in \mathcal{X} \mid l_t^{(G_t)}(\mathbf{x}; \eta) > \alpha - \xi\}, \\ L_t = \{\mathbf{x} \in \mathcal{X} \mid l_t^{(G_t)}(\mathbf{x}; \eta) \leq \alpha - \xi \text{ and } u_t^{(G_t)}(\mathbf{x}; \eta) \leq \alpha\}, \\ M_t = \{\mathbf{x} \in \mathcal{X} \mid l_t^{(G_t)}(\mathbf{x}; \eta) \leq \alpha - \xi \text{ and } u_t^{(G_t)}(\mathbf{x}; \eta) > \alpha\}.$$

3.2. Acquisition Function

We propose an AF to determine the next evaluation point. Our proposed AF is based on the following utility function:

$$\mathbb{E}[\max\{F_t(\mathbf{x}) - c_t^{(\text{best})}, 0\}] \times \mathbb{P}[G_t(\mathbf{x}) > \alpha]. \quad (7)$$

The first term is the expected improvement for $F_t(\mathbf{x})$, and the second term is the probability that the DR probability function $G_t(\mathbf{x})$ is greater than α . In the context of constrained BOs without environmental variables, this utility is known as the expected constrained improvement (ECI) (Gardner et al., 2014). Similarly, in the CCBO framework,

(Amri et al., 2021) proposed the expected feasible improvement (EFI) AF using the same utility. Unfortunately, in the DRCC setup, both the first and second terms cannot be calculated analytically because $F_t(\mathbf{x})$ and $G_t(\mathbf{x})$ do not follow GPs. Moreover, numerical approximation is also expensive because it requires a re-optimization calculation (inf operation) for all generated samples. For this reason, instead of (7), we consider a CI-based utility function which mimics (7).

First, we define the current best point $c_t^{(\text{best})}$ at trial t as

$$c_t^{(\text{best})} = \begin{cases} \max_{\mathbf{x} \in H_t} l_t^{(F_t)}(\mathbf{x}) & \text{if } H_t \neq \emptyset \\ \min_{\mathbf{x} \in M_t} l_t^{(F_t)}(\mathbf{x}) & \text{if } H_t = \emptyset \text{ and } M_t \neq \emptyset \\ \min_{\mathbf{x} \in \mathcal{X}} l_t^{(F_t)}(\mathbf{x}) & \text{if } H_t = \emptyset \text{ and } M_t = \emptyset \end{cases}.$$

Using this, we define the CI-based improvement $a_t^{(F_t)}(\mathbf{x})$ for F_t as

$$a_t^{(F_t)}(\mathbf{x}) = \max\{u_t^{(F_t)}(\mathbf{x}) - c_t^{(\text{best})}, 0\}.$$

Similarly, we define the CI-based probability $a_t^{(G_t)}(\mathbf{x})$ for $G_t > \alpha$ as

$$a_t^{(G_t)}(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} \in H_t \\ \frac{u_t^{(G_t)}(\mathbf{x}; \eta) - (\alpha - \xi)}{u_t^{(G_t)}(\mathbf{x}; \eta) - l_t^{(G_t)}(\mathbf{x}; \eta)} & \text{if } \mathbf{x} \in M_t \\ 0 & \text{if } \mathbf{x} \in L_t \end{cases}.$$

By combining these, we propose the following AF $a_t(\mathbf{x})$:

$$a_t(\mathbf{x}) = a_t^{(F_t)}(\mathbf{x}) \times a_t^{(G_t)}(\mathbf{x}).$$

Then, the next selected point is evaluated as follows:

Definition 3.1. The next design variable \mathbf{x}_{t+1} to evaluate is selected by

$$\mathbf{x}_{t+1} = \operatorname{argmax}_{\mathbf{x} \in H_t \cup M_t} a_t(\mathbf{x}).$$

On the other hand, unlike the uncontrollable setting, we also need to select \mathbf{w}_{t+1} in the simulator setting. One of the reasonable approaches is to focus on large posterior variances at the selected \mathbf{x}_{t+1} . Thus, we propose the following selection rule to evaluate \mathbf{w} :

Definition 3.2. The next environmental variable \mathbf{w}_{t+1} to evaluate is selected by

$$\mathbf{w}_{t+1} = \operatorname{argmax}_{\mathbf{w} \in \Omega} \{\sigma_t^{(f)2}(\mathbf{x}_{t+1}, \mathbf{w}) + \sigma_t^{(g)2}(\mathbf{x}_{t+1}, \mathbf{w})\}.$$

3.3. Stopping Conditions and DRCC-BO

We formulate the stopping condition of the proposed algorithm. If it is identified that the constraint $G_t(\mathbf{x}) > \alpha$ is not satisfied with high confidence for all the design variables \mathbf{x} ,

Algorithm 1 DRCC-BO: BO for DRCC problem

Require: GP priors $\mathcal{GP}(0, k^{(f)})$, $\mathcal{GP}(0, k^{(g)})$, threshold $h \in \mathbb{R}$, level $\alpha \in (0, 1)$, overestimation parameter $\eta > 0$, $\{\beta_{f,t}\}_{t \geq 1}, \{\beta_{g,t}\}_{t \geq 1}$, accuracy parameter $\xi > 0$, \mathcal{A}_t
 $H_0 \leftarrow \emptyset, L_0 \leftarrow \emptyset, M_0 \leftarrow \mathcal{X}, t \leftarrow 1$
while Both (S1) and (S2) are not satisfied **do**
 Compute $Q_t^{(F_t)}(\mathbf{x})$ and $Q_t^{(G_t)}(\mathbf{x}; \eta)$ for any $\mathbf{x} \in \mathcal{X}$
 Select the next evaluation point \mathbf{x}_t
 if Uncontrollable setting **then**
 Generate the next evaluation point $\mathbf{w}_t \sim P^\dagger$
 else
 Select the next evaluation point \mathbf{w}_t
 end if
 Observe $y_t^{(f)} = f(\mathbf{x}_t, \mathbf{w}_t) + \varepsilon_{f,t}$ and $y_t^{(g)} = g(\mathbf{x}_t, \mathbf{w}_t) + \varepsilon_{g,t}$
 Update the GPs by adding observations, and compute H_t, L_t, M_t
 $t \leftarrow t + 1$
end while
Ensure: No solution exists if (S1) is satisfied, and otherwise $\hat{\mathbf{x}}_t = \operatorname{argmax}_{\mathbf{x} \in H_{t-1}} l_{t-1}^{(F_t)}(\mathbf{x})$

then the algorithm should be stopped because there is no solution. Alternatively, the algorithm should also be stopped if the difference between the conservative maximum for $F_t(\mathbf{x})$ in the points that satisfy the constraint and the optimistic maximum for $F_t(\mathbf{x})$ in the points that may satisfy the constraint is sufficiently small. Based on these ideas, we define the following two stopping conditions:

(S1) $L_t = \mathcal{X}$.

(S2) $H_t \neq \emptyset$ and $\max_{\mathbf{x} \in H_t \cup M_t} u_t^{(F_t)}(\mathbf{x}) - \max_{\mathbf{x} \in H_t} l_t^{(F_t)}(\mathbf{x}) < \xi$.

The pseudocode of the proposed method in uncontrollable and simulator settings are given in Algorithm 1.

4. Theoretical Analysis

In this section, we show the theoretical guarantee on the accuracy and convergence in our proposed algorithm. First, we assume that the true black-box functions f and g follow GPs $\mathcal{GP}(0, k^{(f)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')))$, and $\mathcal{GP}(0, k^{(g)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')))$, respectively. Moreover, as a technical condition, we assume that the posterior variances $k^{(f)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}, \mathbf{w})) \equiv \sigma_0^{(f)2}(\mathbf{x}, \mathbf{w})$ and $k^{(g)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}, \mathbf{w})) \equiv \sigma_0^{(g)2}(\mathbf{x}, \mathbf{w})$ satisfy

$$\max_{(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega} \sigma_0^{(f)2}(\mathbf{x}, \mathbf{w}) \leq 1, \quad \max_{(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega} \sigma_0^{(g)2}(\mathbf{x}, \mathbf{w}) \leq 1, \\
 0 < \sigma_{0, \min}^{(g)2} \equiv \min_{(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega} \sigma_0^{(g)2}(\mathbf{x}, \mathbf{w}).$$

Here, let $\kappa_T^{(f)}, \kappa_T^{(g)}$ be the maximum information gain of f and g at trial T , respectively. Note that the maximum information gain is a commonly used complexity measure in the context of the GP-based BO method (see, e.g., (Srinivas et al., 2010)). The value $\kappa_T^{(f)}$ can be expressed as $\kappa_T^{(f)} = \max_{S \subset \mathcal{X} \times \Omega} I(\mathbf{y}_S^{(f)}; f)$, where $I(\mathbf{y}^{(f)}; f)$ is the mutual information between $\mathbf{y}^{(f)}$ and f . Similarly, $\kappa_T^{(g)}$ can be expressed by using $I(\mathbf{y}^{(g)}; g)$. Next, we define the goodness of the estimated $\hat{\mathbf{x}}_t$ as follows:

Definition 4.1. For a given positive constant C , we define the estimated solution $\hat{\mathbf{x}}_t$ as the C -accurate solution if $\hat{\mathbf{x}}_t$ satisfies the following inequalities:

$$F_t(\mathbf{x}_t^*) - F_t(\hat{\mathbf{x}}_t) < C, \quad G_t(\hat{\mathbf{x}}_t) > \alpha - C.$$

Then, the following theorem holds:

Theorem 4.2. Let $h \in \mathbb{R}$, $\alpha \in (0, 1)$, $t \geq 1$, $\delta \in (0, 1)$, and define $\beta_{f,t} = \beta_{g,t} = 2 \log(2|\mathcal{X} \times \Omega| \pi^2 t^2 / (3\delta)) \equiv \beta_t$. For a user-specified accuracy parameter $\xi > 0$, define an overestimation parameter $\eta > 0$ as

$$\eta = \min \left\{ \frac{\xi \sigma_{0, \min}^{(g)}}{2}, \frac{\xi^2 \delta \sigma_{0, \min}^{(g)}}{8|\mathcal{X} \times \Omega|} \right\}.$$

Then, when Algorithm 1 is performed, with a probability of at least $1 - \delta$, the following holds for any t and \mathcal{A}_t :

- If (S1) is satisfied, then $G_t(\mathbf{x}) \leq \alpha$ for all $\mathbf{x} \in \mathcal{X}$, that is, the DRCC problem has no solution.
- If (S2) is satisfied, then $\hat{\mathbf{x}}_t$ is the 2ξ -accurate solution.

Moreover, these results do not depend on whether the simulator or uncontrollable setting is used.

We would like to note that although Theorem 4.2 guarantees the returned solution by Algorithm 1 is good, but does not state whether the stopping conditions are satisfied or not. The sufficient conditions for stopping conditions to be satisfied need to be considered for the simulator and uncontrollable settings, separately. First, we give the sufficient conditions in the simulator setting.

Theorem 4.3. Under the same condition as in Theorem 4.2, let T be the smallest positive integer satisfying the following inequalities:

$$\frac{\beta_{T^2}(C_{1,f} \kappa_T^{(f)} + C_{1,g} \kappa_T^{(g)})}{T} < \min \left\{ \frac{\xi^4}{4}, \frac{\eta^2}{4} \right\}, \quad (8)$$

where $C_{1,f} = 2/\log(1 + \sigma_{f, \text{noise}}^{-2})$ and $C_{1,g} = 2/\log(1 + \sigma_{g, \text{noise}}^{-2})$. Then, under the simulator setting, Algorithm 1 terminates after at most T^2 trials.

Next, we give the sufficient conditions in the uncontrollable setting. In the simulator setting, we can choose any \mathbf{w} and thus the uncertainty of f and g can be reduced sufficiently. In contrast, in the uncontrollable setting, we cannot choose \mathbf{w} , freely. For this reason, it is desirable to be able to make the uncertainty of all points small in probability. However, if $p^\dagger(\mathbf{w}) = 0$ for some \mathbf{w} , the uncertainty at points including this point is not reduced sufficiently. To avoid this problem, in the uncontrollable setting, we assume that the true distribution satisfies

$$\min_{\mathbf{w} \in \Omega} p^\dagger(\mathbf{w}) \equiv p_{\min} > 0.$$

Then, the following theorem holds:

Theorem 4.4. *Under the same condition as in Theorem 4.2, assume that $p_{\min} > 0$. Let T be the smallest positive integer satisfying the following inequalities:*

$$\frac{\beta_{T^2}(C_{2,f}\kappa_T^{(f)} + C_{2,g}\kappa_T^{(g)} + C_3)}{T} < \min \left\{ \frac{\xi^4}{4}, \frac{\eta^2}{4} \right\}, \quad (9)$$

where $C_{2,f} = (4p_{\min}^{-1})/\log(1 + \sigma_{f,\text{noise}}^{-2})$, $C_{2,g} = (4p_{\min}^{-1})/\log(1 + \sigma_{g,\text{noise}}^{-2})$ and $C_3 = 16p_{\min}^{-1}\log(5/\delta)$. Then, under the uncontrollable setting, with a probability of at least $1 - \delta$, Algorithm 1 terminates after at most T^2 trials.

Furthermore, we give a theorem that it is possible to link the DRCC problem to the CC problem by choosing \mathcal{A}_t appropriately. Specifically, it ensures that the solution to the DRCC problem is also a good solution to the CC problem. Here, we consider the following CC problem

$$\begin{aligned} \tilde{F}(\mathbf{x}) &= \sum_{\mathbf{w} \in \Omega} f(\mathbf{x}, \mathbf{w})p^\dagger(\mathbf{w}), \\ \tilde{G}(\mathbf{x}) &= \sum_{\mathbf{w} \in \Omega} \mathbb{1}[g(\mathbf{x}, \mathbf{w}) > h]p^\dagger(\mathbf{w}), \\ \tilde{\mathbf{x}}^* &= \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \tilde{F}(\mathbf{x}) \quad \text{subject to} \quad \tilde{G}(\mathbf{x}) > \alpha, \end{aligned}$$

where if the optimal solution $\tilde{\mathbf{x}}^*$ does not exist, it is formally defined as $\tilde{F}(\tilde{\mathbf{x}}^*) = \min_{\mathbf{x} \in \mathcal{X}} \tilde{F}(\mathbf{x})$. As with the DRCC problem, we define the goodness of the solution to the CC problem as follows:

Definition 4.5. For a given positive constant C , we define the solution $\hat{\mathbf{x}}_t$ as the C -accurate solution to the CC problem if $\hat{\mathbf{x}}_t$ satisfies the following inequalities:

$$\tilde{F}(\hat{\mathbf{x}}_t) - \tilde{F}(\tilde{\mathbf{x}}^*) < C, \quad \tilde{G}(\hat{\mathbf{x}}_t) > \alpha - C.$$

Then, the following theorem holds:

Theorem 4.6. *Under the uncontrollable setting, let $h \in \mathbb{R}$, $\alpha \in (0, 1)$, $t \geq 1$, $\delta \in (0, 1)$, and define $\beta_{f,t} = \beta_{g,t} =$*

$2 \log(2|\mathcal{X} \times \Omega| \pi^2 t^2 / (3\delta)) \equiv \beta_t$. For a user-specified accuracy parameter $\xi > 0$, define $\alpha' = \alpha - \xi$ and an overestimation parameter $\eta > 0$ as

$$\eta = \min \left\{ \frac{\xi \sigma_{0,\min}^{(g)}}{2}, \frac{\xi^2 \delta \sigma_{0,\min}^{(g)}}{8|\mathcal{X} \times \Omega|} \right\}.$$

Furthermore, let $p_t^*(\mathbf{w})$ be an empirical distribution of \mathbf{w} , and let

$$\epsilon_t = |\Omega| \sqrt{\frac{1}{2t} \log \left(\frac{|\Omega| \pi^2 t^2}{3\delta} \right)},$$

$$d(p_1(\mathbf{w}), p_2(\mathbf{w})) = \sum_{\mathbf{w} \in \Omega} |p_1(\mathbf{w}) - p_2(\mathbf{w})|.$$

Then, when Algorithm 1 is performed by using α' , with a probability of at least $1 - 2\delta$, the following holds for any $t \geq T$:

- If (S1) is satisfied, then $G(\mathbf{x}) \leq \alpha$ for all $\mathbf{x} \in \mathcal{X}$, that is, the CC problem has no solution,
- If (S2) is satisfied, then $\hat{\mathbf{x}}_t$ is the 3ξ -accurate solution for the CC problem,

where T is the smallest positive integer satisfying

$$\forall n \geq T, 2(1 + \beta_{f,1}^{1/2})\epsilon_n < \xi.$$

Finally, the results of the theorems obtained in this section are summarized in Table 1.

Note that the order of maximum information gains $\kappa_T^{(f)}$ and $\kappa_T^{(g)}$ is known to be sublinear under mild conditions (Srinivas et al., 2010). Therefore, noting that the order of $\beta_{f,T} = \beta_{g,T}$ is $O(\log T)$, the positive integer T satisfying (8) and (9) exists.

5. Numerical Experiments

In this section, we confirm the performance of the proposed method in simulator and uncontrollable settings using synthetic functions and real-world simulations. In this experiment, both design and environment variables were set to one dimension, and the following Gaussian kernels were used as the kernel functions:

$$\begin{aligned} k^{(f)}((x, w), (x', w')) &= \sigma_{f,\ker}^2 \exp(-\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2 / L_f), \\ k^{(g)}((x, w), (x', w')) &= \sigma_{g,\ker}^2 \exp(-\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2 / L_g), \end{aligned}$$

where $\boldsymbol{\theta} = (x, w)$. We used the $L1$ -norm as the distance between distributions, and set $\epsilon_t = 0.15$. Here, for simplicity, we set the overestimation parameter η to 0 and the accuracy

Table 1. Probabilities of accuracy and algorithm termination in each setting

	Simulator	Uncontrollable	Uncontrollable using Theorem 4.6 setting
2ξ-accuracy for DRCC	1 - δ	1 - δ	1 - δ
3ξ-accuracy for CC	NA	NA	1 - 2δ
Algorithm termination	1	1 - δ	1 - δ
2ξ-accuracy for DRCC and termination	1 - δ	1 - 2δ	1 - 2δ
3ξ-accuracy for CC and termination	NA	NA	1 - 3δ

parameter to $\xi = 10^{-12}$. In all experiments, we evaluated the performance of each algorithm using the following utility gap UG_t :

$$UG_t = \begin{cases} F_t(x_t^*) - F_t(\hat{x}_t) & \text{if } H_t \neq \emptyset \text{ and } G_t(\hat{x}_t) > \alpha \\ F_t(x_t^*) - \min_{x \in \mathcal{X}} F_t(x) & \text{otherwise,} \end{cases}$$

where \hat{x}_t is given by $\hat{x}_t = \operatorname{argmax}_{x \in H_t} l_t^{(F_t)}(x)$. The details of the experimental setting, which are not included in the main body, are described in Appendix B.

5.1. Synthetic Function

We evaluate the performance of the proposed method using a synthetic function. We used the input space $\mathcal{X} \times \Omega$ as a set of grid points divided by $[-10, 10] \times [-10, 10]$ into 50×50 equally spaced. Moreover, we used the following black-box functions f and g :

$$\begin{aligned} f(x, w) &= \exp(-x^2/4) + 0.6 \exp(-(x-8)^2/3) \\ &\quad + 0.3 \exp(-(x+9)^2/5) + \exp(-w^2/4) \\ &\quad + 0.6 \exp(-(w-8)^2/3) + 0.3 \exp(-(w+9)^2/5), \\ g(x, w) &= 0.26(x^2 + w^2) - 0.48xw. \end{aligned}$$

In this experiment, we performed a total of three different experiments, one with the simulator setting and two with the uncontrollable setting:

Simulator: Under the simulator setting, $p_t^*(\mathbf{w}) = 1/50$ was used as the reference distribution.

Fixed: Under the uncontrollable setting, the mixture normal distribution $0.5\mathcal{N}(-5, 10) + 0.5\mathcal{N}(5, 10)$ discretized on Ω was used as the true distribution $p^\dagger(\mathbf{w})$. The reference distribution was set to $p_t^*(\mathbf{w}) = 1/50$.

Data-driven: Under the uncontrollable setting, for the true distribution $p^\dagger(\mathbf{w})$ we used the same as Fixed, and for the reference distribution we used the empirical distribution function of \mathbf{w} .

We compared the following seven methods:

Random: Select (x_{t+1}, w_{t+1}) randomly.

US: Select (x_{t+1}, w_{t+1}) by maximizing the maximum posterior variance of f and g , i.e., (x_{t+1}, w_{t+1}) is given by

$$(x_{t+1}, w_{t+1}) = \operatorname{argmax}_{(x, w) \in \mathcal{X} \times \Omega} \text{US}_t(x, w),$$

where $\text{US}_t(x, w) = \max\{\sigma_t^{(f)2}(x, w), \sigma_t^{(g)2}(x, w)\}$.

DRBO: Use the DRBO method proposed by (Kirschner et al., 2020), i.e., x_{t+1} and w_{t+1} are selected by $x_{t+1} = \operatorname{argmax}_{x \in \mathcal{X}} u_t^{(F_t)}(x)$ and $w_{t+1} = \operatorname{argmax}_{w \in \Omega} \sigma_t^{(f)2}(x_{t+1}, w)$.

DRPTR: Use the DRPTR method proposed by (Inatsu et al., 2021), i.e., (x_{t+1}, w_{t+1}) is selected by $(x_{t+1}, w_{t+1}) = \operatorname{argmax}_{(x, w) \in \mathcal{X} \times \Omega} a_t^{(2)}(\mathbf{x}, \mathbf{w})$, where $a_t^{(2)}(\mathbf{x}, \mathbf{w})$ is given by Definition 3.2 of (Inatsu et al., 2021).

CCBO: Use the CCBO method proposed by (Amri et al., 2021), i.e., x_{t+1} and w_{t+1} are selected by $x_{t+1} = \operatorname{argmax}_{x \in \mathcal{X}} \text{EFI}(x)$ and $w_{t+1} = \operatorname{argmax}_{w \in \Omega} S(x_{t+1}, w)$, where $\text{EFI}(x)$ and $S(x, w)$ are given by (7) and (13) of (Amri et al., 2021).

MVA: Use a modified version of the MVA-BO method proposed by (Iwazaki et al., 2021b), i.e., x_{t+1} is selected by

$$x_{t+1} = \operatorname{argmax}_{x \in P_t} \sqrt{\text{length}_t^{(F_t)}(x)^2 + \text{length}_t^{(G_t)}(x; \eta)^2},$$

$$P_t = (H_t \cup M_t) \cap \{x \in \mathcal{X} \mid u_t^{(F_t)}(x) > \max_{x \in H_t} l_t^{(F_t)}(x)\},$$

$$\text{length}_t^{(F_t)}(x) = u_t^{(F_t)}(x) - l_t^{(F_t)}(x),$$

$$\text{length}_t^{(G_t)}(x) = u_t^{(G_t)}(x; \eta) - l_t^{(G_t)}(x; \eta)$$

and w_{t+1} is selected by Definition 3.2.

Proposed: Use Definition 3.1-3.2.

In the case of uncontrollable setting, we selected only x_{t+1} . On the other hand, because US and DRPTR select x and w simultaneously, we modified them in the uncontrollable setting as follows:

US: $x_{t+1} = \operatorname{argmax}_{x \in \mathcal{X}} \mathbb{E}_w[\text{US}_t(x, w)]$.

DRPTR: $x_{t+1} = \operatorname{argmax}_{x \in \mathcal{X}} \mathbb{E}_w[a_t^{(2)}(x, w)]$.

Here, the expectation is taken with respect to the empirical distribution of w . We would like to emphasize that DRBO focuses only on the maximization of $F_t(x)$, and does not consider whether the constraints are satisfied or not. In contrast, DRPTR focuses only on the identification of variables that satisfy the constraints and does not consider the maximization of $F_t(x)$. As for CCBO, it is the BO method for the CC problem (1a)–(1b), and does not consider the distributionally robustness. Similarly, MVA focuses on only one black-box function $f(x, w)$ and deals the constraint optimization problem between the mean and variance of $f(x, w)$ with respect to w . Therefore, MVA does not consider the (distributionally robust) chance-constraint problem.

With this setup, we took one initial point at random and ran the algorithms until the number of iterations reached 300. The simulation was repeated 100 times and the average value of the utility gap at each iteration was calculated. From Figure 1, it can be confirmed that the proposed method shows high performance.

5.2. Infection Simulation

We then applied the proposed method to the decision-making problem for simulation-based infectious diseases in the real world. Here, we used the SIR model (Kermack & McKendrick, 1927), which is a commonly used model to describe the behavior of infection. The SIR model uses the contact rate $\beta \in [0, 1]$ and the isolation rate $\gamma \in [0, 1]$ to model the behavior of infection over time. In this experiment, we considered the grid points that divide the interval $[0.01, 0.5]$ into 50 equal parts as β and γ , and used them as input. Based on the SIR model, we defined the following two risk functions:

$$\begin{aligned} R_1(\beta, \gamma) &= n_{\text{infected}}(\beta, \gamma) - 450\beta + 800\gamma - C_1, \\ R_2(\beta, \gamma) &= n_{\text{infected}}(\beta, \gamma) - C_2, \end{aligned}$$

where $n_{\text{infected}}(\beta, \gamma)$, which is calculated by using the SIR model with (β, γ) , is the maximum number of infected within a given period. In addition, C_1 (resp. C_2) is a shift constant to match the absolute values of the maximum and minimum of $R_1(\beta, \gamma)$ (resp. $R_2(\beta, \gamma)$). While $R_2(\beta, \gamma)$, which represents the number of infected people, is an intuitive risk function, $R_1(\beta, \gamma)$ can be interpreted as an economic risk function. In fact, as the number of infected people increases, the economic risk increases. In addition, if the contact rate is large, that is, if freedom of action is not restricted, economic activity will not stagnate and the risk will be small. On the other hand, if the isolation rate is

large, economic activity will stagnate and the risk will increase. Although these risk functions should be minimized, we multiplied them by minus one in our experiments in order to match the setting of this paper. Also, $R_1(\beta, \gamma)$ can be interpreted as both an objective function and a constraint function, and the same is true for $R_2(\beta, \gamma)$. Similarly, the contact rate β can be interpreted as both a design variable and an environmental variable, and the same is true for γ . For these reasons, we performed the following experiments:

Case1: Design variable x : β , environmental variable w : γ ,
 $f(x, w) = -R_1(x, w)$, $g(x, w) = -R_2(x, w)$.

Case2: Design variable x : β , environmental variable w : γ ,
 $f(x, w) = -R_2(x, w)$, $g(x, w) = -R_1(x, w)$.

Case3: Design variable x : γ , environmental variable w : β ,
 $f(x, w) = -R_1(x, w)$, $g(x, w) = -R_2(x, w)$.

Case4: Design variable x : γ , environmental variable w : β ,
 $f(x, w) = -R_2(x, w)$, $g(x, w) = -R_1(x, w)$.

In all experiments, the simulator setting was considered, and $p_t^*(w) = 1/50$ was used as the reference distribution.

With this setup, we took one initial point at random and ran the algorithms until the number of iterations reached 100. The simulation was repeated 100 times and the average value of the utility gap at each iteration was calculated. From Figure 2, it can be confirmed that the proposed method performs as well as or better than the comparison methods.

5.3. Hyperparameter Sensitivity

We evaluate the sensitivity of $\beta_{f,t}^{1/2}$ and $\beta_{g,t}^{1/2}$. We used the same setting as in §5.1 except for $\beta_{f,t}^{1/2}$ and $\beta_{g,t}^{1/2}$, and evaluated only the proposed method. We considered the 16 cases of $\beta_{f,t}^{1/2} = \beta_f^{1/2} \in \{0, 1, 2, 3\}$ and $\beta_{g,t}^{1/2} = \beta_g^{1/2} \in \{0, 1, 2, 3\}$. From Figure 3, it can be confirmed that utility gaps do not become zero when $\beta_f^{1/2}$ or $\beta_g^{1/2}$ is zero. On the other hand, when $\beta_f^{1/2}$ and $\beta_g^{1/2}$ are non-zero, the best performance is achieved when $\beta_f^{1/2} = 2$ and $\beta_g^{1/2} = 1$ in both settings. Thus, it can be confirmed that good performance can actually be achieved without the large $\beta_{f,t}^{1/2}$ and $\beta_{g,t}^{1/2}$ values for theoretical guarantees given in §4. Experimental results of the uncontrollable setting are described in Appendix.

6. Conclusion

In this paper, we proposed the BO method for efficiently finding the optimal solution to the DRCC problem for the simulator and uncontrollable settings. We showed that the proposed method can return an arbitrary accurate solution

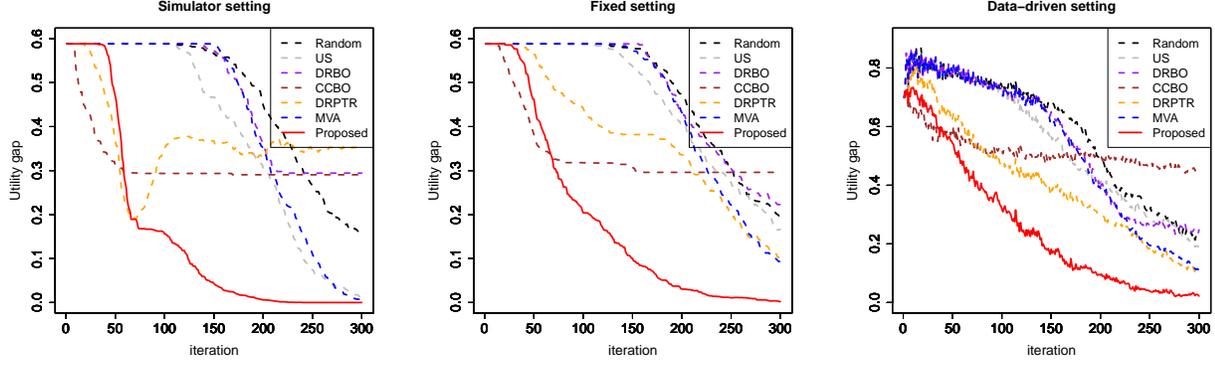


Figure 1. Average utility gap for each method in simulator and uncontrollable settings.

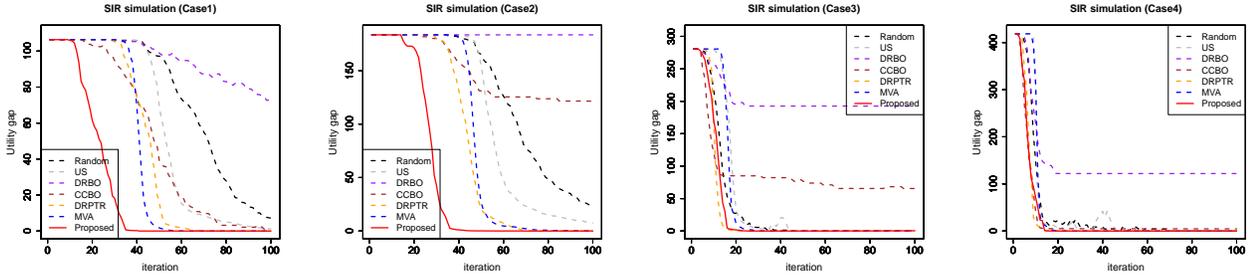


Figure 2. Average utility gap for SIR simulation experiments in the simulator setting.

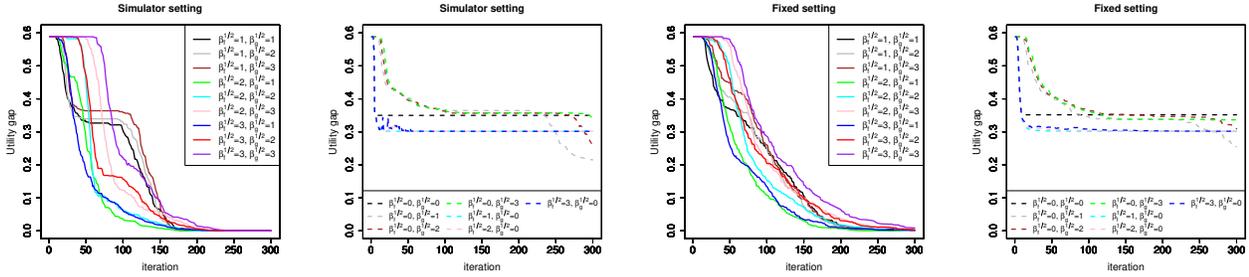


Figure 3. Hyperparameter sensitivity for the proposed method in simulator and fixed settings.

with high probability in a finite number of trials. Furthermore, through numerical experiments, we confirmed that the performance of the proposed method is superior to other comparison methods.

Acknowledgements

This work was partially supported by MEXT KAKENHI (22H00300, 21H03498, 20H00601, 17H04694, 16H06538), JSPS KAKENHI (JP21J14673), JST CREST (JPMJCR21D3), JST Moonshot R&D (JPMJMS2033-05), JST AIP Acceleration Research (JPMJCR21U2), NEDO (JPNP18002, JPNP20006) and RIKEN Center for Advanced

Intelligence Project.

References

- Amri, R. E., Riche, R. L., Helbert, C., Blanchet-Scalliet, C., and Da Veiga, S. A sampling criterion for constrained bayesian optimization with uncertainties. *arXiv preprint arXiv:2103.05706*, 2021.
- Bogunovic, I., Scarlett, J., Jegelka, S., and Cevher, V. Adversarially robust optimization with Gaussian processes. In *Advances in neural information processing systems*, pp. 5760–5770, 2018.

- Bryan, B., Nichol, R. C., Genovese, C. R., Schneider, J., Miller, C. J., and Wasserman, L. Active learning for identifying function threshold boundaries. In *Advances in neural information processing systems*, pp. 163–170, 2006.
- Fang, X., Hodge, B.-M., Li, F., Du, E., and Kang, C. Adjustable and distributionally robust chance-constrained economic dispatch considering wind power uncertainty. *Journal of Modern Power Systems and Clean Energy*, 7(3):658–664, 2019.
- Gardner, J., Kusner, M., Weinberger, K., Cunningham, J., et al. Bayesian optimization with inequality constraints. In *International Conference on Machine Learning*, pp. 937–945. PMLR, 2014.
- Gotovos, A., Casati, N., Hitz, G., and Krause, A. Active learning for level set estimation. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pp. 1344–1350. AAAI Press, 2013. ISBN 978-1-57735-633-2. URL <http://dl.acm.org/citation.cfm?id=2540128.2540322>.
- 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*, 17:1–53, 2016.
- Ho-Nguyen, N., Kılınc-Karzan, F., Küçükyavuz, S., and Lee, D. Distributionally robust chance-constrained programs with right-hand side uncertainty under wasserstein ambiguity. *Mathematical Programming*, pp. 1–32, 2021.
- Inatsu, Y., Karasuyama, M., Inoue, K., Kandori, H., and Takeuchi, I. Active learning of Bayesian linear models with high-dimensional binary features by parameter confidence-region estimation. *Neural Computation*, 32(10):1998–2031, 2020a.
- Inatsu, Y., Karasuyama, M., Inoue, K., and Takeuchi, I. Active learning for level set estimation under input uncertainty and its extensions. *Neural Computation*, 32(12):2486–2531, 2020b.
- Inatsu, Y., Iwazaki, S., and Takeuchi, I. Active learning for distributionally robust level-set estimation. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 4574–4584. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/inatsu21a.html>.
- Iwazaki, S., Inatsu, Y., and Takeuchi, I. Bayesian experimental design for finding reliable level set under input uncertainty. *IEEE Access*, 8:203982–203993, 2020.
- Iwazaki, S., Inatsu, Y., and Takeuchi, I. Bayesian Quadrature Optimization for Probability Threshold Robustness Measure. *Neural Computation*, 33(12):3413–3466, 11 2021a. ISSN 0899-7667. doi: 10.1162/neco_a_01442. URL https://doi.org/10.1162/neco_a_01442.
- Iwazaki, S., Inatsu, Y., and Takeuchi, I. Mean-variance analysis in bayesian optimization under uncertainty. In Banerjee, A. and Fukumizu, K. (eds.), *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 973–981. PMLR, 13–15 Apr 2021b.
- Kermack, W. O. and McKendrick, A. G. A contribution to the mathematical theory of epidemics. *Proceedings of the royal society of london. Series A, Containing papers of a mathematical and physical character*, 115(772):700–721, 1927.
- Kirschner, J. and Krause, A. Information directed sampling and bandits with heteroscedastic noise. In *Conference On Learning Theory*, pp. 358–384. PMLR, 2018.
- Kirschner, J., Bogunovic, I., Jegelka, S., and Krause, A. Distributionally robust Bayesian optimization. In Chiappa, S. and Calandra, R. (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2174–2184. PMLR, 26–28 Aug 2020.
- Nguyen, Q. P., Dai, Z., Low, B. K. H., and Jaillet, P. Optimizing conditional value-at-risk of black-box functions. *Advances in Neural Information Processing Systems*, 34, 2021a.
- Nguyen, Q. P., Dai, Z., Low, B. K. H., and Jaillet, P. Value-at-risk optimization with gaussian processes. In Meila, M. and Zhang, T. (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8063–8072. PMLR, 18–24 Jul 2021b. URL <https://proceedings.mlr.press/v139/nguyen21b.html>.
- Nguyen, T., Gupta, S., Ha, H., Rana, S., and Venkatesh, S. Distributionally robust Bayesian quadrature optimization. In Chiappa, S. and Calandra, R. (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1921–1931. PMLR, 26–28 Aug 2020.

- Rahimian, H. and Mehrotra, S. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.
- Scarf, H. A min-max solution of an inventory problem. *Studies in the mathematical theory of inventory and production*, 10:201–209, 1958.
- Settles, B. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- Srinivas, N., Krause, A., Kakade, S., and Seeger, M. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on International Conference on Machine Learning, ICML'10*, pp. 1015–1022, USA, 2010. Omnipress. ISBN 978-1-60558-907-7. URL <http://dl.acm.org/citation.cfm?id=3104322.3104451>.
- Sui, Y., Gotovos, A., Burdick, J., and Krause, A. Safe exploration for optimization with Gaussian processes. In *International Conference on Machine Learning*, pp. 997–1005, 2015.
- Sui, Y., Burdick, J., Yue, Y., et al. Stagesafe Bayesian optimization with Gaussian processes. In *International Conference on Machine Learning*, pp. 4781–4789, 2018.
- Turchetta, M., Berkenkamp, F., and Krause, A. Safe exploration in finite Markov decision processes with Gaussian processes. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 4312–4320, 2016.
- Wachi, A., Sui, Y., Yue, Y., and Ono, M. Safe exploration and optimization of constrained MDPs using Gaussian processes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Williams, C. K. and Rasmussen, C. E. Gaussian processes for machine learning. *the MIT Press*, 2(3):4, 2006.
- Xie, W. On distributionally robust chance constrained programs with wasserstein distance. *Mathematical Programming*, 186(1):115–155, 2021.
- Xie, W. and Ahmed, S. Distributionally robust chance constrained optimal power flow with renewables: A conic reformulation. *IEEE Transactions on Power Systems*, 33(2):1860–1867, 2017.
- Zanette, A., Zhang, J., and Kochenderfer, M. J. Robust super-level set estimation using Gaussian processes. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 276–291. Springer, 2018.

A. Proofs

A.1. Proof of Theorem 4.2

From the proof of Theorem 4.1 in (Inatsu et al., 2021), with a probability of at least $1 - 3\delta/4$ the following holds for any $\mathbf{x} \in \mathcal{X}$ and $t \geq 1^2$:

$$G_t(\mathbf{x}) \leq u_t^{(G_t)}(\mathbf{x}; \eta), \quad G_t(\mathbf{x}) \geq l_t^{(G_t)}(\mathbf{x}; \eta) - \xi.$$

Here, if the stopping condition (S1) holds, then $u_t^{(G_t)}(\mathbf{x}; \eta) \leq \alpha$ for any \mathbf{x} . By combining this and $G_t(\mathbf{x}) \leq u_t^{(G_t)}(\mathbf{x}; \eta)$, we have $G_t(\mathbf{x}) \leq \alpha$. This implies that the DRCC problem has no solution. On the other hand, if the stopping condition (S2) holds, $\hat{\mathbf{x}}_t$ satisfies that

$$l_t^{(G_t)}(\mathbf{x}; \eta) > \alpha - \xi.$$

By using this and $G_t(\mathbf{x}) \geq l_t^{(G_t)}(\mathbf{x}; \eta) - \xi$, we obtain $G_t(\hat{\mathbf{x}}_t) \geq \alpha - 2\xi$. Here, if the optimal solution \mathbf{x}_t^* does not exist, from the definition it follows that $F_t(\mathbf{x}_t^*) - F_t(\hat{\mathbf{x}}_t) \leq 0 < \xi < 2\xi$. Therefore, \mathbf{x}_t^* is a 2ξ -accurate solution. Next, we consider the case where the optimal solution \mathbf{x}_t^* exists. From Lemma 5.1 in (Srinivas et al., 2010), under the assumption on Theorem 4.2, with a probability of at least $1 - \delta/4$ the following inequality holds for any $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$ and $t \geq 1$:

$$f(\mathbf{x}, \mathbf{w}) \in Q_t^{(f)}(\mathbf{x}, \mathbf{w}).$$

Hence, it follows that $F_t(\mathbf{x}) \in Q_t^{(F_t)}(\mathbf{x})$. Moreover, because \mathbf{x}_t^* satisfies $G_t(\mathbf{x}_t^*) > \alpha$, then $G_t(\mathbf{x}_t^*) \leq u_t^{(G_t)}(\mathbf{x}_t^*; \eta)$ with a probability of at least $1 - 3\delta/4$. Thus, we get $\mathbf{x}_t^* \in H_t \cup M_t$. Hence, the following holds:

$$F_t(\mathbf{x}_t^*) \leq u_t^{(F_t)}(\mathbf{x}_t^*) \leq \max_{\mathbf{x} \in H_t \cup M_t} u_t^{(F_t)}(\mathbf{x}).$$

Similarly, noting that $F_t(\hat{\mathbf{x}}_t) \geq l_t^{(F_t)}(\hat{\mathbf{x}}_t) = \max_{\mathbf{x} \in H_t} l_t^{(F_t)}(\mathbf{x})$, from the stopping condition (S2) it follows that

$$2\xi > \xi > \max_{\mathbf{x} \in H_t \cup M_t} u_t^{(F_t)}(\mathbf{x}) - \max_{\mathbf{x} \in H_t} l_t^{(F_t)}(\mathbf{x}) \geq F_t(\mathbf{x}_t^*) - F_t(\hat{\mathbf{x}}_t).$$

Therefore, $\hat{\mathbf{x}}_t$ is a 2ξ -accurate solution.

A.2. Proof of Theorem 4.3

Let T be the smallest positive integer satisfying (8). Also let $(\mathbf{x}_1, \mathbf{w}_1), \dots, (\mathbf{x}_{T^2}, \mathbf{w}_{T^2})$ be points selected by the algorithm. Here, one of the following holds for T^2 :

Case1 There exists a positive integer $t \leq T^2$ such that $L_t = \mathcal{X}$.

Case2 For any positive integer $t \leq T^2$, $L_t \neq \mathcal{X}$.

If Case1 holds, then from the stopping condition (S1) the algorithm terminates. Next, we consider Case2. Let

$$T_1 = \operatorname{argmin}_{1 \leq t \leq T} \{\sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w}_t) + \sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w}_t)\}.$$

Then, the following inequality holds:

$$\begin{aligned} T\{\sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w}_{T_1}) + \sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w}_{T_1})\} &\leq \sum_{t=1}^T \sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w}_t) + \sum_{t=1}^T \sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w}_t) \\ &\leq C_{1,f} \kappa_T^{(f)} + C_{1,g} \kappa_T^{(g)}, \end{aligned}$$

where the last inequality can be derived from Lemma 5.3 and 5.4 in (Srinivas et al., 2010). Thus, it follows that

$$\sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w}_{T_1}) + \sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w}_{T_1}) \leq \frac{C_{1,f} \kappa_T^{(f)} + C_{1,g} \kappa_T^{(g)}}{T}.$$

²They only consider the fixed candidate family \mathcal{A} , but the same argument also holds in the case of $\mathcal{A}_t \neq \mathcal{A}_t$.

In addition, from Definition 3.2, the following holds for any $\mathbf{w} \in \Omega$:

$$\sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w}) \leq \max_{\mathbf{w} \in \Omega} (\sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w}) + \sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w})) = \sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w}_{T_1}) + \sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w}_{T_1}) \leq \frac{C_{1,f}\kappa_T^{(f)} + C_{1,g}\kappa_T^{(g)}}{T}.$$

Hence, we have

$$\beta_{g,T_1} \sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w}) \leq \frac{\beta_{g,T_1}(C_{1,f}\kappa_T^{(f)} + C_{1,g}\kappa_T^{(g)})}{T} \leq \frac{\beta_{g,T^2}(C_{1,f}\kappa_T^{(f)} + C_{1,g}\kappa_T^{(g)})}{T}.$$

Here, from the theorem's assumption, it holds that

$$\frac{\beta_{g,T^2}(C_{1,f}\kappa_T^{(f)} + C_{1,g}\kappa_T^{(g)})}{T} < \eta^2/4.$$

Therefore, we have $\beta_{g,T_1}^{1/2} \sigma_{T_1-1}^{(g)}(\mathbf{x}_{T_1}, \mathbf{w}) < \eta/2$. Furthermore, by combining this and Lemma A.3 in (Inatsu et al., 2021), we get $u_{T_1-1}^{(G_{T_1})}(\mathbf{x}_{T_1}; \eta) = l_{T_1-1}^{(G_{T_1})}(\mathbf{x}_{T_1}; \eta)$. Using this and the definition of M_t , it holds that $\mathbf{x}_{T_1} \notin M_{T_1}$. Moreover, from Definition 3.1, it follows that $\mathbf{x}_{T_1} \in H_{T_1} \cup M_{T_1}$. Thus, we have $\mathbf{x}_{T_1} \in H_{T_1}$. Similarly, we consider $(\mathbf{x}_{T+1}, \mathbf{w}_{T+1}), \dots, (\mathbf{x}_{2T}, \mathbf{w}_{2T})$. As with T_1 , let

$$T_2 = \operatorname{argmin}_{T+1 \leq t \leq 2T} \{\sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w}_t) + \sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w}_t)\}.$$

Then, the following inequality holds:

$$T\{\sigma_{T_2-1}^{(f)2}(\mathbf{x}_{T_2}, \mathbf{w}_{T_2}) + \sigma_{T_2-1}^{(g)2}(\mathbf{x}_{T_2}, \mathbf{w}_{T_2})\} \leq \sum_{t=T+1}^{2T} \sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w}_t) + \sum_{t=T+1}^{2T} \sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w}_t).$$

Furthermore, let $\sigma_0^{(f)2}(\mathbf{x}, \mathbf{w} | \mathbf{x}_{m:n}, \mathbf{w}_{m:n})$ be a posterior variance of $f(\mathbf{x}, \mathbf{w})$ after adding $(\mathbf{x}_m, \mathbf{w}_m), \dots, (\mathbf{x}_n, \mathbf{w}_n)$. Then, it holds that

$$\sum_{t=T+1}^{2T} \sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w}_t) \leq \sigma_0^{(f)2}(\mathbf{x}_{T+1}, \mathbf{w}_{T+1}) + \sum_{t=2}^T \sigma_0^{(f)2}(\mathbf{x}_{T+t}, \mathbf{w}_{T+t} | \mathbf{x}_{(T+1):(T+t-1)}, \mathbf{w}_{(T+1):(T+t-1)}) \leq C_{1,f}\kappa_T^{(f)}.$$

As with $\sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w}_t)$, the following holds for $\sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w}_t)$:

$$\sum_{t=T+1}^{2T} \sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w}_t) \leq C_{1,g}\kappa_T^{(g)}.$$

Thus, the following inequality holds for T_2 :

$$T\{\sigma_{T_2-1}^{(f)2}(\mathbf{x}_{T_2}, \mathbf{w}_{T_2}) + \sigma_{T_2-1}^{(g)2}(\mathbf{x}_{T_2}, \mathbf{w}_{T_2})\} \leq C_{1,f}\kappa_T^{(f)} + C_{1,g}\kappa_T^{(g)}.$$

Hence, from the same argument as before, we obtain $\mathbf{x}_{T_2} \in H_{T_2}$. By repeating this procedure up to T^2 , we get the sequence $\mathbf{x}_{T_1}, \mathbf{x}_{T_2}, \dots, \mathbf{x}_{T_T}$ satisfying $\mathbf{x}_{T_i} \in H_{T_i}$.

Next, from $\mathbf{x}_{T_i} \in H_{T_i}$, it follows that $a_{T_i-1}^{(G_{T_i-1})}(\mathbf{x}_{T_i}) = 1$ and $a_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) \leq u_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) - l_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i})$. Therefore, it holds that

$$a_{T_i-1}(\mathbf{x}_{T_i}) \leq u_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) - l_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}).$$

Here, let $\tilde{p}(\mathbf{w}) \in \mathcal{A}_{T_i-1}$ be a probability function satisfying

$$l_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) = \sum_{\mathbf{w} \in \Omega} l_{T_i-1}^{(f)}(\mathbf{x}_{T_i}, \mathbf{w}) \tilde{p}(\mathbf{w}).$$

Then, from the definition of $u_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i})$, the following holds:

$$u_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) \leq \sum_{\mathbf{w} \in \Omega} u_{T_i-1}^{(f)}(\mathbf{x}_{T_i}, \mathbf{w}) \tilde{p}(\mathbf{w}).$$

Thus, we get

$$u_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) - l_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) \leq \sum_{\mathbf{w} \in \Omega} 2\beta_{f,T_i}^{1/2} \sigma_{T_i-1}^{(f)}(\mathbf{x}_{T_i}, \mathbf{w}) \tilde{p}(\mathbf{w}) \leq 2\beta_{f,T_i}^{1/2} \max_{\mathbf{w} \in \Omega} \sigma_{T_i-1}^{(f)}(\mathbf{x}_{T_i}, \mathbf{w}).$$

Hence, from Definition 3.2 it follows that

$$a_{T_i-1}^2(\mathbf{x}_{T_i}) \leq 4\beta_{f,T_i} (\max_{\mathbf{w} \in \Omega} \sigma_{T_i-1}^{(f)}(\mathbf{x}_{T_i}, \mathbf{w}))^2 \leq 4\beta_{f,T^2} \max_{\mathbf{w} \in \Omega} \sigma_{T_i-1}^{(f)2}(\mathbf{x}_{T_i}, \mathbf{w}) \leq 4\beta_{f,T^2} (\sigma_{T_i-1}^{(f)2}(\mathbf{x}_{T_i}, \mathbf{w}_{T_i}) + \sigma_{T_i-1}^{(g)2}(\mathbf{x}_{T_i}, \mathbf{w}_{T_i})).$$

Furthermore, let

$$\tilde{T} = \operatorname{argmin}_{T \in \{T_1, \dots, T_T\}} a_{T-1}^2(\mathbf{x}_T).$$

Then, the following inequality holds:

$$T a_{\tilde{T}-1}^2(\mathbf{x}_{\tilde{T}}) \leq 4\beta_{f,T^2} \sum_{i=1}^T (\sigma_{T_i-1}^{(f)2}(\mathbf{x}_{T_i}, \mathbf{w}_{T_i}) + \sigma_{T_i-1}^{(g)2}(\mathbf{x}_{T_i}, \mathbf{w}_{T_i})) \leq 4\beta_{f,T^2} (C_{1,f} \kappa_T^{(f)} + C_{1,g} \kappa_T^{(g)}).$$

This implies that

$$a_{\tilde{T}-1}^2(\mathbf{x}_{\tilde{T}}) \leq \frac{4\beta_{f,T^2} (C_{1,f} \kappa_T^{(f)} + C_{1,g} \kappa_T^{(g)})}{T}.$$

On the other hand, for any $\mathbf{x} \in H_{\tilde{T}-1} \cup M_{\tilde{T}-1}$, from the definition of $a_{\tilde{T}-1}^{(G_{\tilde{T}-1})}(\mathbf{x})$ it follows that $a_{\tilde{T}-1}^{(G_{\tilde{T}-1})}(\mathbf{x}) \geq \xi$. Hence, $a_{\tilde{T}-1}(\mathbf{x})$ can be bounded as

$$(u_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}) - \max_{\mathbf{x} \in H_{\tilde{T}-1}} l_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x})) \xi = (u_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}) - c_{\tilde{T}-1}^{(\text{best})}) \xi \leq a_{\tilde{T}-1}(\mathbf{x}).$$

Therefore, by using the theorem's assumption, we obtain

$$(u_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}) - \max_{\mathbf{x} \in H_{\tilde{T}-1}} l_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}))^2 \leq \frac{4\beta_{f,T^2} (C_{1,f} \kappa_T^{(f)} + C_{1,g} \kappa_T^{(g)})}{T} \xi^{-2} < \xi^2.$$

Hence, $\mathbf{x}_{\tilde{T}-1} \in H_{\tilde{T}-1}$ and $u_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}) - \max_{\mathbf{x} \in H_{\tilde{T}-1}} l_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}) < \xi$ for any $\mathbf{x} \in H_{\tilde{T}-1} \cup M_{\tilde{T}-1}$. Thus, the stopping condition (S2) holds.

A.3. Proof of Theorem 4.4

The proof is almost the same as the proof of Theorem 4.3. Assume that there exists a positive integer $t \leq T^2$ such that $L_t = \mathcal{X}$. Then, the stopping condition (S1) holds.

Next, we consider the case where $L_t \neq \mathcal{X}$ for any $t \leq T^2$. For each $i \in \{1, \dots, T\}$, let

$$T_i = \operatorname{argmin}_{(i-1)T+1 \leq t \leq iT} \{\mathbb{E}_{\mathbf{w}}[\sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w})] + \mathbb{E}_{\mathbf{w}}[\sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w})]\}.$$

Then, T_1 satisfies that

$$T \{\mathbb{E}_{\mathbf{w}}[\sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w})] + \mathbb{E}_{\mathbf{w}}[\sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w})]\} \leq \sum_{t=1}^T \mathbb{E}_{\mathbf{w}}[\sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w})] + \sum_{t=1}^T \mathbb{E}_{\mathbf{w}}[\sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w})].$$

Furthermore, from Lemma 3 in (Kirschner & Krause, 2018), the following uniform bound holds with a probability of at least $1 - \delta$:

$$\begin{aligned} \sum_{t=1}^T \mathbb{E}_{\mathbf{w}}[\sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w})] + \sum_{t=1}^T \mathbb{E}_{\mathbf{w}}[\sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w})] &\leq 2 \sum_{t=1}^T \sigma_{t-1}^{(f)2}(\mathbf{x}_t, \mathbf{w}_t) + 2 \sum_{t=1}^T \sigma_{t-1}^{(g)2}(\mathbf{x}_t, \mathbf{w}_t) + 16 \log(5/\delta) \\ &\leq 2C_{1,f}\kappa_T^{(f)} + 2C_{1,g}\kappa_T^{(g)} + 16 \log(5/\delta). \end{aligned}$$

By combining these, we have

$$\mathbb{E}_{\mathbf{w}}[\sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w})] + \mathbb{E}_{\mathbf{w}}[\sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w})] \leq \frac{2C_{1,f}\kappa_T^{(f)} + 2C_{1,g}\kappa_T^{(g)} + 16 \log(5/\delta)}{T}.$$

In addition, noting that $p_{\min} > 0$, the following inequality holds for any $\mathbf{w} \in \Omega$:

$$\begin{aligned} \sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w}) &\leq \max_{\mathbf{w} \in \Omega}(\sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w}) + \sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w})) \\ &\leq p_{\min}^{-1}(\mathbb{E}_{\mathbf{w}}[\sigma_{T_1-1}^{(f)2}(\mathbf{x}_{T_1}, \mathbf{w})] + \mathbb{E}_{\mathbf{w}}[\sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w})]) \leq \frac{C_{2,f}\kappa_T^{(f)} + C_{2,g}\kappa_T^{(g)} + C_3}{T}. \end{aligned}$$

Hence, we get

$$\beta_{g,T_1}\sigma_{T_1-1}^{(g)2}(\mathbf{x}_{T_1}, \mathbf{w}) \leq \frac{\beta_{g,T^2}(C_{2,f}\kappa_T^{(f)} + C_{2,g}\kappa_T^{(g)} + C_3)}{T}.$$

Moreover, from the theorem's assumption, it follows that

$$\frac{\beta_{g,T^2}(C_{2,f}\kappa_T^{(f)} + C_{2,g}\kappa_T^{(g)} + C_3)}{T} < \eta^2/4.$$

Thus, by using the same argument as the proof of Theorem 4.3, we obtain $\mathbf{x}_{T_1} \in H_{T_1}$. By repeating this procedure up to T^2 , we have the sequence $\mathbf{x}_{T_1}, \mathbf{x}_{T_2}, \dots, \mathbf{x}_{T_T}$ satisfying $\mathbf{x}_{T_i} \in H_{T_i}$. From $\mathbf{x}_{T_i} \in H_{T_i}$, it follows that $a_{T_i-1}^{(G_{T_i-1})}(\mathbf{x}_{T_i}) = 1$ and $a_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) \leq u_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) - l_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i})$. Therefore, from the definition of the proposed AF, $a_{T_i-1}(\mathbf{x}_{T_i})$ can be bounded as

$$a_{T_i-1}(\mathbf{x}_{T_i}) \leq u_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) - l_{T_i-1}^{(F_{T_i-1})}(\mathbf{x}_{T_i}) \leq 2\beta_{f,T_i}^{1/2} \max_{\mathbf{w} \in \Omega} \sigma_{T_i-1}^{(f)}(\mathbf{x}_{T_i}, \mathbf{w}).$$

In addition, noting that $p_{\min} > 0$ we get

$$\begin{aligned} a_{T_i-1}^2(\mathbf{x}_{T_i}) &\leq 4\beta_{f,T_i}(\max_{\mathbf{w} \in \Omega} \sigma_{T_i-1}^{(f)}(\mathbf{x}_{T_i}, \mathbf{w}))^2 \leq 4\beta_{f,T^2} \max_{\mathbf{w} \in \Omega} \sigma_{T_i-1}^{(f)2}(\mathbf{x}_{T_i}, \mathbf{w}) \\ &\leq 4p_{\min}^{-1}\beta_{f,T^2}(\mathbb{E}_{\mathbf{w}}[\sigma_{T_i-1}^{(f)2}(\mathbf{x}_{T_i}, \mathbf{w})] + \mathbb{E}_{\mathbf{w}}[\sigma_{T_i-1}^{(g)2}(\mathbf{x}_{T_i}, \mathbf{w})]). \end{aligned}$$

Let \tilde{T} be an positive integer satisfying

$$\tilde{T} = \operatorname{argmin}_{T \in \{T_1, \dots, T_T\}} a_{T-1}^2(\mathbf{x}_T).$$

Then, it follows that

$$T a_{\tilde{T}-1}^2(\mathbf{x}_{\tilde{T}}) \leq 4p_{\min}^{-1}\beta_{f,T^2} \sum_{i=1}^T (\mathbb{E}_{\mathbf{w}}[\sigma_{T_i-1}^{(f)2}(\mathbf{x}_{T_i}, \mathbf{w})] + \mathbb{E}_{\mathbf{w}}[\sigma_{T_i-1}^{(g)2}(\mathbf{x}_{T_i}, \mathbf{w})]) \leq 4\beta_{f,T^2}(C_{2,f}\kappa_T^{(f)} + C_{2,g}\kappa_T^{(g)} + C_3).$$

This implies that

$$a_{\tilde{T}-1}^2(\mathbf{x}_{\tilde{T}}) \leq \frac{4\beta_{f,T^2}(C_{2,f}\kappa_T^{(f)} + C_{2,g}\kappa_T^{(g)} + C_3)}{T}.$$

Moreover, for any $\mathbf{x} \in H_{\tilde{T}-1} \cup M_{\tilde{T}-1}$, from the definition of $a_{\tilde{T}-1}^{(G_{\tilde{T}-1})}(\mathbf{x})$ it holds that $a_{\tilde{T}-1}^{(G_{\tilde{T}-1})}(\mathbf{x}) \geq \xi$. Hence, the following holds:

$$(u_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}) - \max_{\mathbf{x} \in H_{\tilde{T}-1}} l_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}))\xi = (u_{\tilde{T}-1}^{(F_{\tilde{T}-1})}(\mathbf{x}) - c_{\tilde{T}-1}^{(\text{best})})\xi \leq a_{\tilde{T}-1}(\mathbf{x}).$$

Thus, from the theorem's assumption, it follows that

$$(u_{\hat{T}-1}^{(F_{\hat{T}-1})}(\mathbf{x}) - \max_{\mathbf{x} \in H_{\hat{T}-1}} l_{\hat{T}-1}^{(F_{\hat{T}-1})}(\mathbf{x}))^2 \leq \frac{4\beta_{f,T^2}(C_{2,f}\kappa_T^{(f)} + C_{2,g}\kappa_T^{(g)} + C_3)}{T} \xi^{-2} < \xi^2.$$

Therefore, from $\mathbf{x}_{\hat{T}-1} \in H_{\hat{T}-1}$ and $u_{\hat{T}-1}^{(F_{\hat{T}-1})}(\mathbf{x}) - \max_{\mathbf{x} \in H_{\hat{T}-1}} l_{\hat{T}-1}^{(F_{\hat{T}-1})}(\mathbf{x}) < \xi$ for any $\mathbf{x} \in H_{\hat{T}-1} \cup M_{\hat{T}-1}$, the stopping condition (S2) holds.

A.4. Proof of Theorem 4.6

Let $p_t^*(\mathbf{w})$ be an empirical distribution of \mathbf{w} . Then, from the Hoeffding's inequality, the following holds for any $\mathbf{w} \in \Omega$:

$$\mathbb{P}(|p_t^*(\mathbf{w}) - p^\dagger(\mathbf{w})| \geq \lambda) \leq 2 \exp(-2t\lambda^2).$$

By letting

$$\lambda = \sqrt{\frac{1}{2t} \log \left(\frac{|\Omega| \pi^2 t^2}{3\delta} \right)},$$

with a probability of at least $1 - \delta$, the following inequality holds for any $t \leq 1$ and $\mathbf{w} \in \Omega$:

$$|p_t^*(\mathbf{w}) - p^\dagger(\mathbf{w})| \leq \lambda.$$

Moreover, from the theorem's assumption, the distance between distributions can be expressed as

$$d(p_t^*(\mathbf{w}), p^\dagger(\mathbf{w})) = \sum_{\mathbf{w} \in \Omega} |p_t^*(\mathbf{w}) - p^\dagger(\mathbf{w})| \leq |\Omega| \lambda = \epsilon_t.$$

Thus, it follows that $p^\dagger(\mathbf{w}) \in \mathcal{A}_t$. Here, if the stopping condition (S1) is satisfied, from Theorem 4.2, with a probability of at least $1 - \delta$ the inequality $G_t(\mathbf{x}) \leq \alpha'$ holds for any $\mathbf{x} \in \mathcal{X}$. Therefore, we get

$$\tilde{G}(\mathbf{x}) = \{\tilde{G}(\mathbf{x}) - G_t(\mathbf{x})\} + G_t(\mathbf{x}) \leq |\tilde{G}(\mathbf{x}) - G_t(\mathbf{x})| + \alpha'.$$

Furthermore, let $p_t(\mathbf{w})$ be a probability function satisfying

$$G_t(\mathbf{x}) = \sum_{\mathbf{w} \in \Omega} \mathbb{1}[g(\mathbf{x}, \mathbf{w}) > h] p_t(\mathbf{w}).$$

Then, noting that $p_t(\mathbf{w}), p^\dagger(\mathbf{w}) \in \mathcal{A}_t$, $|\tilde{G}(\mathbf{x}) - G_t(\mathbf{x})|$ can be expressed as follows:

$$\begin{aligned} |\tilde{G}(\mathbf{x}) - G_t(\mathbf{x})| &= \left| \sum_{\mathbf{w} \in \Omega} \mathbb{1}[g(\mathbf{x}, \mathbf{w}) > h] \{p^\dagger(\mathbf{w}) - p_t(\mathbf{w})\} \right| \\ &\leq \sum_{\mathbf{w} \in \Omega} |p^\dagger(\mathbf{w}) - p_t(\mathbf{w})| \\ &= \sum_{\mathbf{w} \in \Omega} |p^\dagger(\mathbf{w}) - p_t^*(\mathbf{w}) + p_t^*(\mathbf{w}) - p_t(\mathbf{w})| \\ &\leq \sum_{\mathbf{w} \in \Omega} |p^\dagger(\mathbf{w}) - p_t^*(\mathbf{w})| + \sum_{\mathbf{w} \in \Omega} |p_t^*(\mathbf{w}) - p_t(\mathbf{w})| \\ &= d(p^\dagger(\mathbf{w}), p_t^*(\mathbf{w})) + d(p_t^*(\mathbf{w}), p_t(\mathbf{w})) \leq 2\epsilon_t \leq 2(1 + \beta_{f,1}^{1/2})\epsilon_t < \xi. \end{aligned}$$

Hence, we have

$$\tilde{G}(\mathbf{x}) \leq |\tilde{G}(\mathbf{x}) - G_t(\mathbf{x})| + \alpha' < \xi + \alpha' = \xi + (\alpha - \xi) = \alpha.$$

Thus, it holds that $\tilde{G}(\mathbf{x}) < \alpha$ with a probability of at least $1 - 2\delta$. This implies that the CC problem has no solution.

Next, if the stopping condition (S2) is satisfied, $\hat{\mathbf{x}}_t$ satisfies the following inequality with a probability of at least $1 - \delta$:

$$G_t(\hat{\mathbf{x}}_t) \geq \alpha' - 2\xi = \alpha - 3\xi.$$

Noting that $p^\dagger(\mathbf{w}) \in \mathcal{A}_t$, we obtain

$$\tilde{G}(\hat{\mathbf{x}}_t) \geq G_t(\hat{\mathbf{x}}_t) \geq \alpha' - 2\xi = \alpha - 3\xi.$$

Here, if the CC problem has no solution, then from the definition the following holds:

$$\tilde{F}(\tilde{\mathbf{x}}^*) - \tilde{F}(\hat{\mathbf{x}}_t) \leq 0 < 3\xi.$$

Hence, $\hat{\mathbf{x}}_t$ is a 3ξ -accurate solution for the CC problem. Similarly, if the optimal solution $\tilde{\mathbf{x}}^*$ to the CC problem exists, we get

$$\begin{aligned} & \tilde{F}(\tilde{\mathbf{x}}^*) - \tilde{F}(\hat{\mathbf{x}}_t) \\ &= \tilde{F}(\tilde{\mathbf{x}}^*) - F_t(\tilde{\mathbf{x}}^*) + F_t(\tilde{\mathbf{x}}^*) - F_t(\hat{\mathbf{x}}_t) + F_t(\hat{\mathbf{x}}_t) - \tilde{F}(\hat{\mathbf{x}}_t). \end{aligned} \quad (10)$$

Because $\tilde{\mathbf{x}}^*$ is the optimal solution to the CC problem, we have $\tilde{G}(\tilde{\mathbf{x}}^*) > \alpha$. Hence, by using this we obtain

$$\begin{aligned} G_t(\tilde{\mathbf{x}}^*) &= \tilde{G}(\tilde{\mathbf{x}}^*) + \{G_t(\tilde{\mathbf{x}}^*) - \tilde{G}(\tilde{\mathbf{x}}^*)\} \\ &> \alpha - |\tilde{G}(\tilde{\mathbf{x}}^*) - G_t(\tilde{\mathbf{x}}^*)| > \alpha - \xi = \alpha'. \end{aligned}$$

Therefore, from the definition of \mathbf{x}_t^* , it follows that

$$F_t(\tilde{\mathbf{x}}^*) \leq F_t(\mathbf{x}_t^*). \quad (11)$$

In addition, from $p^\dagger(\mathbf{w}) \in \mathcal{A}_t$ and the definition of $F_t(\mathbf{x})$ and $\tilde{F}(\mathbf{x})$, the following inequality holds:

$$F_t(\hat{\mathbf{x}}_t) - \tilde{F}(\hat{\mathbf{x}}_t) \leq 0. \quad (12)$$

Moreover, let $\tilde{p}_t(\mathbf{w}) \in \mathcal{A}_t$ be a probability function satisfying

$$F_t(\tilde{\mathbf{x}}^*) = \sum_{\mathbf{w} \in \Omega} f(\tilde{\mathbf{x}}^*, \mathbf{w}) \tilde{p}_t(\mathbf{w}).$$

Then, we get

$$|\tilde{F}(\tilde{\mathbf{x}}^*) - F_t(\tilde{\mathbf{x}}^*)| \leq \sum_{\mathbf{w} \in \Omega} |f(\tilde{\mathbf{x}}^*, \mathbf{w})| |p^\dagger(\mathbf{w}) - \tilde{p}_t(\mathbf{w})|.$$

From Lemma 5.1 in (Srinivas et al., 2010), the following holds with a probability of at least $1 - \delta$:

$$|f(\tilde{\mathbf{x}}^*, \mathbf{w})| \leq \beta_{f,1}^{1/2} \sigma_0(\tilde{\mathbf{x}}^*, \mathbf{w}) \leq \beta_{f,1}^{1/2}.$$

By using this, we have

$$\begin{aligned} \tilde{F}(\tilde{\mathbf{x}}^*) - F_t(\tilde{\mathbf{x}}^*) &\leq |\tilde{F}(\tilde{\mathbf{x}}^*) - F_t(\tilde{\mathbf{x}}^*)| \\ &\leq \beta_{f,1}^{1/2} (d(p^\dagger(\mathbf{w}), p_t^*(\mathbf{w})) + d(p_t^*(\mathbf{w}), \tilde{p}_t(\mathbf{w}))) \leq 2\beta_{f,1}^{1/2} \epsilon_t \leq 2(1 + \beta_{f,1}^{1/2}) \epsilon_t < \xi. \end{aligned} \quad (13)$$

By substituting (11), (12) and (13) into (10), we obtain

$$\tilde{F}(\tilde{\mathbf{x}}^*) - \tilde{F}(\hat{\mathbf{x}}_t) < \xi + F_t(\mathbf{x}_t^*) - F_t(\hat{\mathbf{x}}_t).$$

Finally, from Theorem 4.2, noting that the \mathbf{x}_t^* is a 2ξ -accurate solution for the DRCC problem, we get

$$F_t(\mathbf{x}_t^*) - F_t(\hat{\mathbf{x}}_t) < 2\xi.$$

Therefore, we get $\tilde{F}(\tilde{\mathbf{x}}^*) - \tilde{F}(\hat{\mathbf{x}}_t) < 3\xi$.

B. Experimental Details

In this section, we give the details of the experiments conducted in Section 5.

Table 2. Experimental parameters for each setting in the synthetic function

	Parameter
Simulator	
Fixed	$\sigma_{f,\ker}^2 = 1, L_f = 3, \sigma_{f,\text{noise}}^2 = 10^{-8}, \beta_{f,t}^{1/2} = 3, \sigma_{g,\ker}^2 = 2500, L_g = 4, \sigma_{g,\text{noise}}^2 = 10^{-4}, \beta_{g,t}^{1/2} = 2, h = 5, \alpha = 0.53$
Data-driven	

Table 3. Experimental parameters for each setting in the SIR model simulation

	Parameter
Case1	$\sigma_{f,\ker}^2 = 5000, L_f = 0.1, \sigma_{f,\text{noise}}^2 = 10^{-8}, \beta_{f,t}^{1/2} = 3, \sigma_{g,\ker}^2 = 10^5, L_g = 0.01, \sigma_{g,\text{noise}}^2 = 10^{-4}, \beta_{g,t}^{1/2} = 2, h = 320, \alpha = 0.85$
Case2	$\sigma_{f,\ker}^2 = 5000, L_f = 0.1, \sigma_{f,\text{noise}}^2 = 10^{-8}, \beta_{f,t}^{1/2} = 3, \sigma_{g,\ker}^2 = 10^5, L_g = 0.01, \sigma_{g,\text{noise}}^2 = 10^{-4}, \beta_{g,t}^{1/2} = 2, h = 320, \alpha = 0.85$
Case3	$\sigma_{f,\ker}^2 = 10^4, L_f = 0.1, \sigma_{f,\text{noise}}^2 = 10^{-3}, \beta_{f,t}^{1/2} = 2, \sigma_{g,\ker}^2 = 10^5, L_g = 0.1, \sigma_{g,\text{noise}}^2 = 10^{-3}, \beta_{g,t}^{1/2} = 3, h = 100, \alpha = 0.69$
Case4	$\sigma_{f,\ker}^2 = 10^4, L_f = 0.1, \sigma_{f,\text{noise}}^2 = 10^{-3}, \beta_{f,t}^{1/2} = 2, \sigma_{g,\ker}^2 = 10^5, L_g = 0.1, \sigma_{g,\text{noise}}^2 = 10^{-3}, \beta_{g,t}^{1/2} = 3, h = 100, \alpha = 0.69$

Experimental Parameter The experimental parameters used in each experiment are given in Table 2 and 3.

True Distribution of Environmental Variables We give the details of the true distribution $p^\dagger(w)$ considered in the uncontrollable setting in the synthetic function experiment. Let $h(w; \mu, \sigma^2)$ be a probability density function of Normal distribution with mean μ and variance σ^2 , and let $h(w) = 0.5h(w; -5, 10) + 0.5h(w; 5, 10)$. Then, $p^\dagger(w)$ is given by

$$p^\dagger(w) = \frac{h(w)}{\sum_{w \in \Omega} h(w)}.$$

DRPTR The DRPTR AF is based on the expected classification improvement for $G_t(\mathbf{x})$ after adding new data $(\mathbf{x}^*, \mathbf{w}^*, y^{(g)*})$. Let $l_t^{(G_t)}(\mathbf{x}; \eta | \mathbf{x}^*, \mathbf{w}^*, y^{(g)*})$ be a lower of the credible interval of $G_t(\mathbf{x})$ at \mathbf{x} after adding $(\mathbf{x}^*, \mathbf{w}^*, y^{(g)*})$. Then, the expected classification improvement is given by

$$a_t(\mathbf{x}^*, \mathbf{w}^*) = \sum_{\mathbf{x} \in M_t} \mathbb{E}_{y^{(g)*}} [\mathbb{1}[l_t^{(G_t)}(\mathbf{x}; \eta | \mathbf{x}^*, \mathbf{w}^*, y^{(g)*}) > \alpha]]. \quad (14)$$

In (Inatsu et al., 2021), they suggest combining (14) and RMILE AF proposed by (Zanette et al., 2018). The RMILE is based on the expected classification improvement for $g(\mathbf{x}, \mathbf{w})$ after adding $(\mathbf{x}^*, \mathbf{w}^*, y^{(g)*})$. Let $l_t^{(g)}(\mathbf{x}, \mathbf{w} | \mathbf{x}^*, \mathbf{w}^*, y^{(g)*})$ be a lower of the credible interval of $g(\mathbf{x}, \mathbf{w})$ at (\mathbf{x}, \mathbf{w}) after adding $(\mathbf{x}^*, \mathbf{w}^*, y^{(g)*})$. In our experiments, we used the following modified RMILE function:

$$\text{RMILE}_t(\mathbf{x}^*, \mathbf{w}^*) = \sum_{(\mathbf{x}, \mathbf{w}) \in M_t \times \Omega} \mathbb{E}_{y^{(g)*}} [\mathbb{1}[l_t^{(g)}(\mathbf{x}, \mathbf{w} | \mathbf{x}^*, \mathbf{w}^*, y^{(g)*}) > h]]. \quad (15)$$

Then, the DRPTR AF is defined as

$$\text{DRPTR}_t(\mathbf{x}^*, \mathbf{w}^*) = \max\{a_t(\mathbf{x}^*, \mathbf{w}^*), \gamma \text{RMILE}_t(\mathbf{x}^*, \mathbf{w}^*)\}, \quad (16)$$

where γ is a trade-off parameter. In all experiments, we set $\gamma = 0.1$. From GP properties, (15) can be calculated analytically (Zanette et al., 2018). In contrast, (14) can be represented in an exact form (see, (Inatsu et al., 2021)), but its computational cost is high. In Lemma 3.3 in (Inatsu et al., 2021), an arbitrary-accurate approximation method for calculating (14) is proposed. For all experiments, we used its lemma with approximation parameter $\zeta = 0.005(|\Omega| + 1)$. This implies that the calculation error between the true (14) and approximated one is at most ζ .

CCBO The CCBO AF is based on the expected feasible improvement for the following CC problem:

$$\max_{\mathbf{x} \in \mathcal{X}} Z^{(F)}(\mathbf{x}) \quad \text{s.t.} \quad Z^{(G)}(\mathbf{x}) > \alpha,$$

where $Z^{(F)}(\mathbf{x})$ and $Z^{(G)}(\mathbf{x})$ are given by

$$Z^{(F)}(\mathbf{x}) = \sum_{\mathbf{w} \in \Omega} f(\mathbf{x}, \mathbf{w}) p^\dagger(\mathbf{w}), Z^{(G)}(\mathbf{x}) = \sum_{\mathbf{w} \in \Omega} \mathbb{1}[g(\mathbf{x}, \mathbf{w}) > h] p^\dagger(\mathbf{w}).$$

In our experiments, to define $Z^{(F)}(\mathbf{x})$ and $Z^{(G)}(\mathbf{x})$, we used the reference distribution instead of $p^\dagger(\mathbf{w})$. Let $Z_t^{(F)}(\mathbf{x})$ and $Z_t^{(G)}(\mathbf{x})$ be posterior distributions of $Z^{(F)}(\mathbf{x})$ and $Z^{(G)}(\mathbf{x})$, respectively. Here, the calculation of posterior distribution is based on GP posteriors of f and g . Then, the CCBO AF is given by

$$\text{CCBO}_t(\mathbf{x}) = \mathbb{E}[\max\{Z_t^{(F)}(\mathbf{x}) - c_t^{(\text{feas})}, 0\}] \times \mathbb{P}(Z^{(G)}(\mathbf{x}) > \alpha), \quad (17)$$

where $c_t^{(\text{feas})}$ is given by

$$c_t^{(\text{feas})} = \begin{cases} \max_{\mathbf{x} \in S_t} \mathbb{E}[Z_t^{(F)}(\mathbf{x})] & \text{if } S_t \equiv \{\mathbf{x} \in \mathcal{X} \mid \mathbb{E}[Z^{(G)}(\mathbf{x})] > \alpha\} \neq \emptyset, \\ \mathbb{E}[Z_t^{(F)}(\tilde{\mathbf{x}})], \tilde{\mathbf{x}} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \mathbb{E}[Z^{(G)}(\mathbf{x})] & \text{otherwise} \end{cases}.$$

We select \mathbf{x}_{t+1} by maximizing $\text{CCBO}_t(\mathbf{x})$, that is,

$$\mathbf{x}_{t+1} = \operatorname{argmax}_{\mathbf{x} \in \mathcal{X}} \text{CCBO}_t(\mathbf{x}).$$

In CCBO, the selection of \mathbf{w} is based on the variance of $\text{CCBO}_t(\mathbf{x}_{t+1})$ after adding $(\mathbf{x}_{t+1}, \mathbf{w}^*, y^{(g)*})$. Let $\text{CCBO}_t(\mathbf{x}_{t+1} | \mathbf{x}_{t+1}, \mathbf{w}^*, y^{(g)*})$ be a value of $\text{CCBO}_t(\mathbf{x}_{t+1})$ after adding $(\mathbf{x}_{t+1}, \mathbf{w}^*, y^{(g)*})$. Then, we consider the variance of $\text{CCBO}_t(\mathbf{x}_{t+1} | \mathbf{x}_{t+1}, \mathbf{w}^*, y^{(g)*})$ with respect to $y^{(g)*}$:

$$\text{Var}_{y^{(g)*}}[\text{CCBO}_t(\mathbf{x}_{t+1} | \mathbf{x}_{t+1}, \mathbf{w}^*, y^{(g)*})] \equiv v_t(\mathbf{w}^*). \quad (18)$$

Using (18) we select \mathbf{w}_{t+1} as

$$\mathbf{w}_{t+1} = \operatorname{argmin}_{\mathbf{w}^* \in \Omega} v_t(\mathbf{w}^*).$$

Note that a part of the calculation of (17) and (18) requires a Monte Carlo approximation, we took 1000 samples and approximated them.

Hyperparameter Sensitivity of Uncontrollable Setting The experimental results of the hyperparameter sensitivity in the uncontrollable setting are given in Figure 4.

SIR Model Simulation The SIR model is often used in infectious disease modeling and is given as the following differential equation using the contact rate β and isolation rate γ :

$$\begin{cases} \frac{dS}{dT} = -\frac{\beta IS}{N}, \\ \frac{dI}{dT} = \frac{\beta IS}{N} - \gamma I, \\ \frac{dR}{dT} = \gamma I, \end{cases}$$

where $N = S + I + R$, and S , I and R are the number of susceptible, infected and removed people, respectively. In our experiment, we considered simulations from time $T = 0$ to time $T = 15$, and initial S_0 , I_0 , and R_0 were set to 990, 10, and 0, respectively. By considering $dT \approx 0.005$ and discrete approximation of the differential equation, we calculated the number of I_T for each $T \in \{0, 0.005, \dots, 15\} \equiv \mathcal{T}$. Using this we defined the maximum number of infected people $n_{\text{infected}}(\beta, \gamma)$ as

$$n_{\text{infected}}(\beta, \gamma) = \max_{T \in \mathcal{T}} I_T.$$

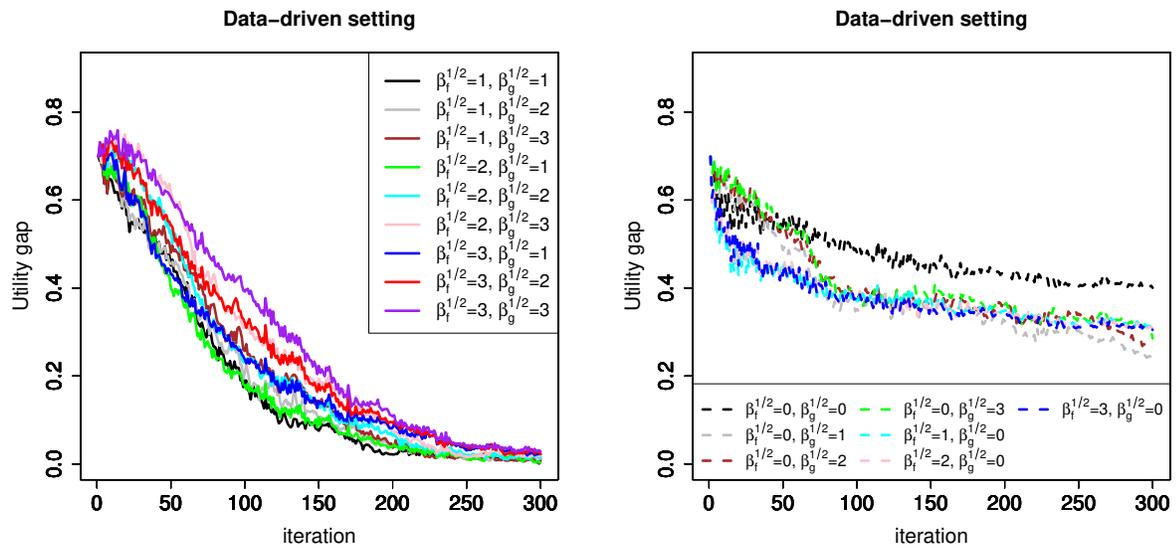


Figure 4. Hyperparameter sensitivity for the proposed method in the uncontrollable setting.