
Bounding Box-based Multi-objective Bayesian Optimization of Risk Measures under Input Uncertainty

Yu Inatsu
Nagoya Institute of Technology

Shion Takeno
RIKEN

Hiroyuki Hanada
RIKEN

Kazuki Iwata
Nagoya Institute of Technology

Ichiro Takeuchi
Nagoya University/RIKEN

Abstract

In this study, we propose a novel multi-objective Bayesian optimization (MOBO) method to efficiently identify the Pareto front (PF) defined by risk measures for black-box functions under the presence of input uncertainty (IU). Existing BO methods for Pareto optimization in the presence of IU are risk-specific or without theoretical guarantees, whereas our proposed method addresses general risk measures and has theoretical guarantees. The basic idea of the proposed method is to assume a Gaussian process (GP) model for the black-box function and to construct high-probability bounding boxes for the risk measures using the GP model. Furthermore, in order to reduce the uncertainty of non-dominated bounding boxes, we propose a method of selecting the next evaluation point using a maximin distance defined by the maximum value of a quasi distance based on bounding boxes. As theoretical analysis, we prove that the algorithm can return an arbitrary-accurate solution in a finite number of iterations with high probability, for various risk measures such as Bayes risk, worst-case risk, and value-at-risk. We also give a theoretical analysis that takes into account approximation errors because there exist non-negligible approximation errors (e.g., finite approximation of PFs and sampling-based approximation of bounding boxes) in practice. We confirm that the

proposed method performs as well or better than existing methods not only in the setting with IU but also in the setting of ordinary MOBO through numerical experiments.

1 INTRODUCTION

In this study, we treat a multi-objective Pareto optimization problem under input uncertainty (IU). In many real-world applications such as engineering, industry and computer simulations, it is often desired to simultaneously optimize an expensive-to-evaluate multi-objective black-box function. Because there is typically no point at which all functions are simultaneously optimal, the multi-objective optimization problem is often formulated as a Pareto optimization problem to identify the Pareto front (PF). The black-box functions actually handled often have IU. Our motivating example in this study is an expensive-to-evaluate docking simulation for real-world chemical compounds. The purpose of this simulation is to evaluate the inhibitory performance of candidate compounds on specific sites of some target protein. Because each compound has uncertain isomers, this simulation is expressed as the Pareto optimization problem under IU.

We consider a multi-objective black-box function optimization problem with M objective functions under IU with $m \in \{1, 2, \dots, M\}$. Let $f^{(m)}(\mathbf{x}, \mathbf{w})$ be the m -th objective function, where $\mathbf{x} \in \mathcal{X}$ and $\mathbf{w} \in \Omega$ are called *design variables* and *environmental variables*, respectively. The variable \mathbf{x} is an input that can be completely controlled, whereas \mathbf{w} is a random variable that cannot be controlled and follows some probability distribution. When considering a Pareto optimization problem in the presence of IU, it is necessary to consider optimization by taking into account the uncertainty of \mathbf{w} that cannot be controlled. A *risk measure*

Table 1: Characteristics of the proposed method and the representative existing methods

	Proposed	SABBa	BQ	MVA	MOBO without IU
IU setting	Yes	Yes	Yes	Yes	No
General risk setting	Yes	Yes	No	No	No
Theoretical guarantees	Yes	No	No	Yes	Yes/No
Approximation error setting	Yes	No	No	No	No

is the widely used function that is determined based on only \mathbf{x} while considering the uncertainty of \mathbf{w} . Various risk measures, for example, Bayes risk, worst-case risk and value-at-risk, are used depending on the problem. Given a risk measure $F^{(m)}(\mathbf{x}) \equiv \rho^{(m)}(f^{(m)}(\mathbf{x}, \mathbf{w}))$, the problem that we treat in this study is formulated as

$$\text{maximize } (F^{(1)}(\mathbf{x}), \dots, F^{(M)}(\mathbf{x})) \text{ s.t. } \mathbf{x} \in \mathcal{X}. \quad (1)$$

Bayesian optimization (BO) (Shahriari et al., 2015) using Gaussian processes (GPs) (Rasmussen and Williams, 2005) is a powerful tool for optimizing black-box functions. Many BO methods have been proposed for both single-objective and multi-objective black-box functions without IU. In contrast, designing BO methods for risk measures under the presence of IU is challenging. This is because risk measures cannot be observed directly and do not generally follow GPs even if black-box functions follow GPs. The main way to solve this problem is to design a predicted region that may contain a black-box function, then compute the risk measure on the region and use the lower and upper bounds of this to construct a predicted interval for the risk measure (Nguyen et al., 2021b,a; Kirschner et al., 2020). As an exception, special risk measures such as Bayes risk are known to follow GP in practice, allowing Bayesian quadrature (BQ)-based inference (Beland and Nair, 2017). Recently, multi-objective Bayesian optimization (MOBO) methods under IU have been proposed, which apply the BQ-based or predicted interval-based method (Qing et al., 2023; Iwazaki et al., 2021b; Rivier and Congedo, 2022). However, the BQ-based method proposed by Qing et al. (2023) and Mean-variance-analysis (MVA)-based method proposed by Iwazaki et al. (2021b) can only be applied to specific risk measures. On the other hand, the surrogate-assisted bounding-box approach (SABBa) proposed by Rivier and Congedo (2022) can be applied to general risk measures, but this method has no theoretical guarantees for constructing credible intervals (CIs) and is consequently a heuristic.

In this study, we propose a novel MOBO method based on *high-probability bounding boxes (HPBBs)* for risk measures using GP surrogate models, which solves the above problem. The basic idea of the proposed method is to design a high-probability credible region

(HPCR) that contains a black-box function with high probability. We use the fact that many risk measures can be expressed as a composite of a tractable function and some monotonic function, and construct high-probability CIs (HPCIs) of risk measures as a transformation of the lower and upper bounds of the tractable function. We also propose a method for computing a sampling-based CI of risk measures on the HPCR. Furthermore, we provide theoretical guarantees for these two methods in the case with/without various approximation errors that may occur in the practical computation. Through these results, we can propose a theoretically guaranteed MOBO methods for general risk measures. The characteristics of the proposed method and the representative existing methods are given in Table 1. Our contributions can be summarized as follows:

- We develop a general method for designing HPBB that can be applied to various risk measures.
- We propose a novel acquisition function (AF) for MOBO under IU, which effectively incorporates the quantified uncertainty of Pareto optimal solutions using HPBB.
- We provide theoretical guarantees of accuracy and termination based on HPBB and the proposed AF, as well as a theoretical error analysis that accounts for various types of approximation errors that can occur in the practical computation.

Related Work In the optimization of expensive-to-evaluate black-box functions, BO has gained popularity and has been the subject of active research. A variety of AFs for BO and MOBO settings have been introduced (Moćkus, 1975; Srinivas et al., 2010; Wang and Jegelka, 2017; Emmerich and Klinkenberg, 2008; Svenson and Santner, 2010; Zuluaga et al., 2016; Knowles, 2006; Suzuki et al., 2020). Moreover, multi-objective optimization has also been extensively studied in the evolutionary computation community (Deb et al., 2002). However, methodologies based on evolutionary computation often necessitate several thousand to tens of thousands of function evaluations (Deb and Gupta, 2005; Zhou et al., 2018), which can be prohibitively costly.

Studies on Pareto optimization under IU have also been gradually proposed in recent years, mainly in the development of BO methods to efficiently identify the PF defined by risk measures. Considered risk measures are, for example, Bayes risk (Qing et al., 2023), mean and negative standard deviation (Iwazaki et al., 2021b), and general risk measures (Rivier and Congedo, 2022). However, as mentioned earlier, these are methods that risk-specific or without theoretical guarantees. A BO method for a multivariate value-at-risk (MVaR) has also been proposed (Daulton et al., 2022). This method is similar to other MOBO methods, but is very different in essence. In general, in Pareto optimization, the PF is defined as the boundary defined by the points satisfying Pareto optimality, i.e., the points define the PF. On the other hand, MVaR is itself a PF, and the PF considered in Daulton et al. (2022) is defined as the boundary of the union of MVaR. Therefore, in Daulton et al. (2022), although the problem setup is Pareto optimization, the final PF is defined by PFs (MVaR). Thus, we only introduce it here because it differs from Pareto optimization in the essential point.

2 PRELIMINARY

Let $f^{(m)} : \mathcal{X} \times \Omega \rightarrow \mathbb{R}$ be an expensive-to-evaluate black-box function, where $m \in \{1, 2, \dots, M\} \equiv [M]$. Assume that the set of design variables \mathcal{X} and set of environmental variables Ω are compact and convex. For each $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$, $f^{(m)}(\mathbf{x}, \mathbf{w})$ is observed with noise as $y^{(m)} = f^{(m)}(\mathbf{x}, \mathbf{w}) + \varepsilon^{(m)}$, where $\varepsilon^{(m)}$ follows normal distribution with mean 0 and variance ς_m^2 , and the sequence of noises $(\varepsilon_i^{(m)})_{i \in \mathbb{N}, m \in [M]}$ is independent. In this study, $\mathbf{w} \in \Omega$ follows some distribution P_w , and $(\varepsilon_i^{(m)})_{i \in \mathbb{N}, m \in [M]}$ and $(\mathbf{w}_i)_{i \in \mathbb{N}}$ are mutually independent. In the BO framework including environment variables, two different settings for \mathbf{w} exist called the *simulator setting* and the *uncontrollable setting* (Kirschner et al., 2020; Iwazaki et al., 2021b; Inatsu et al., 2022). In the simulator setting, \mathbf{w} is fully controllable during optimization, whereas in the uncontrollable setting, \mathbf{w} is not controllable even during optimization. In the main body, only the simulator setting is treated, and the uncontrollable setting is discussed in Appendix A. Let $\rho^{(m)}(f^{(m)}(\mathbf{x}, \mathbf{w})) \equiv F^{(m)}(\mathbf{x})$ be a risk measure. For example, the widely used Bayes and worst-case risks are given by $F^{(m)}(\mathbf{x}) = \mathbb{E}[f^{(m)}(\mathbf{x}, \mathbf{w})]$ and $F^{(m)}(\mathbf{x}) = \inf_{\mathbf{w} \in \Omega} f^{(m)}(\mathbf{x}, \mathbf{w})$, respectively, where the expectation is taken with respect to \mathbf{w} . The purpose of this study is to efficiently identify the PF defined based on $F^{(m)}(\mathbf{x})$. For any $\mathbf{x} \in \mathcal{X}$ and $E \subset \mathcal{X}$, let $\mathbf{F}(\mathbf{x}) = (F^{(1)}(\mathbf{x}), \dots, F^{(M)}(\mathbf{x}))$ and $\mathbf{F}(E) = \{\mathbf{F}(\mathbf{x}) \mid \mathbf{x} \in E\}$. Then, for any $B \subset \mathbb{R}^M$, the dominated

region $\text{Dom}(B)$ and PF $\text{Par}(B)$ of B are defined as $\text{Dom}(B) = \{\mathbf{s} \in \mathbb{R}^M \mid \exists \mathbf{s}' \in B \text{ s.t. } \mathbf{s} \leq \mathbf{s}'\}$ and $\text{Par}(B) = \partial(\text{Dom}(B))$. Here, for any vector $\mathbf{a} = (a_1, \dots, a_M), \mathbf{b} = (b_1, \dots, b_M) \in \mathbb{R}^M$ and set C , $\mathbf{a} \leq \mathbf{b}$ represents $a_m \leq b_m$ for any $m \in [M]$, and $\partial(C)$ represents the boundary of C . Let Z^* be our target PF. Then, Z^* can be expressed as $Z^* = \text{Par}(\mathbf{F}(\mathcal{X}))$. Note that although we nominally use the term ‘‘risk measure’’ for $F^{(m)}(\mathbf{x})$, we are requesting only the property that $F^{(m)}(\mathbf{x})$ is a function based on only \mathbf{x} by removing the uncertainty of \mathbf{w} from $f^{(m)}(\mathbf{x}, \mathbf{w})$. That is, it allows for both functions that should be minimized and maximized. Since the minimization problem is equivalent to the maximization problem by multiplying minus one, it is formulated in a unified manner as a maximization problem in this study, as in (1).

We introduce a *regularity assumption* for $f^{(m)}$. For each $m \in [M]$, let $k^{(m)} : (\mathcal{X} \times \Omega) \times (\mathcal{X} \times \Omega) \rightarrow \mathbb{R}$ be a positive-definite kernel, where $k^{(m)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}, \mathbf{w})) \leq 1$ for any $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$. Also let $\mathcal{H}(k^{(m)})$ be a reproducing kernel Hilbert space corresponding to $k^{(m)}$. We assume that $f^{(m)}$ is the element of $\mathcal{H}(k^{(m)})$ and has the bounded Hilbert norm $\|f^{(m)}\|_{\mathcal{H}(k^{(m)})} \leq B_m < \infty$.

In this study, we use a GP model for the black-box function $f^{(m)}$. We assume the GP $\mathcal{GP}(0, k^{(m)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')))$ as the prior of $f^{(m)}$. For $m \in [M]$, given a dataset $\{(\mathbf{x}_i, \mathbf{w}_i, y_i^{(m)})\}_{i=1}^t$, where t is the number of queried instances, the posterior of $f^{(m)}$ is a GP. Then, its posterior mean $\mu_t^{(m)}(\mathbf{x}, \mathbf{w})$ and posterior variance $\sigma_t^{(m)2}(\mathbf{x}, \mathbf{w})$ can be calculated analytically (Rasmussen and Williams, 2005).

3 PROPOSED METHOD

In this section, we propose a BO method to efficiently identify Z^* . In Section 3.1, we provide a method for computing the CI of $F^{(m)}(\mathbf{x})$ using the CI of $f^{(m)}(\mathbf{x}, \mathbf{w})$. We also give a bounding box for $\mathbf{F}(\mathbf{x})$, which is the direct product of CIs.

3.1 Credible Interval and Bounding Box

For each input $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$ and $t \geq 1$, the CI of $f^{(m)}(\mathbf{x}, \mathbf{w})$ is denoted by $Q_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}) = [l_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}), u_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})]$, where $l_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ and $u_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ are given by $l_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}) = \mu_{t-1}^{(m)}(\mathbf{x}, \mathbf{w}) - \beta_{m,t}^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}, \mathbf{w})$ and $u_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}) = \mu_{t-1}^{(m)}(\mathbf{x}, \mathbf{w}) + \beta_{m,t}^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}, \mathbf{w})$. Here, $\beta_{m,t}^{1/2} \geq 0$ is a user-specified tradeoff parameter. If we set $\beta_{m,t}^{1/2}$ appropriately, $Q_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ becomes a HPCI which contains $f^{(m)}(\mathbf{x}, \mathbf{w})$ with high probability (details are de-

scribed in Section 4). For $\mathbf{x} \in \mathcal{X}$, $t \geq 1$ and $m \in [M]$, we define the set of functions $G_{t-1}^{(m)}(\mathbf{x}) = \{g(\mathbf{x}, \mathbf{w}) \mid \forall \mathbf{w} \in \Omega, g(\mathbf{x}, \mathbf{w}) \in Q_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})\}$. Let $Q_{t-1}^{(F^{(m)})}(\mathbf{x}) = [\text{lcb}_{t-1}^{(m)}(\mathbf{x}), \text{ucb}_{t-1}^{(m)}(\mathbf{x})]$ be a CI of $F^{(m)}(\mathbf{x})$. Also let $B_{t-1}(\mathbf{x}) = Q_{t-1}^{(F^{(1)})}(\mathbf{x}) \times \cdots \times Q_{t-1}^{(F^{(M)})}(\mathbf{x})$ be a bounding box of $\mathbf{F}(\mathbf{x})$. Then, when $Q_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ is HPCI for all $m \in [M]$, $t \geq 1$, $\mathbf{x} \in \mathcal{X}$ and $\mathbf{w} \in \Omega$, a sufficient condition for $Q_{t-1}^{(F^{(m)})}(\mathbf{x})$ to also be HPCI is given as follows:

$$\begin{aligned} \forall g(\mathbf{x}, \mathbf{w}) \in G_{t-1}^{(m)}(\mathbf{x}), \\ \text{lcb}_{t-1}^{(m)}(\mathbf{x}) \leq \rho^{(m)}(g(\mathbf{x}, \mathbf{w})) \leq \text{ucb}_{t-1}^{(m)}(\mathbf{x}). \end{aligned} \quad (2)$$

If (2) holds, then $B_{t-1}(\mathbf{x})$ is also a HPBB. Next, we provide computation methods for $\text{lcb}_{t-1}^{(m)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x})$. First, we provide a *generalized method* for $\text{lcb}_{t-1}^{(m)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x})$ to satisfy (2). The $\text{lcb}_{t-1}^{(m)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x})$ by the generalized method are calculated with $\text{lcb}_{t-1}^{(m)}(\mathbf{x}) = \inf_{g(\mathbf{x}, \mathbf{w}) \in G_{t-1}^{(m)}(\mathbf{x})} \rho^{(m)}(g(\mathbf{x}, \mathbf{w}))$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x}) = \sup_{g(\mathbf{x}, \mathbf{w}) \in G_{t-1}^{(m)}(\mathbf{x})} \rho^{(m)}(g(\mathbf{x}, \mathbf{w}))$. We emphasize that although the condition (2) holds by using the generalized method, the inf and sup calculations in the generalized method are not always easy. Therefore, in this study, we give additional two computation methods for $\text{lcb}_{t-1}^{(m)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x})$, (i) *decomposition method* and (ii) *sampling method*. In the decomposition method, $\text{lcb}_{t-1}^{(m)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x})$ in (2) are calculated directly. Let $\rho(\cdot)$ be a risk measure. In many cases, $\rho(\cdot)$ can be decomposed as $\rho(\cdot) = \tilde{\rho} \circ h(\cdot)$, where $\tilde{\rho}(\cdot)$ and $h(\cdot)$ are respectively monotonic and tractable functions. The basic idea of the decomposition method is to compute the infimum and supremum of $h(g(\mathbf{x}, \mathbf{w}))$ on $G_{t-1}^{(m)}(\mathbf{x})$, and then to compute $\text{lcb}_{t-1}^{(m)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x})$ by taking $\tilde{\rho}(\cdot)$ to these. Calculated values for several risk measures are given in Appendix A. Note that many widely used risk measures such as the Bayes risk, value-at-risk and standard deviation can be calculated by the decomposition method. In the sampling method, we generate S sample paths $f_1^{(m)}(\mathbf{x}, \mathbf{w}), \dots, f_S^{(m)}(\mathbf{x}, \mathbf{w})$ of $f^{(m)}(\mathbf{x}, \mathbf{w})$ independently from the GP posterior and compute $\text{lcb}_{t-1}^{(m)}(\mathbf{x}) = \min_{j \in [S], f_j^{(m)}(\mathbf{x}, \mathbf{w}) \in G_{t-1}^{(m)}(\mathbf{x})} \rho^{(m)}(f_j^{(m)}(\mathbf{x}, \mathbf{w}))$ and $\text{ucb}_{t-1}^{(m)}(\mathbf{x}) = \max_{j \in [S], f_j^{(m)}(\mathbf{x}, \mathbf{w}) \in G_{t-1}^{(m)}(\mathbf{x})} \rho^{(m)}(f_j^{(m)}(\mathbf{x}, \mathbf{w}))$. However, in all cases of generalized, decomposition and sampling methods, there is a case that (2) is not satisfied due to approximation errors that may occur in practice, e.g., approximation errors in the expected value computation or insufficient approximation due to a small number of sample paths. These problems

are discussed in Section 4.

3.2 Pareto Front Estimation

For any input $\mathbf{x} \in \mathcal{X}$ and subset $E \subset \mathcal{X}$, we define $\mathbf{LCB}_{t-1}(\mathbf{x}) = (\text{lcb}_{t-1}^{(1)}(\mathbf{x}), \dots, \text{lcb}_{t-1}^{(M)}(\mathbf{x}))$, $\mathbf{UCB}_{t-1}(\mathbf{x}) = (\text{ucb}_{t-1}^{(1)}(\mathbf{x}), \dots, \text{ucb}_{t-1}^{(M)}(\mathbf{x}))$ and $\mathbf{LCB}_{t-1}(E) = \{\mathbf{LCB}_{t-1}(\mathbf{x}) \mid \mathbf{x} \in E\}$. The estimated Pareto solution set $\hat{\Pi}_{t-1} \subset \mathcal{X}$ for the design variables is then defined as follows: $\hat{\Pi}_{t-1} = \{\mathbf{x} \in \mathcal{X} \mid \mathbf{LCB}_{t-1}(\mathbf{x}) \in \text{Par}(\mathbf{LCB}_{t-1}(\mathcal{X}))\}$. Figure 1 (a) shows a conceptual diagram of $\mathbf{LCB}_{t-1}(\mathbf{x})$ and $\mathbf{UCB}_{t-1}(\mathbf{x})$, and (b) shows a conceptual diagram of $\text{Par}(\mathbf{LCB}_{t-1}(\mathcal{X}))$ and $\hat{\Pi}_{t-1}$. Here, in order to actually compute $\hat{\Pi}_{t-1}$, we need to compute the PF defined by $\mathbf{LCB}_{t-1}(\mathbf{x})$. However, if \mathcal{X} is an infinite set, then $\hat{\Pi}_{t-1}$ may also be an infinite set. In this case, since the exact calculation of $\hat{\Pi}_{t-1}$ is difficult, it is necessary to make a finite approximation using an approximation solver such as NSGA-II (Deb et al., 2002). The effects on this finite approximation are discussed in Section 4.

3.3 Acquisition Function

We propose an AF for determining the next point to be evaluated. First, for each point $\mathbf{a} \in \mathbb{R}^m$ and subset $B \subset \mathbb{R}^m$, we denote the quasi distance between them as $\text{dist}(\mathbf{a}, B) = \min_{\mathbf{b} \in B} d_\infty(\mathbf{a}, \mathbf{b})$, where $d_\infty(\mathbf{a}, \mathbf{b})$ denotes the metric function given by $d_\infty(\mathbf{a}, \mathbf{b}) = \max\{|a_1 - b_1|, \dots, |a_m - b_m|\}$. Using this, we define AF $a_t^{(\mathcal{X})}(\mathbf{x})$ for $\mathbf{x} \in \mathcal{X}$ as $a_t^{(\mathcal{X})}(\mathbf{x}) = \text{dist}(\mathbf{UCB}_t(\mathbf{x}), \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t)))$. Then, the next design variable, \mathbf{x}_{t+1} , to be evaluated is selected by $\mathbf{x}_{t+1} = \text{argmax}_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x})$. Hence, the value of $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1})$ is equal to the following maximin distance: $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) = \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{b} \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{UCB}_t(\mathbf{x}), \mathbf{b})$. Figure 1 (c) shows a conceptual diagram of the AF $a_t^{(\mathcal{X})}(\mathbf{x})$. The value of $a_t^{(\mathcal{X})}(\mathbf{x})$ can be computed analytically using the following lemma when $\hat{\Pi}_t$ is finite:

Lemma 3.1. Let $\mathbf{UCB}_t(\mathbf{x}) = (u_1, \dots, u_M)$ and $\mathbf{LCB}_t(\hat{\Pi}_t) = \{(l_1^{(i)}, \dots, l_M^{(i)}) \mid 1 \leq i \leq k\}$. Then, $a_t^{(\mathcal{X})}(\mathbf{x})$ can be computed by $a_t^{(\mathcal{X})}(\mathbf{x}) = \max\{\tilde{a}_t(\mathbf{x}), 0\}$, where $\tilde{a}_t(\mathbf{x}) = \min_{1 \leq i \leq k} \max\{u_1 - l_1^{(i)}, \dots, u_M - l_M^{(i)}\}$.

The proposed AF is based on the bounding box as well as existing bounding box-based AFs (Iwazaki et al., 2021b; Zuluaga et al., 2016; Belakaria et al., 2020), but differs in the following points. Most of the existing methods focus only on reducing the size of the non-dominated bounding box¹ (e.g., diagonal length and

¹The bounding box $B_t(\mathbf{x})$ with $\mathbf{UCB}_t(\mathbf{x}) \notin \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))$.

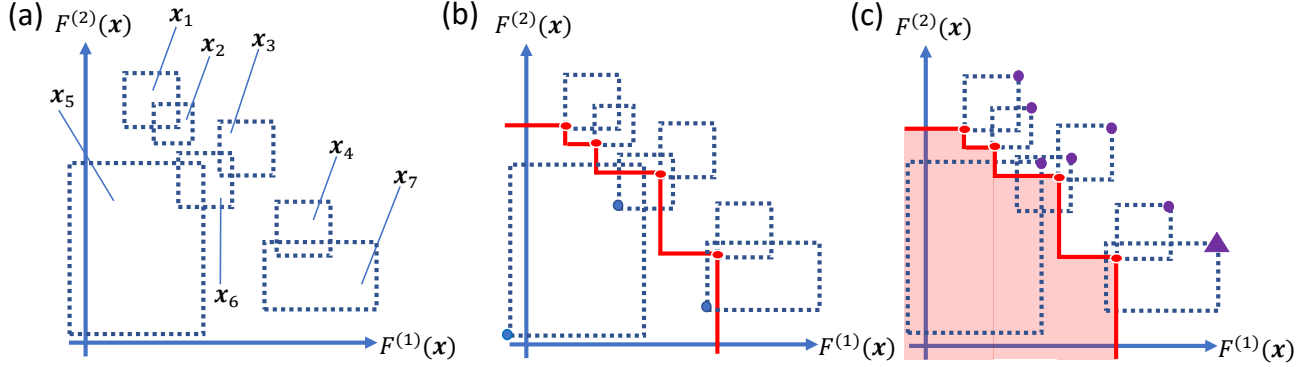


Figure 1: Conceptual diagrams of $\mathbf{LCB}_t(\mathbf{x})$, $\mathbf{UCB}_t(\mathbf{x})$, $\text{Par}(\mathbf{LCB}_t(\mathcal{X}))$, $\hat{\Pi}_t$ and AFs for seven input points $\mathbf{x}_1, \dots, \mathbf{x}_7$. At each point \mathbf{x} in the left figure (a), $\mathbf{LCB}_t(\mathbf{x})$ and $\mathbf{UCB}_t(\mathbf{x})$ indicate the lower left point and the upper right point of the dashed rectangular region, respectively. In (b), the PF (red line) computed using each $\mathbf{LCB}_t(\mathbf{x})$ is $\text{Par}(\mathbf{LCB}_t(\mathcal{X}))$, and because it is constructed by $\mathbf{LCB}_t(\mathbf{x}_1), \mathbf{LCB}_t(\mathbf{x}_2), \mathbf{LCB}_t(\mathbf{x}_3), \mathbf{LCB}_t(\mathbf{x}_4)$, $\hat{\Pi}_t$ is given by $\hat{\Pi}_t = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4\}$. In (c), the light red region indicates $\text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))$, the region dominated by the red points ($\mathbf{LCB}_t(\hat{\Pi}_t)$), and $a_t^{(\mathcal{X})}(\mathbf{x})$ is the closeness between the light red region and $\mathbf{UCB}_t(\mathbf{x})$ (purple point). The furthest point is represented by the purple triangle, $\mathbf{UCB}_t(\mathbf{x}_7)$. Thus, the next design variable to be evaluated is \mathbf{x}_7 .

hypervolume), and therefore do not aim to improve the estimated PF (size-based AFs choose x_5 in Fig. 1, but the room for improvement is small). Hence, these AFs focus on exploration. In contrast, the proposed AF focuses on the non-dominated bounding box with the largest maximin distance to the estimated PF. In this sense, the proposed AF focuses on exploration, but also exploitation.

Next, we consider the choice of the environment variable \mathbf{w}_{t+1} . The variable \mathbf{w}_{t+1} should be determined based on the uncertainty of the chosen bounding box $B_t(\mathbf{x}_{t+1})$. We define the uncertainty of $B_t(\mathbf{x}_{t+1})$ by the maximum length of each edge $\|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty$. In many risk measures including Bayes risk, the following inequality holds:

$$\|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty \leq q(\zeta_{t+1}), \quad (3)$$

where $q(\cdot) : [0, \infty) \rightarrow [0, \infty)$ is a strictly increasing function defined by risk measures and satisfies $q(0) = 0$, and $\zeta_{t+1} = \max_{\mathbf{w} \in \Omega} \sum_{m=1}^M 2\beta_{m,t+1}^{1/2} \sigma_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w})$. Then, we choose \mathbf{w}_{t+1} based on (3). The next environmental variable, \mathbf{w}_{t+1} , to be evaluated is selected by $\mathbf{w}_{t+1} = \text{argmax}_{\mathbf{w} \in \Omega} a_t^{(\Omega)}(\mathbf{w})$, where $a_t^{(\Omega)}(\mathbf{w}) = \sum_{m=1}^M 2\beta_{m,t+1}^{1/2} \sigma_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w})$.

3.4 Stopping Condition

We describe the stopping conditions of the proposed algorithm. From Fig. 1 (c), AF $a_t^{(\mathcal{X})}(\mathbf{x})$ represents the closeness of the pessimistic PF and the optimistic predictive value of $\mathbf{F}(\mathbf{x})$. That is, if this value is suf-

ficiently small, there is little room for improvement in the PF; therefore, it is reasonable to use it as the stopping condition. Let $\epsilon > 0$ be a predetermined desired accuracy parameter. Then the algorithm is terminated if $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq \epsilon$ is satisfied. The pseudocode of the proposed algorithm is described in Algorithm 1.

4 THEORETICAL ANALYSIS

In this section, we give the theorems for the accuracy and termination of the proposed algorithm. The details of the proofs are presented in Appendix B. First, we quantify the goodness of the estimated $\hat{\Pi}_t$. If $\hat{\Pi}_t$ is a good estimate, the following two indicators defined by $\hat{\Pi}_t$ should be small: $I_t^{(i)} = \max_{\mathbf{y} \in Z^*} \text{dist}(\mathbf{y}, \text{Par}(\mathbf{F}(\hat{\Pi}_t)))$, $I_t^{(ii)} = \max_{\mathbf{y} \in \mathbf{F}(\hat{\Pi}_t)} \text{dist}(\mathbf{y}, Z^*)$. Here, $I_t^{(i)}$ and $I_t^{(ii)}$ have similar meanings as recall and precision in the classification problem, respectively. For example, if $\hat{\Pi}_t$ is estimated as $\hat{\Pi}_t = \mathcal{X}$, $\hat{\Pi}_t$ contains all of true Pareto optimal design variables. In this case, since $\text{Par}(\mathbf{F}(\hat{\Pi}_t)) = \text{Par}(\mathbf{F}(\mathcal{X})) = Z^*$, $I_t^{(i)} = 0$. Similarly, when $\hat{\Pi}_t$ is estimated as $\hat{\Pi}_t = \{\mathbf{x}_1^*\}$, where \mathbf{x}_1^* is one of true Pareto optimal design variables, $\hat{\Pi}_t$ does not have unnecessary points, and $I_t^{(ii)} = 0$. As with recall and precision in ordinary classification problems, over (under)-estimation makes $I_t^{(ii)}$ ($I_t^{(i)}$) larger. For this reason, we define the *inference discrepancy* $I_t = \max\{I_t^{(i)}, I_t^{(ii)}\}$ for $\hat{\Pi}_t$ as the goodness measure. Next, in order to show the theoretical validity of the proposed algorithm, we introduce the maximum information gain

Algorithm 1 Bounding box-based MOBO of general risk measures

Require: GP priors $\mathcal{GP}(0, k^{(m)})$, tradeoff parameters $\{\beta_{m,t}\}_{t \geq 0}$, $m \in [M]$, accuracy parameter $\epsilon > 0$
for $t = 0, 1, 2, \dots$ **do**
 Compute $Q_t^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ for each $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$
 Compute $Q_t^{(F^{(m)})}(\mathbf{x})$ for each $\mathbf{x} \in \mathcal{X}$ by the generalized, decomposition or sampling method
 Compute $B_t(\mathbf{x}) = Q_t^{(F^{(1)})}(\mathbf{x}) \times \dots \times Q_t^{(F^{(M)})}(\mathbf{x})$ for each $\mathbf{x} \in \mathcal{X}$
 Estimate $\hat{\Pi}_t$ by $B_t(\mathbf{x})$
 Select the next evaluation point \mathbf{x}_{t+1} by $a_t^{(\mathcal{X})}(\mathbf{x})$
 if $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq \epsilon$ **then**
 break
 end if
 Select the next evaluation point \mathbf{w}_{t+1} by $a_t^{(\Omega)}(\mathbf{w})$
 Observe $y_{t+1}^{(m)} = f^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}) + \varepsilon_{t+1}^{(m)}$ at the point $(\mathbf{x}_{t+1}, \mathbf{w}_{t+1})$ for all $m \in [M]$
 Update GPs by adding observed points
end for
Ensure: Return $\hat{\Pi}_t$ as the estimated set of design variables

$\kappa_t^{(m)}$. This indicator is frequently used in theoretical analysis in the context of GP-based BOs and can be expressed as $\kappa_t^{(m)} = 2^{-1} \max_{(\tilde{\mathbf{x}}_1, \tilde{\mathbf{w}}_1), \dots, (\tilde{\mathbf{x}}_t, \tilde{\mathbf{w}}_t)} \log \det(\mathbf{I}_t + \zeta_m^{-2} \tilde{\mathbf{K}}_t^{(m)})$, where \mathbf{I}_t is the $t \times t$ identity matrix, and $\tilde{\mathbf{K}}_t^{(m)}$ is the $t \times t$ matrix whose (j, k) -th element is $k^{(m)}((\tilde{\mathbf{x}}_j, \tilde{\mathbf{w}}_j), (\tilde{\mathbf{x}}_k, \tilde{\mathbf{w}}_k))$. It is known that the order of $\kappa_t^{(m)}$ with respect to commonly used kernels such as linear, Gaussian and Matérn kernels is sublinear under mild conditions (see, e.g., Theorem 5 in Srinivas et al. (2010)). Then, the following theorem holds:

Lemma 4.1 (Theorem 3.11 in Abbasi-Yadkori (2013)). Suppose that the regularity assumption holds. Let $\delta \in (0, 1)$, and define $\beta_{m,t}^{1/2} = B_m + (2(\kappa_t^{(m)} + \log(M/\delta)))^{1/2}$. Then, with probability at least $1 - \delta$, the following inequality holds for any $t \geq 1$, $m \in [M]$ and $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$: $|f^{(m)}(\mathbf{x}, \mathbf{w}) - \mu_{t-1}^{(m)}(\mathbf{x}, \mathbf{w})| \leq \beta_{m,t}^{1/2} \sigma_{t-1}^{(m)}(\mathbf{x}, \mathbf{w})$.

Using this, we give the following theorems for the inference discrepancy, stopping condition and $q(a)$:

Theorem 4.1. Suppose that the assumption of Lemma 4.1 and the inequality (2) hold. Let $t \geq 0$, $m \in [M]$, $\delta \in (0, 1)$, and let $\beta_{m,t+1}^{1/2}$ be defined as in Lemma 4.1. In addition, let $\epsilon > 0$ be a predetermined desired accuracy parameter. Then, with probability at least $1 - \delta$, the inequality $I_t \leq a_t^{(\mathcal{X})}(\mathbf{x}_{t+1})$ holds for any $t \geq 0$ and \mathbf{x}_{t+1} . Therefore, if the stopping condition satisfies at T iterations, the inference discrepancy I_T

satisfies $I_T \leq \epsilon$ with probability at least $1 - \delta$.

Theorem 4.2. Suppose that the assumption in Theorem 4.1 holds. Let $q : [0, \infty) \rightarrow [0, \infty)$ be a strictly increasing function satisfying $q(0) = 0$ and (3). Also let $s_t = (\sum_{m=1}^M C_m \beta_{m,t+1} \kappa_{t+1}^{(m)})^{1/2} (t+1)^{-1/2}$, where $C_m = 8M(\log(1 + \zeta_m^{-2}))^{-1}$. Then, the inequality $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1}) \leq q(s_t)$ holds for any $t \geq 0$ and some $\hat{t} \leq t$. Therefore, the algorithm terminates after at most T iterations, where T is the smallest positive integer satisfying $q(s_T) \leq \epsilon$.

Theorem 4.3. Suppose that the assumption in Theorem 4.1 holds. Also assume that there exist strictly increasing functions $q^{(m)} : [0, \infty) \rightarrow [0, \infty)$ satisfying $q^{(m)}(0) = 0$ and $|\text{ucb}_t^{(m)}(\mathbf{x}_{t+1}) - \text{lcb}_t^{(m)}(\mathbf{x}_{t+1})| \leq q^{(m)}(\tilde{s}_t)$ for any $t \geq 0$, $m \in [M]$, and $\mathbf{x}_{t+1} \in \mathcal{X}$, where $\tilde{s}_t = \max_{\mathbf{w} \in \Omega} 2\beta_{m,t+1}^{1/2} \sigma_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w})$. Then, $q(a) \equiv \max_{m \in [M]} q^{(m)}(a)$ is the strictly increasing function and satisfies $q(0) = 0$ and (3).

Specific forms of $q^{(m)}(a)$ for commonly used risk measures are described in Appendix A.

4.1 Theoretical Error Analysis

In this subsection, we give an extension of Theorem 4.1 and 4.2 when approximation errors are included in the algorithm. In practice, the algorithm includes the following approximation errors: (i) Errors in the computation of $\text{lcb}_{t-1}^{(m)}(\mathbf{x})$, $\text{ucb}_{t-1}^{(m)}(\mathbf{x})$, (ii) errors in computing $\hat{\Pi}_{t-1}$ due to the finite approximation of the estimated PF, and (iii) computational errors in maximizing the AFs $a_{t-1}^{(\mathcal{X})}(\mathbf{x})$ and $a_{t-1}^{(\Omega)}(\mathbf{w})$. Let $\epsilon_{\text{lcb}}, \epsilon_{\text{ucb}}, \epsilon_{\text{PF}}, \epsilon_{\mathcal{X}}, \epsilon_{\Omega}$ be non-negative error parameters that represent the errors in these approximations, respectively. We consider the case that the following four error inequalities hold for any $t \geq 0$, $m \in [M]$, $\mathbf{x}, \mathbf{x}_{t+1} \in \mathcal{X}$, $\mathbf{w}_{t+1} \in \Omega$ and $g(\mathbf{x}, \mathbf{w}) \in G_t^{(m)}(\mathbf{x})$: $\text{lcb}_t^{(m)}(\mathbf{x}) - \epsilon_{\text{lcb}} \leq \rho^{(m)}(g(\mathbf{x}, \mathbf{w})) \leq \text{ucb}_t^{(m)}(\mathbf{x}) + \epsilon_{\text{ucb}}$; $\max_{\mathbf{y} \in \text{Par}(\mathbf{LCB}_t(\hat{\Pi}_t))} \text{dist}(\mathbf{y}, \text{Par}(\mathbf{LCB}_t(\mathcal{X}))) \leq \epsilon_{\text{PF}}$; $\max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}) - a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq \epsilon_{\mathcal{X}}$; $\max_{\mathbf{w} \in \Omega} a_t^{(\Omega)}(\mathbf{w}) - a_t^{(\Omega)}(\mathbf{w}_{t+1}) \leq \epsilon_{\Omega}$. These inequalities imply that the difference between the desired and actual calculated values is less than the error parameter. In this case, a desirable property is that these error parameters simply add to the inequalities in Theorem 4.1 and 4.2. Here, we must emphasize that it is not obvious whether the above is true or not. This is because the inference discrepancy is defined by the combination of operations such as the computation of bounding box and the estimation of $\hat{\Pi}_{t-1}$, and it is not obvious how the approximation error affects the inequality. The next theorem shows how these approximation errors affect the inequalities:

Theorem 4.4. Suppose that the assumption in Lemma 4.1 holds. Let $t \geq 0$, $m \in [M]$, $\delta \in (0, 1)$, and let $\beta_{m,t+1}^{1/2}$ be defined as in Lemma 4.1. In addition, let $\epsilon > 0$ be a predetermined desired accuracy parameter. Moreover, let $\epsilon_{\text{lcB}}, \epsilon_{\text{ucB}}, \epsilon_{\text{PF}}, \epsilon_{\mathcal{X}}, \epsilon_{\Omega}$ be non-negative error parameters satisfying the error inequalities. Then, with probability at least $1 - \delta$, the inequality $I_t \leq a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) + \epsilon_{\text{lcB}} + \epsilon_{\text{ucB}} + \epsilon_{\mathcal{X}}$ holds for any $t \geq 0$ and \mathbf{x}_{t+1} . Therefore, if the stopping condition satisfies at T iterations, the inference discrepancy I_T satisfies $I_T \leq \epsilon + \epsilon_{\text{lcB}} + \epsilon_{\text{ucB}} + \epsilon_{\mathcal{X}}$ with probability at least $1 - \delta$.

Theorem 4.5. Suppose that the assumption in Theorem 4.4 holds. Let $q : [0, \infty) \rightarrow [0, \infty)$ be a strictly increasing function satisfying $q(0) = 0$ and (3). Then, the inequality $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1}) \leq \epsilon_{\text{PF}} + q(\epsilon_{\Omega} + s_t)$ holds for any $t \geq 0$ and some $\hat{t} \leq t$, where C_m and s_t are given by Theorem 4.2. Therefore, the algorithm terminates after at most T iterations, where T is the smallest positive integer satisfying $\epsilon_{\text{PF}} + q(\epsilon_{\Omega} + s_T) \leq \epsilon$.

Note that for Theorem 4.5, the integer T satisfying the theorem’s last inequality does not always exist. However, the left hand side in this inequality is merely an upper bound of $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1})$. Thus, in some cases the actual value of $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1})$ satisfies $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1}) \leq \epsilon$ and the stopping condition is satisfied.

5 NUMERICAL EXPERIMENTS

In this section, we confirm the performance of the proposed method using synthetic functions and real-world docking simulations. For all experiments, we used Gaussian kernels and GP models. Experimental details and additional experiments are described in Appendix C.

5.1 Synthetic Function

We confirm the performance of the proposed method through synthetic functions. Although the proposed method is constructed under the presence of IU, the algorithm itself can be applied even when there is no IU. Therefore, in the synthetic function experiments, we compared the proposed method with existing MOBO methods without (with) IU.

In the experiments under no IU, the input space \mathcal{X} was a set of grid points divided into $[-5, 5] \times [-5, 5]$ equally spaced at 50×50 . For black-box functions, we used Booth, Matyas, Himmelblau’s and McCormick benchmark functions. We performed a two-objective optimization using the first two and a four-objective optimization using all four. As evaluation indicators, we used the simple Pareto hypervolume (PHV) re-

gret, which is a commonly used indicator in the context of MOBOs, and inference discrepancy. As AFs, we considered the random sampling (Random), uncertainty sampling (US), EHVI (Emmerich and Klöckner, 2008), EMmI (Svenson and Santner, 2010), ePAL (Zuluaga et al., 2016), ParEGO (Knowles, 2006), PFES (Suzuki et al., 2020) and proposed AF (Proposed). We also compared the commonly used evolutionary computation-based method NSGA-II (Deb et al., 2002). Under this setup, one initial point was taken at random and the algorithm was run until the number of iterations reached 300. This simulation repeated 100 times and the average simple PHV regret and inference discrepancy at each iteration were calculated. From the top of Fig. 2, it can be confirmed that the performance at the end of 300 iterations is comparable or better than the existing methods except for the simple PHV regret in the four-objective setting. In particular, the proposed method significantly outperforms other methods for inference discrepancy in the four-objective setting after about 180 iterations.

In the experiment under IU, the input space $\mathcal{X} \times \Omega$ was a compact subset, and we considered infinite and finite set settings. We set $\mathcal{X} \times \Omega = [0.25, 0.75]^2 \times [-0.25, 0.25]^2$ in the infinite set setting. In the finite setting, $\mathcal{X} \times \Omega$ was a set of grid points divided into $[-1, 1]^3 \times [-1, 1]^3$ equally spaced at $7^3 \times 7^3 = 117649$. The black-box function in the infinite setting was used the ZDT1 benchmark function $\mathbf{ZDT1}(\mathbf{a}) \in \mathbb{R}^2$ with a two-dimensional input \mathbf{a} , and the environmental variable \mathbf{w} was used as the input noise for \mathbf{x} . Thus, our considered black-box function was defined by $\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})$. We assumed \mathbf{w} was the uniform distribution on Ω and used the Bayes risk $\mathbb{E}[\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})]$. On the other hand, the black-box function in the finite setting was used the six-dimensional Rosenbrock function $f(w_1, w_2, x_1, x_2, x_3, w_3) \in \mathbb{R}$. We assume that \mathbf{w} was a discretized normal distribution on Ω . As risk measures, we used the expectation and negative standard deviation. As comparison methods, we considered the BQ-based method (Qing et al., 2023), MVA-based method (Iwazaki et al., 2021b) and SABBa-based method (Rivier and Congedo, 2022). Furthermore, four naive methods, Naive-random, Naive-US, Naive-EMmI and Naive-ePAL, were used for comparison. In the naive methods, \mathbf{w} was generated five times from the same \mathbf{x} in one iteration t , and the sample mean and the negative square root of the sample variance of the black-box function values were calculated. By using \mathbf{x} and these values, the experiments in naive four methods were performed as a usual MOBO. The name after “Naive-” means the name of the used AF. We used the inference discrepancy as the evaluation indicator. Under this setup, one initial point was taken at random and the algorithm was run until the num-

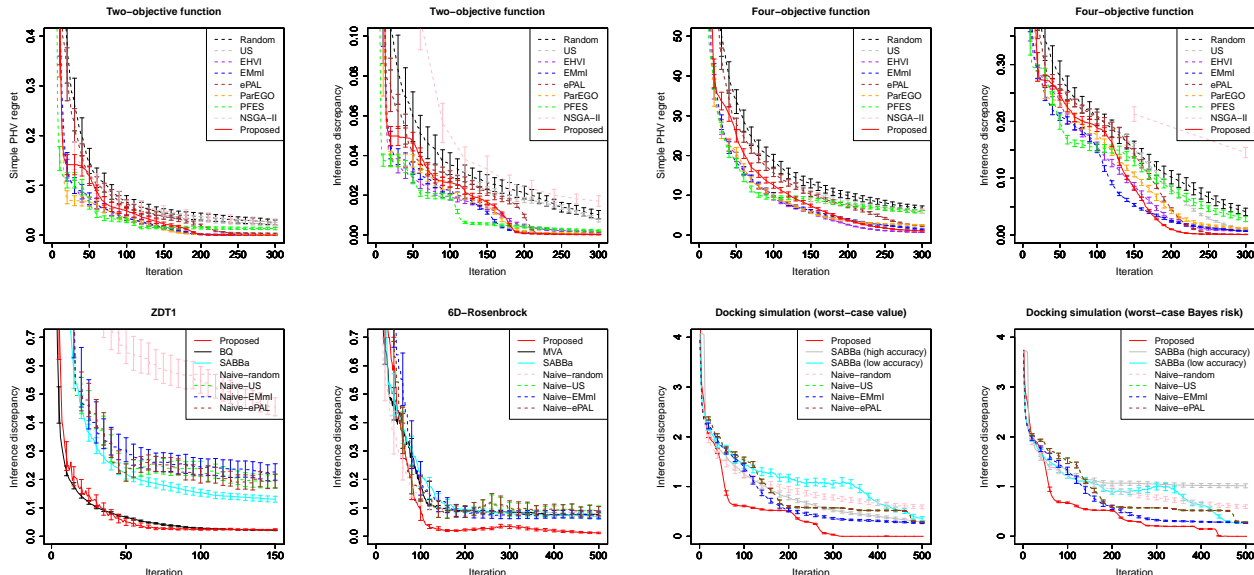


Figure 2: Comparison with MOBO methods. Solid (and dashed) lines are averages of the evaluation measures (simple PHV regret and inference discrepancy) for each iteration in 100, 920 or 429 trials. Each error bar length represents the six times the standard error. In the top row, the left two columns represent the two-objective setup and the right two columns represent the four-objective setup. In the bottom row, the left two columns respectively represent the ZDT1 and six-dimensional Rosenbrock setups in the synthetic experiment, and the right two columns respectively represent WC and WCBR setups in the real-world docking simulation.

ber of iterations reached 150 and 500. This simulation repeated 100 times and the average inference discrepancy at each iteration were calculated. From the bottom of Fig. 2, it can be confirmed that the proposed method achieves the same or better performance as the existing methods. In particular, the results are comparable to those of BQ, which is a limited method applicable only to the Bayes risk case.

5.2 Real-world Docking Simulation

In this subsection, we applied the proposed method to docking simulations for real-world chemical compounds. The purpose of this simulation is to evaluate the inhibitory performance of candidate compounds on two specific sites of the target protein “KAT1”, the structure of this protein is available at <https://pdj.org/mine/summary/6v1x>, and to enumerate the Pareto optimal compounds in the presence of structural uncertainty (isomers). We used the software suite *Schrödinger* (Schrödinger LLC, 2021) to calculate docking scores and explanatory variables in the compounds. Each compound C_i may have an isomer S_{ij} , and in this simulation the maximum number of isomers was limited to 10. For each i , we computed a 51-dimensional isomer-independent design variables \mathbf{x}_i and a 51-dimensional environment variable \mathbf{w}_{ij} that can vary with isomers, using explanatory variables of

(C_i, S_{ij}) . Thus, the black-box functions, the docking scores in two sites, can be expressed as $f^{(1)}(\mathbf{x}_i, \mathbf{w}_{ij})$ and $f^{(2)}(\mathbf{x}_i, \mathbf{w}_{ij})$, respectively. We emphasize that the number of isomers N_i was not same for all i . As risk measures for C_i , we considered the worst-case (WC) and worst-case Bayes risk (WCBR). For each compound, WC is defined as the minimum docking scores, and WCBR is defined as the minimum weighted average of docking scores in predefined candidate weights. The total number of compounds was 429, and the total number of data including isomers was 920. We compared the SABBa, Proposed and naive four methods. In the SABBa method, we considered two different accuracy parameter settings, a high accuracy model and a low accuracy model. In addition, in the naive four methods, we calculated docking scores for all isomers in the compound C_i at iteration t and determined the exact risk values. In this experiment, the observation noise was zero. Under this setup, one initial point was taken from the data and the algorithm was run until the number of iterations reached 500. In SABBa and Proposed, by changing the initial point, this simulation repeated 920 times. Similarly, in naive methods, by changing the initial compound, this simulation repeated 429 times. We calculated the average inference discrepancy at each iteration. From the bottom in Fig.2, we can confirm that the proposed method is superior to other methods. In addition, only the pro-

Table 2: Computational time (second) for obtaining $(\mathbf{x}_{t+1}, \mathbf{w}_{t+1})$ in the proposed method

Experimental setup	Average	Standard deviation
Two-objective optimization without IU	0.93	0.60
Four-objective optimization without IU	2.03	1.26
ZDT1 with IU	5.60	2.25
6D-Rosenbrock with IU	0.73	0.16
Docking simulation (WC)	0.0161	0.0030
Docking simulation (WCBR)	0.0236	0.0089

posed method correctly identifies the true PF at the end of 500 iterations for all risk measures at all 920 different initial points. Specifically, after 425 iterations for WC and 465 iterations for WCBR, the true PF is identified for all 920 different initial points. Therefore, compared to the exhaustive search, the number of iterations required to find the true PF can be reduced to about half. Thus, the sample efficient decision making was achieved in our motivating example.

5.3 Computational Time Experiments

Here, we measured the computational time required to obtain $(\mathbf{x}_{t+1}, \mathbf{w}_{t+1})$ for each iteration t in the proposed method, where the time required for modeling GP is not included in the measurement time because all MOBO methods, including the comparison methods, perform GP modeling. We measured the computational time for each iteration in a single trial and calculated its average and standard deviation for the iterations in the experiments performed in Section 5.1 and 5.2. From Table 2, the computational time for AFs in the proposed method is acceptable even in a 6D-Rosenbrock setting with more than 100,000 candidate points. In contrast, the reason why the computational time in the ZDT1 setting is larger than the others is due to the finite approximation of PF by the NSGA-II algorithm. Therefore, the computational time can be reduced if the population size n_p or number of generations n_g is reduced. Nevertheless, the computational time is acceptable even for our experimental setup, $n_p = n_g = 50$.

6 CONCLUSION

In this study, we proposed the efficient MOBO method for identifying the PF defined by general risk measures. The proposed method can work with (and without) IU and has theoretical guarantees. In various risk measures, we proved that the algorithm can return an arbitrary-accurate solution with high probability in a finite number of iterations. Through numerical experiments, we confirmed that the proposed method outperforms existing methods. Moreover, from the real-

world docking simulation that is our motivating example, we confirmed that by using the proposed method, the number of function evaluations required to identify the true PF has been successfully reduced to about half that of the exhaustive search.

The proposed method has two limitations. First, although we have given a theoretical analysis of how approximation errors in the proposed method affect the final results, we have not mentioned an estimate of the degree of approximation errors in the first place. Thus, as a practical matter, it is difficult to estimate the final accuracy of the proposed method considering the approximation error in advance. Second, the proposed method does not consider constraint conditions. In actual applications, Pareto optimization under some constraints is often considered. We can apply the proposed method to this setting directly by designing a HPBB for the constraint function. However, it is not obvious whether theoretical results derived in this study can be derived in the same way in such a case. The above problems are left for future work.

Acknowledgements

This work was partially supported by JSPS KAKENHI (JP20H00601, JP23K16943, JP23K19967), JST ACT-X (JPMJAX23CD), JST CREST (JPMJCR21D3, JPMJCR22N2), JST Moonshot R&D (JPMJMS2033-05), JST AIP Acceleration Research (JPMJCR21U2), NEDO (JPNP18002, JPNP20006) and RIKEN Center for Advanced Intelligence Project.

References

- Abbasi-Yadkori, Y. (2013). Online learning for linearly parametrized control problems.
- Belakaria, S., Deshwal, A., Jayakodi, N. K., and Doppa, J. R. (2020). Uncertainty-aware search framework for multi-objective bayesian optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 10044–10052.
- Beland, J. J. and Nair, P. B. (2017). Bayesian optimization under uncertainty. In *NIPS BayesOpt 2017 workshop*, volume 2.

- Daulton, S., Cakmak, S., Balandat, M., Osborne, M. A., Zhou, E., and Bakshy, E. (2022). Robust multi-objective bayesian optimization under input noise. In *International Conference on Machine Learning*, pages 4831–4866. PMLR.
- Deb, K. and Gupta, H. (2005). Searching for robust pareto-optimal solutions in multi-objective optimization. In *International conference on evolutionary multi-criterion optimization*, pages 150–164. Springer.
- Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation*, 6(2):182–197.
- Emmerich, M. and Klinkenberg, J.-w. (2008). The computation of the expected improvement in dominated hypervolume of pareto front approximations. *Rapport technique, Leiden University*, 34:7–3.
- Inatsu, Y., Iwazaki, S., and Takeuchi, I. (2021). Active learning for distributionally robust level-set estimation. In *International Conference on Machine Learning*, pages 4574–4584. PMLR.
- Inatsu, Y., Takeno, S., Karasuyama, M., and Takeuchi, I. (2022). Bayesian optimization for distributionally robust chance-constrained problem. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S., editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 9602–9621. PMLR.
- Iwazaki, S., Inatsu, Y., and Takeuchi, I. (2021a). Bayesian quadrature optimization for probability threshold robustness measure. *Neural Computation*, 33(12):3413–3466.
- Iwazaki, S., Inatsu, Y., and Takeuchi, I. (2021b). Mean-variance analysis in bayesian optimization under uncertainty. In Banerjee, A. and Fukumizu, K., editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 973–981. PMLR.
- Kirschner, J., Bogunovic, I., Jegelka, S., and Krause, A. (2020). Distributionally robust bayesian optimization. In Chiappa, S. and Calandra, R., editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2174–2184. PMLR.
- Kirschner, J. and Krause, A. (2018). Information directed sampling and bandits with heteroscedastic noise. In Bubeck, S., Perchet, V., and Rigollet, P., editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 358–384. PMLR.
- Knowles, J. (2006). Parego: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE transactions on evolutionary computation*, 10(1):50–66.
- Kusakawa, S., Takeno, S., Inatsu, Y., Kutsukake, K., Iwazaki, S., Nakano, T., Ujihara, T., Karasuyama, M., and Takeuchi, I. (2022). Bayesian optimization for cascade-type multistage processes. *Neural Computation*, 34(12):2408–2431.
- Makarova, A., Usmanova, I., Bogunovic, I., and Krause, A. (2021). Risk-averse heteroscedastic bayesian optimization. In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W., editors, *Advances in Neural Information Processing Systems*, volume 34, pages 17235–17245. Curran Associates, Inc.
- Moćkus, J. (1975). On bayesian methods for seeking the extremum. In *Optimization Techniques IFIP Technical Conference: Novosibirsk, July 1–7, 1974*, pages 400–404. Springer.
- Nguyen, Q. P., Dai, Z., Low, B. K. H., and Jaillet, P. (2021a). Optimizing conditional value-at-risk of black-box functions. *Advances in Neural Information Processing Systems*, 34:4170–4180.
- Nguyen, Q. P., Dai, Z., Low, B. K. H., and Jaillet, P. (2021b). Value-at-risk optimization with gaussian processes. In *International Conference on Machine Learning*, pages 8063–8072. PMLR.
- Qing, J., Couckuyt, I., and Dhaene, T. (2023). A robust multi-objective bayesian optimization framework considering input uncertainty. *Journal of Global Optimization*, 86(3):693–711.
- Rahimi, A. and Recht, B. (2007). Random features for large-scale kernel machines. *Advances in neural information processing systems*, 20.
- Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Rivier, M. and Congedo, P. M. (2022). Surrogate-assisted bounding-box approach applied to constrained multi-objective optimisation under uncertainty. *Reliability Engineering & System Safety*, 217:108039.
- Schrödinger LLC (2021). Schrödinger release 2021-2.
- Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. (2015). Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175.

Srinivas, N., Krause, A., Kakade, S. M., and Seeger, M. W. (2010). Gaussian process optimization in the bandit setting: No regret and experimental design. In Fürnkranz, J. and Joachims, T., editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 1015–1022. Omnipress.

Suzuki, S., Takeno, S., Tamura, T., Shitara, K., and Karasuyama, M. (2020). Multi-objective bayesian optimization using pareto-frontier entropy. In *International Conference on Machine Learning*, pages 9279–9288. PMLR.

Svenson, J. D. and Santner, T. J. (2010). Multiobjective optimization of expensive black-box functions via expected maximin improvement. *The Ohio State University, Columbus, Ohio*, 32.

Takeno, S., Inatsu, Y., and Karasuyama, M. (2023). Randomized Gaussian process upper confidence bound with tighter Bayesian regret bounds. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., editors, *Proceedings of the 40th International Conference on Machine Learning Research*, volume 202 of *Proceedings of Machine Learning Research*, pages 33490–33515. PMLR.

Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.

Wang, Z. and Jegelka, S. (2017). Max-value entropy search for efficient bayesian optimization. In *International Conference on Machine Learning*, pages 3627–3635. PMLR.

Zhou, Q., Jiang, P., Huang, X., Zhang, F., and Zhou, T. (2018). A multi-objective robust optimization approach based on gaussian process model. *Structural and Multidisciplinary Optimization*, 57:213–233.

Zuluaga, M., Krause, A., and Püschel, M. (2016). ε -pal: an active learning approach to the multi-objective optimization problem. *The Journal of Machine Learning Research*, 17(1):3619–3650.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes. See Section 2 and 3.]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes. Time and space complexities in GP models are generally $O(n^3)$ and $O(n^2)$, and theoretical results in sample complexity are shown in Theorem 4.1 and 4.2.]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes. We will submit experimental codes as supplementary materials.]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes. See, e.g., Theorem 4.1.]
 - (b) Complete proofs of all theoretical results. [Yes. We provide them in Appendix B.]
 - (c) Clear explanations of any assumptions. [Yes. See, e.g., Section 2 and theorems’ assumptions.]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes. We provide experimental details and codes in Appendix C and supplemental materials.]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes. See Appendix C.]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes. See Section 5 and the caption in Fig. 2.]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Yes. We downloaded the target protein structure used in our real data experiments from public databases, and the URL is described in Section 5.2.]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]

5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Appendix

This appendix describes the extensions of the proposed method, proofs of all theoretical results, and details of experiments and additional experiments. Note that the symbols used in this appendix are independent of the main body.

A EXTENSION

In this section, we extend the proposed method. We consider the following four extensions:

- The number of black-box functions and the number of risk measures are different.
- The true noise distribution follows some heteroscedastic sub-Gaussian distribution.
- The distribution of \mathbf{w} depends on the design variable \mathbf{x} .
- We consider the uncontrollable setting, that is, \mathbf{w} cannot be controlled even during optimization.

A.1 Extension of Problem Setup

Preliminary Let $f^{(m)} : \mathcal{X} \times \Omega \rightarrow \mathbb{R}$ be an expensive-to-evaluate black-box function, where $m \in [M_f]$ and $M_f \geq 1$. Assume that the set of design variables \mathcal{X} and set of environmental variables Ω are compact and convex. For each design variable $\mathbf{x} \in \mathcal{X}$, the environmental variable \mathbf{w} follows some probability distribution $P_{\mathbf{w}}(\mathbf{x})$, which depends on \mathbf{x} , and takes values in a compact and convex subset $\Omega_{\mathbf{x}} \subset \Omega$. For each iteration t , input $(\mathbf{x}_t, \mathbf{w}_t) \in \mathcal{X} \times \Omega$, and $m \in [M_f]$, the value of the black-box function $f^{(m)}$ is observed with noise as $y_t^{(m)} = f^{(m)}(\mathbf{x}_t, \mathbf{w}_t) + \eta^{(m)}(\mathbf{x}_t, \mathbf{w}_t)$, where $\eta^{(m)}(\mathbf{x}_t, \mathbf{w}_t)$ is zero-mean noise independent across different iteration t , $m \in [M_f]$ and \mathbf{w}_t . In this section, we assume that $\eta^{(m)}(\mathbf{x}_t, \mathbf{w}_t)$ is a sub-Gaussian heteroscedastic noise that depends on $(\mathbf{x}, \mathbf{w}, m)$.

Definition A.1. Let η be a zero-mean real-valued random variable. Then, η is τ -sub-Gaussian if there exists a positive constant τ^2 such that

$$\forall a \in \mathbb{R}, \quad \mathbb{E}[e^{a\eta}] \leq \exp\left(\frac{a^2\tau^2}{2}\right).$$

Commonly used distributions such as Gaussian, Bernoulli and uniform are sub-Gaussian (Vershynin, 2018). We assume that the random variables $\{\mathbf{w}_t, \eta^{(m)}(\mathbf{x}_t, \mathbf{w}_t)\}_{t \geq 1, m \in [M_f]}$ are mutually independent. For \mathbf{w} , we consider the both simulator and uncontrollable settings. Let $\rho^{(m,l)}(f^{(m)}(\mathbf{x}, \mathbf{w})) \equiv F^{(m,l)}(\mathbf{x})$ be a risk measure, where $l \in \{1, \dots, L_m\}$ and $L_1 + \dots + L_{M_f} \equiv L \geq 2$. The purpose of this study is to efficiently identify the PF defined based on $F^{(m,l)}(\mathbf{x})$. For any $\mathbf{x} \in \mathcal{X}$ and $E \subset \mathcal{X}$, let

$$\mathbf{F}(\mathbf{x}) = (F^{(1,1)}(\mathbf{x}), \dots, F^{(1,L_1)}(\mathbf{x}), \dots, F^{(M_f,1)}(\mathbf{x}), \dots, F^{(M_f,L_{M_f})}(\mathbf{x}))$$

and $\mathbf{F}(E) = \{\mathbf{F}(\mathbf{x}) \mid \mathbf{x} \in E\}$. Then, for any $B \subset \mathbb{R}^L$, the dominated region $\text{Dom}(B)$ and PF $\text{Par}(B)$ of B are defined as

$$\text{Dom}(B) = \{\mathbf{s} \in \mathbb{R}^L \mid \exists \mathbf{s}' \in B \text{ s.t. } \mathbf{s} \leq \mathbf{s}'\}, \quad \text{Par}(B) = \partial(\text{Dom}(B)).$$

Let Z^* be our target PF. Then, Z^* can be expressed as

$$Z^* = \text{Par}(\mathbf{F}(\mathcal{X})).$$

Regularity Assumption We introduce a regularity assumption for $f^{(m)}$. For each $m \in [M_f]$, let $k^{(m)} : (\mathcal{X} \times \Omega) \times (\mathcal{X} \times \Omega) \rightarrow \mathbb{R}$ be a positive-definite kernel, where $k^{(m)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}, \mathbf{w})) \leq 1$ for any $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$. Also let $\mathcal{H}(k^{(m)})$ be a reproducing kernel Hilbert space (RKHS) corresponding to $k^{(m)}$. We assume that $f^{(m)}$ is the element of $\mathcal{H}(k^{(m)})$ and has the bounded Hilbert norm $\|f^{(m)}\|_{\mathcal{H}(k^{(m)})} \leq B_m < \infty$. Moreover, we assume that the noise $\eta^{(m)}(\mathbf{x}, \mathbf{w})$ is $\tau(\mathbf{x}, \mathbf{w}, m)$ -sub-Gaussian, where $\tau(\mathbf{x}, \mathbf{w}, m) \equiv \tau_{\mathbf{x}, \mathbf{w}, m}$ satisfies $\tau_{\mathbf{x}, \mathbf{w}, m} \in [\underline{\tau}, \bar{\tau}]$ for some $\bar{\tau} \geq \underline{\tau} > 0$.

Gaussian Process Model We use a GP model for the black-box function $f^{(m)}$. Let $\lambda_1, \dots, \lambda_{M_f}$ be positive numbers. We assume the GP $\mathcal{GP}(0, \tilde{k}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')))$ as the prior of $f^{(m)}$, where $\tilde{k}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}'))$ is given by

$$\tilde{k}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')) = \frac{1}{\lambda_m} k^{(m)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}', \mathbf{w}')).$$

Furthermore, we consider the zero-mean normal distribution with variance $\tau_{\mathbf{x}, \mathbf{w}, m}^2$, as the error distribution in the GP model. For $m \in [M_f]$, given a dataset $\{(\mathbf{x}_i, \mathbf{w}_i, y_i^{(m)})\}_{i=1}^t$, where t is the number of queried instances, the posterior of $f^{(m)}$ is a GP. Then, its posterior mean $\tilde{\mu}_t^{(m)}(\mathbf{x}, \mathbf{w})$ and posterior variance $\tilde{\sigma}_t^{(m)2}(\mathbf{x}, \mathbf{w})$ can be calculated as follows:

$$\begin{aligned} \tilde{\mu}_t^{(m)}(\mathbf{x}, \mathbf{w}) &= \tilde{\mathbf{k}}_t^{(m)}(\mathbf{x}, \mathbf{w})^\top (\tilde{\mathbf{K}}_t^{(m)} + \mathbf{\Sigma}_t^{(m)})^{-1} \mathbf{y}_t^{(m)}, \\ \tilde{\sigma}_t^{(m)2}(\mathbf{x}, \mathbf{w}) &= \tilde{k}^{(m)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}, \mathbf{w})) - \tilde{\mathbf{k}}_t^{(m)}(\mathbf{x}, \mathbf{w})^\top (\tilde{\mathbf{K}}_t^{(m)} + \mathbf{\Sigma}_t^{(m)})^{-1} \tilde{\mathbf{k}}_t^{(m)}(\mathbf{x}, \mathbf{w}), \end{aligned}$$

where $\tilde{\mathbf{k}}_t^{(m)}(\mathbf{x}, \mathbf{w})$ is the t -dimensional vector, whose j -th element is $\tilde{k}^{(m)}((\mathbf{x}, \mathbf{w}), (\mathbf{x}_j, \mathbf{w}_j))$, $\mathbf{y}_t^{(m)} = (y_1^{(m)}, \dots, y_t^{(m)})^\top$, $\mathbf{\Sigma}_t^{(m)}$ is the $t \times t$ diagonal matrix whose (j, j) -th element is $\tau_{\mathbf{x}_t, \mathbf{w}_t, m}^2$, $\tilde{\mathbf{K}}_t^{(m)}$ is the $t \times t$ matrix whose (j, k) -th element is $\tilde{k}^{(m)}((\mathbf{x}_j, \mathbf{w}_j), (\mathbf{x}_k, \mathbf{w}_k))$, with a superscript \top indicating the transpose of vectors or matrices.

A.2 Extension of Proposed Method

Credible Interval and Bounding Box For each input $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$ and $t \geq 1$, the CI of $f^{(m)}(\mathbf{x}, \mathbf{w})$ is denoted by $\tilde{Q}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}) = [\tilde{l}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}), \tilde{u}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})]$, where $\tilde{l}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ and $\tilde{u}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ are given by

$$\begin{aligned} \tilde{l}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}) &= \tilde{\mu}_{t-1}^{(m)}(\mathbf{x}, \mathbf{w}) - \tilde{\beta}_{m,t}^{1/2} \tilde{\sigma}_{t-1}^{(m)}(\mathbf{x}, \mathbf{w}), \\ \tilde{u}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w}) &= \tilde{\mu}_{t-1}^{(m)}(\mathbf{x}, \mathbf{w}) + \tilde{\beta}_{m,t}^{1/2} \tilde{\sigma}_{t-1}^{(m)}(\mathbf{x}, \mathbf{w}). \end{aligned}$$

For $\mathbf{x} \in \mathcal{X}$, $t \geq 1$ and $m \in [M_f]$, we define the set of functions $\tilde{G}_{t-1}^{(m)}(\mathbf{x})$ as

$$\tilde{G}_{t-1}^{(m)}(\mathbf{x}) = \{g(\mathbf{x}, \mathbf{w}) \mid \forall \mathbf{w} \in \Omega, g(\mathbf{x}, \mathbf{w}) \in \tilde{Q}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})\}.$$

Let $\tilde{Q}_{t-1}^{(F^{(m,l)})}(\mathbf{x}) = [\text{lcb}_{t-1}^{(m,l)}(\mathbf{x}), \text{ucb}_{t-1}^{(m,l)}(\mathbf{x})]$ be a CI of $F^{(m,l)}(\mathbf{x})$. Also let $\tilde{B}_{t-1}(\mathbf{x}) = \prod_{m=1}^{M_f} \prod_{l=1}^{L_m} \tilde{Q}_{t-1}^{(F^{(m,l)})}(\mathbf{x})$ be a bounding box of $\mathbf{F}(\mathbf{x})$. Then, when $\tilde{Q}_{t-1}^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ is HPCI, a sufficient condition for $\tilde{Q}_{t-1}^{(F^{(m,l)})}(\mathbf{x})$ to also be HPCI is given as follows:

$$\forall g(\mathbf{x}, \mathbf{w}) \in \tilde{G}_{t-1}^{(m)}(\mathbf{x}), \text{lcb}_{t-1}^{(m,l)}(\mathbf{x}) \leq \rho^{(m,l)}(g(\mathbf{x}, \mathbf{w})) \leq \text{ucb}_{t-1}^{(m,l)}(\mathbf{x}). \quad (4)$$

If (4) holds, then $\tilde{B}_{t-1}(\mathbf{x})$ is also a HPBB. Next, we provide computation methods for $\text{lcb}_{t-1}^{(m,l)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m,l)}(\mathbf{x})$. First, we provide a generalized method for $\text{lcb}_{t-1}^{(m,l)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m,l)}(\mathbf{x})$ to satisfy (4). The $\text{lcb}_{t-1}^{(m,l)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m,l)}(\mathbf{x})$ by the generalized method are calculated with

$$\begin{aligned} \text{lcb}_{t-1}^{(m,l)}(\mathbf{x}) &= \inf_{g(\mathbf{x}, \mathbf{w}) \in \tilde{G}_{t-1}^{(m)}(\mathbf{x})} \rho^{(m,l)}(g(\mathbf{x}, \mathbf{w})), \\ \text{ucb}_{t-1}^{(m,l)}(\mathbf{x}) &= \sup_{g(\mathbf{x}, \mathbf{w}) \in \tilde{G}_{t-1}^{(m)}(\mathbf{x})} \rho^{(m,l)}(g(\mathbf{x}, \mathbf{w})). \end{aligned}$$

The condition (4) holds by using the generalized method, the inf and sup calculations in the generalized method are not always easy. Therefore, we give additional two computation methods for $\text{lcb}_{t-1}^{(m,l)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m,l)}(\mathbf{x})$, the decomposition method and sampling method. Let $\rho(\cdot)$ be a risk measure. In many cases, $\rho(\cdot)$ can be decomposed as $\rho(\cdot) = \tilde{\rho} \circ h(\cdot)$, where $\tilde{\rho}(\cdot)$ and $h(\cdot)$ are respectively monotonic and tractable functions. The basic idea of the decomposition method is to compute the infimum and supremum of $h(g(\mathbf{x}, \mathbf{w}))$ on $\tilde{G}_{t-1}^{(m)}(\mathbf{x})$, and then to compute $\text{lcb}_{t-1}^{(m,l)}(\mathbf{x})$ and $\text{ucb}_{t-1}^{(m,l)}(\mathbf{x})$ by taking $\tilde{\rho}(\cdot)$ to these. Calculated values for several risk measures are

Table 3: The values of $\text{lcb}_t^{(m)}(\mathbf{x})$ and $\text{ucb}_t^{(m)}(\mathbf{x})$ for commonly used risk measures

Risk measure	Definition	$\text{lcb}_t^{(m)}(\mathbf{x})$	$\text{ucb}_t^{(m)}(\mathbf{x})$
Bayes risk	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}]$	$\mathbb{E}[l_{t,\mathbf{x},\mathbf{w}}^{(m)}]$	$\mathbb{E}[u_{t,\mathbf{x},\mathbf{w}}^{(m)}]$
Worst-case	$\inf_{\mathbf{w} \in \Omega} f_{\mathbf{x},\mathbf{w}}^{(m)}$	$\inf_{\mathbf{w} \in \Omega} l_{t,\mathbf{x},\mathbf{w}}^{(m)}$	$\inf_{\mathbf{w} \in \Omega} u_{t,\mathbf{x},\mathbf{w}}^{(m)}$
Best-case	$\sup_{\mathbf{w} \in \Omega} f_{\mathbf{x},\mathbf{w}}^{(m)}$	$\sup_{\mathbf{w} \in \Omega} l_{t,\mathbf{x},\mathbf{w}}^{(m)}$	$\sup_{\mathbf{w} \in \Omega} u_{t,\mathbf{x},\mathbf{w}}^{(m)}$
α -value-at-risk	$\inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(f_{\mathbf{x},\mathbf{w}}^{(m)} \leq b)\}$	$\inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(l_{t,\mathbf{x},\mathbf{w}}^{(m)} \leq b)\}$	$\inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(u_{t,\mathbf{x},\mathbf{w}}^{(m)} \leq b)\}$
α -conditional value-at-risk	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} \mid f_{\mathbf{x},\mathbf{w}}^{(m)} \leq v_{f^{(m)}}(\mathbf{x}; \alpha)]$	$\frac{1}{\alpha} \int_0^\alpha v_{l^{(m)}}(\mathbf{x}; \alpha') d\alpha'$	$\frac{1}{\alpha} \int_0^\alpha v_{u^{(m)}}(\mathbf{x}; \alpha') d\alpha'$
Mean absolute deviation	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}]]$	$\mathbb{E}[\min\{ j_{t,\mathbf{x},\mathbf{w}}^{(m)} , \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)} \} - \text{STR}(j_{t,\mathbf{x},\mathbf{w}}^{(m)}, \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)})]$	$\mathbb{E}[\max\{ j_{t,\mathbf{x},\mathbf{w}}^{(m)} , \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)} \}]$
Standard deviation	$\sqrt{\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}] ^2]}$	$\sqrt{\mathbb{E}[\min\{ j_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2, \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2\} - \text{STR}^2(j_{t,\mathbf{x},\mathbf{w}}^{(m)}, \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)})]}$	$\sqrt{\mathbb{E}[\max\{ j_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2, \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2\}]}$
Variance	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}] ^2]$	$\mathbb{E}[\min\{ j_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2, \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2\} - \text{STR}^2(j_{t,\mathbf{x},\mathbf{w}}^{(m)}, \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)})]$	$\mathbb{E}[\max\{ j_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2, \tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)} ^2\}]$
Distributionally robust	$\inf_{P \in \mathcal{A}} F^{(m)}(\mathbf{x}; P)$	$\inf_{P \in \mathcal{A}} \text{lcb}_t^{(m)}(\mathbf{x}; P)$	$\inf_{P \in \mathcal{A}} \text{ucb}_t^{(m)}(\mathbf{x}; P)$
Monotonic Lipschitz map	$\mathcal{M}(F^{(m)}(\mathbf{x}))$	$\min\{\mathcal{M}(\text{lcb}_t^{(m)}(\mathbf{x})), \mathcal{M}(\text{ucb}_t^{(m)}(\mathbf{x}))\}$	$\max\{\mathcal{M}(\text{lcb}_t^{(m)}(\mathbf{x})), \mathcal{M}(\text{ucb}_t^{(m)}(\mathbf{x}))\}$
Weighted sum	$\alpha_1 F^{(m_1)}(\mathbf{x}) + \alpha_2 F^{(m_2)}(\mathbf{x})$	$\alpha_1 \text{lcb}_t^{(m_1)}(\mathbf{x}) + \alpha_2 \text{lcb}_t^{(m_2)}(\mathbf{x})$	$\alpha_1 \text{ucb}_t^{(m_1)}(\mathbf{x}) + \alpha_2 \text{ucb}_t^{(m_2)}(\mathbf{x})$
Probabilistic threshold	$\mathbb{P}(f_{\mathbf{x},\mathbf{w}}^{(m)} \geq \theta)$	$\mathbb{P}(l_{t,\mathbf{x},\mathbf{w}}^{(m)} \geq \theta)$	$\mathbb{P}(u_{t,\mathbf{x},\mathbf{w}}^{(m)} \geq \theta)$

$f_{\mathbf{x},\mathbf{w}}^{(m)} \equiv f^{(m)}(\mathbf{x}, \mathbf{w})$, $j_{t,\mathbf{x},\mathbf{w}}^{(m)} \equiv l_t^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$, $u_{t,\mathbf{x},\mathbf{w}}^{(m)} \equiv u_t^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$, $v_{f^{(m)}}(\mathbf{x}; \alpha) \equiv \inf\{b \in \mathbb{R} \mid \mathbb{P}(f_{\mathbf{x},\mathbf{w}}^{(m)} \leq b) \geq \alpha\}$
 $v_{l^{(m)}}(\mathbf{x}; \alpha) \equiv \inf\{b \in \mathbb{R} \mid \mathbb{P}(l_{t,\mathbf{x},\mathbf{w}}^{(m)} \leq b) \geq \alpha\}$, $v_{u^{(m)}}(\mathbf{x}; \alpha) \equiv \inf\{b \in \mathbb{R} \mid \mathbb{P}(u_{t,\mathbf{x},\mathbf{w}}^{(m)} \leq b) \geq \alpha\}$, $\alpha \in (0, 1)$
 $j_{t,\mathbf{x},\mathbf{w}}^{(m)} \equiv l_{t,\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[l_{t,\mathbf{x},\mathbf{w}}^{(m)}]$, $\tilde{u}_{t,\mathbf{x},\mathbf{w}}^{(m)} \equiv u_{t,\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[u_{t,\mathbf{x},\mathbf{w}}^{(m)}]$, $\text{STR}(a, b) \equiv \max\{\min\{-a, b\}, 0\}$
 $F^{(m)}(\mathbf{x}; P)$: Risk measure $F^{(m)}(\mathbf{x})$ defined based on the distribution P
 $\text{lcb}_t^{(m)}(\mathbf{x}; P)$, $\text{ucb}_t^{(m)}(\mathbf{x}; P)$: $\text{lcb}_t^{(m)}(\mathbf{x})$ and $\text{ucb}_t^{(m)}(\mathbf{x})$ for $F^{(m)}(\mathbf{x}; P)$
 $q^{(m)}(a; F^{(m)})$: a function $q^{(m)}(a)$ for $F^{(m)}(\mathbf{x})$, does not depend on P
 $\mathcal{M}(\cdot)$: Monotonic Lipschitz continuous map with a Lipschitz constant K
 $\alpha_1, \alpha_2 \geq 0$
 α -value-at-risk is the same meaning as α -quantile

given in Table 3, where we omit the notation $\tilde{\cdot}$ and l in the table for simplicity. Note that by combining several risk measures such as the Bayes risk, standard deviation, monotonic Lipschitz map and weighted sum, we can obtain the result for mixed risk measures such as $0.7F^{(m_1)}(\mathbf{x}) - 0.3F^{(m_2)}(\mathbf{x})$, where $F^{(m_1)}(\mathbf{x})$ and $F^{(m_2)}(\mathbf{x})$ are the Bayes risk and standard deviation, respectively. In the sampling method, we generate S sample paths $f_1^{(m)}(\mathbf{x}, \mathbf{w}), \dots, f_S^{(m)}(\mathbf{x}, \mathbf{w})$ of $f^{(m)}(\mathbf{x}, \mathbf{w})$ independently from the GP posterior and compute

$$\text{lcb}_{t-1}^{(m,l)}(\mathbf{x}) = \min_{j \in [S], f_j^{(m)}(\mathbf{x}, \mathbf{w}) \in \tilde{G}_{t-1}^{(m)}(\mathbf{x})} \rho^{(m,l)}(f_j^{(m)}(\mathbf{x}, \mathbf{w})),$$

$$\text{ucb}_{t-1}^{(m,l)}(\mathbf{x}) = \max_{j \in [S], f_j^{(m)}(\mathbf{x}, \mathbf{w}) \in \tilde{G}_{t-1}^{(m)}(\mathbf{x})} \rho^{(m,l)}(f_j^{(m)}(\mathbf{x}, \mathbf{w})).$$

Pareto Front Estimation For any input $\mathbf{x} \in \mathcal{X}$ and subset $E \subset \mathcal{X}$, we define $\mathbf{LCB}_{t-1}(\mathbf{x})$, $\mathbf{UCB}_{t-1}(\mathbf{x})$ and $\mathbf{LCB}_{t-1}(E)$ as

$$\mathbf{LCB}_{t-1}(\mathbf{x}) = (\text{lcb}_{t-1}^{(1,1)}(\mathbf{x}), \dots, \text{lcb}_{t-1}^{(M_f, L_{M_f})}(\mathbf{x})), \mathbf{UCB}_{t-1}(\mathbf{x}) = (\text{ucb}_{t-1}^{(1,1)}(\mathbf{x}), \dots, \text{ucb}_{t-1}^{(M_f, L_{M_f})}(\mathbf{x})),$$

$$\mathbf{LCB}_{t-1}(E) = \{\mathbf{LCB}_{t-1}(\mathbf{x}) \mid \mathbf{x} \in E\}.$$

The estimated Pareto solution set $\hat{\Pi}_{t-1} \subset \mathcal{X}$ for the design variables is then defined as follows:

$$\hat{\Pi}_{t-1} = \{\mathbf{x} \in \mathcal{X} \mid \mathbf{LCB}_{t-1}(\mathbf{x}) \in \text{Par}(\mathbf{LCB}_{t-1}(\mathcal{X}))\}.$$

Here, in order to actually compute $\hat{\Pi}_{t-1}$, we need to compute the PF defined by $\mathbf{LCB}_{t-1}(\mathbf{x})$. However, if \mathcal{X} is an infinite set, then $\hat{\Pi}_{t-1}$ may also be an infinite set. In this case, since the exact calculation of $\hat{\Pi}_{t-1}$ is difficult, it is necessary to make a finite approximation using an approximation solver such as NSGA-II (Deb et al., 2002).

Acquisition Function We propose an AF for determining the next point to be evaluated. We define AF $a_t^{(\mathcal{X})}(\mathbf{x})$ for $\mathbf{x} \in \mathcal{X}$ as

$$a_t^{(\mathcal{X})}(\mathbf{x}) = \text{dist}(\mathbf{UCB}_t(\mathbf{x}), \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t)))$$

Then, the next design variable, \mathbf{x}_{t+1} , to be evaluated is selected by

$$\mathbf{x}_{t+1} = \underset{\mathbf{x} \in \mathcal{X}}{\text{argmax}} a_t^{(\mathcal{X})}(\mathbf{x}).$$

Hence, the value of $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1})$ is equal to the following maximin distance:

$$a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) = \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{b} \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{UCB}_t(\mathbf{x}), \mathbf{b}).$$

The value of $a_t^{(\mathcal{X})}(\mathbf{x})$ can be computed analytically using the following lemma when $\hat{\Pi}_t$ is finite:

Lemma A.1. Let $\mathbf{UCB}_t(\mathbf{x}) = (u_1, \dots, u_L)$ and $\mathbf{LCB}_t(\hat{\Pi}_t) = \{(l_1^{(i)}, \dots, l_L^{(i)}) \mid 1 \leq i \leq k\}$. Then, $a_t^{(\mathcal{X})}(\mathbf{x})$ can be computed by

$$a_t^{(\mathcal{X})}(\mathbf{x}) = \max\{\tilde{a}_t(\mathbf{x}), 0\}, \quad \tilde{a}_t(\mathbf{x}) = \min_{1 \leq i \leq k} \max\{u_1 - l_1^{(i)}, \dots, u_L - l_L^{(i)}\}.$$

Next, we consider the simulator setting. In this case, we have to select the environment variable \mathbf{w}_{t+1} . Based on the fact that many risk measures including Bayes risk satisfy

$$\|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty \leq q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\beta_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right), \quad (5)$$

where $q(\cdot) : [0, \infty) \rightarrow [0, \infty)$ is a strictly increasing function defined by risk measures and satisfies $q(0) = 0$, we choose \mathbf{w}_{t+1} as follows:

$$\mathbf{w}_{t+1} = \underset{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}}{\text{argmax}} a_t^{(\Omega_{\mathbf{x}_{t+1}})}(\mathbf{w}), \quad a_t^{(\Omega_{\mathbf{x}_{t+1}})}(\mathbf{w}) = \sum_{m=1}^{M_f} 2\beta_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}).$$

On the other hand, in the uncontrollable setting, since we cannot control \mathbf{w} , \mathbf{w}_{t+1} is defined as the sample from $\Omega_{\mathbf{x}}$.

A.3 Stopping Condition

We describe the stopping conditions of the proposed algorithm. Let $\epsilon > 0$ be a predetermined desired accuracy parameter. Then the algorithm is terminated if $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq \epsilon$ is satisfied. The pseudocode of the proposed algorithm is described in Algorithm 2.

A.4 Theoretical Analysis

In this subsection, we give the theorems for the accuracy and termination of the proposed algorithm. First, we quantify the goodness of the estimated $\hat{\Pi}_t$. If $\hat{\Pi}_t$ is a good estimate, the following two indicators defined by $\hat{\Pi}_t$ should be small:

$$I_t^{(i)} = \max_{\mathbf{y} \in Z^*} \text{dist}(\mathbf{y}, \text{Par}(\mathbf{F}(\hat{\Pi}_t))),$$

$$I_t^{(ii)} = \max_{\mathbf{y} \in \mathbf{F}(\hat{\Pi}_t)} \text{dist}(\mathbf{y}, Z^*).$$

Using these, we define the inference discrepancy $I_t = \max\{I_t^{(i)}, I_t^{(ii)}\}$ for $\hat{\Pi}_t$ as the goodness measure. Next, in order to show the theoretical validity of the proposed algorithm, we introduce the maximum information gain $\tilde{\kappa}_t^{(m)}$. The maximum information gain $\tilde{\kappa}_T^{(m)}$ under the heteroscedastic sub-Gaussian setting can be expressed as follows (Makarova et al., 2021):

$$\tilde{\kappa}_T^{(m)} = \max_{(\mathbf{x}_1, \mathbf{w}_1), \dots, (\mathbf{x}_T, \mathbf{w}_T)} \frac{1}{2} \sum_{t=1}^T \log \left(1 + \frac{\tilde{\sigma}_{t-1}^{(m)2}(\mathbf{x}_t, \mathbf{w}_t)}{\tau_{\mathbf{x}_t, \mathbf{w}_t, m}^2} \right).$$

The order of $\tilde{\kappa}_T^{(m)}$ with respect to widely used kernels such as linear and squared exponential kernels is derived by Makarova et al. (2021). Then, the following theorem holds:

Algorithm 2 Bounding box-based MOBO of general risk measures under extended problem setup

Require: GP priors $\mathcal{GP}(0, k^{(m)})$, tradeoff parameters $\{\beta_{m,t}\}_{t \geq 0}$, accuracy parameter $\epsilon > 0$, $M_f \geq 1$, $L_m \geq 1$, $L \geq 2$

for $t = 0, 1, 2, \dots$ **do**

Compute $\tilde{Q}_t^{(f^{(m)})}(\mathbf{x}, \mathbf{w})$ for each $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$

Compute $\tilde{Q}_t^{(F^{(m,l)})}(\mathbf{x})$ for each $\mathbf{x} \in \mathcal{X}$ by the generalized, decomposition or sampling method

Compute $\tilde{B}_t(\mathbf{x}) = \prod_{m=1}^{M_f} \prod_{l=1}^{L_m} Q_t^{(F^{(m,l)})}(\mathbf{x})$ for each $\mathbf{x} \in \mathcal{X}$

Estimate $\hat{\Pi}_t$ by $\tilde{B}_t(\mathbf{x})$

Select the next evaluation point \mathbf{x}_{t+1} by $a_t^{(\mathcal{X})}(\mathbf{x})$

if $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq \epsilon$ **then**

break

end if

if simulator setting **then**

Select the next evaluation point \mathbf{w}_{t+1} by $a_t^{(\Omega_{\mathbf{x}_{t+1}})}(\mathbf{w})$

else {uncontrollable setting}

\mathbf{w}_{t+1} is generated from $P_{\mathbf{w}}(\mathbf{x}_{t+1})$

end if

Observe $y_{t+1}^{(m)} = f^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}) + \eta^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1})$ at the point $(\mathbf{x}_{t+1}, \mathbf{w}_{t+1})$

Update GPs by adding observed points

end for

Ensure: Return $\hat{\Pi}_t$ as the estimated set of design variables

Lemma A.2 (Lemma 7 in Kirschner and Krause (2018)). Suppose that the regularity assumption holds. Let $\delta \in (0, 1)$, $\lambda_1, \dots, \lambda_{M_f} > 0$ and define

$$\tilde{\beta}_{m,t}^{1/2} = B_m \sqrt{\lambda_m} + \sqrt{2 \log \left(\frac{\det(\lambda_m \Sigma_t^{(m)} + \tilde{\mathbf{K}}_t^{(m)})^{1/2}}{M_f^{-1} \delta \det(\lambda_m \Sigma_t^{(m)})^{1/2}} \right)}.$$

Then, with probability at least $1 - \delta$, the following inequality holds for any $t \geq 1$, $m \in [M_f]$ and $(\mathbf{x}, \mathbf{w}) \in \mathcal{X} \times \Omega$:

$$|f^{(m)}(\mathbf{x}, \mathbf{w}) - \tilde{\mu}_{t-1}^{(m)}(\mathbf{x}, \mathbf{w})| \leq \tilde{\beta}_{m,t}^{1/2} \tilde{\sigma}_{t-1}^{(m)}(\mathbf{x}, \mathbf{w}).$$

Note that from the definition of the maximum information gain, when $\lambda_m = 1$, $M_f = M \geq 2$, $L_m = 1$ and $\tau_{\mathbf{x}, \mathbf{w}, m}^2 = \varsigma_m^2$ and the true noise distribution is Gaussian, the inequality $\tilde{\beta}_{m,t}^{1/2} \leq \beta_{m,t}^{1/2}$ holds, where $\beta_{m,t}^{1/2}$ is given by Lemma 4.1. Using this, we give the theorems for the accuracy, termination, $q(a)$ and approximation errors under both the simulator and uncontrollable settings.

Theorem A.1 (Simulator and uncontrollable settings). Suppose that the assumption of Lemma A.2 and the inequality (4) hold. Let $t \geq 0$, $m \in [M_f]$, $\delta \in (0, 1)$, and let $\tilde{\beta}_{m,t+1}^{1/2}$ be defined as in Lemma A.2. In addition, let $\epsilon > 0$ be a predetermined desired accuracy parameter. Then, with probability at least $1 - \delta$, the inequality $I_t \leq a_t^{(\mathcal{X})}(\mathbf{x}_{t+1})$ holds for any $t \geq 0$ and \mathbf{x}_{t+1} . Therefore, if the stopping condition satisfies at T iterations, the inference discrepancy I_T satisfies $I_T \leq \epsilon$ with probability at least $1 - \delta$.

Theorem A.2 (Simulator setting). Suppose that the assumption in Theorem A.1 holds. Let $q : [0, \infty) \rightarrow [0, \infty)$ be a strictly increasing function satisfying $q(0) = 0$ and (5). Also let

$$s_t = \sqrt{\frac{\sum_{m=1}^{M_f} \tilde{C}_m \tilde{\beta}_{m,t+1} \tilde{\kappa}_{t+1}^{(m)}}{t+1}},$$

where $\tilde{C}_m = \frac{8M_f}{\log(1+\lambda_m^{-1}\tau^{-2})}$. Then, the inequality $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1}) \leq q(s_t)$ holds for any $t \geq 0$ and some $\hat{t} \leq t$. Therefore, Algorithm 2 terminates after at most T iterations, where T is the smallest positive integer satisfying $q(s_T) \leq \epsilon$.

Theorem A.3 (Simulator and uncontrollable settings). Suppose that the assumption in Theorem A.1 holds. Also assume that there exist strictly increasing functions $q^{(m,l)} : [0, \infty) \rightarrow [0, \infty)$ satisfying $q^{(m,l)}(0) = 0$ and

$$|\text{ucb}_t^{(m,l)}(\mathbf{x}_{t+1}) - \text{lcb}_t^{(m,l)}(\mathbf{x}_{t+1})| \leq q^{(m,l)}(\tilde{s}_t)$$

for any $t \geq 0$, $m \in [M_f]$, $l \in [L_m]$ and $\mathbf{x}_{t+1} \in \mathcal{X}$, where

$$\tilde{s}_t = \max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}).$$

Then, $q(a) \equiv \max_{m \in [M_f], l \in [L_m]} q^{(m,l)}(a)$ is the strictly increasing function and satisfies $q(0) = 0$ and (5).

Specific forms of $q^{(m,l)}(a)$ for commonly used risk measures are described in Table 4. For simplicity, we omitted l in the table. From Table 4, the probabilistic threshold measure does not satisfy the inequality in Theorem A.3. For example, if $f^{(m)}(\mathbf{x}, \mathbf{w}) = \theta$, then with high probability $\text{ucb}_t^{(m)}(\mathbf{x}_{t+1})$ and $\text{lcb}_t^{(m)}(\mathbf{x}_{t+1})$ are respectively close to one and zero even when \tilde{s}_t is close to zero. Iwazaki et al. (2021a); Inatsu et al. (2021) proposed BO methods for the (distributionally robust) probabilistic threshold measure and confronted the same problem. They solved this problem by assuming the condition that the probability of a black-box function accumulating in the neighborhood of the threshold is small, and derived $\text{ucb}_t^{(m)}(\mathbf{x}_{t+1}) - \text{lcb}_t^{(m)}(\mathbf{x}_{t+1}) \leq \tilde{q}(\tilde{s}_t) + \xi$, where $\tilde{q}(a) = 0$ if $a \leq c$ and otherwise $\tilde{q}(a) = 1$, and c is some positive constant.

Next, we consider the approximation error setting. Let $\epsilon_{\text{lcb}}, \epsilon_{\text{ucb}}, \epsilon_{\text{PF}}, \epsilon_{\mathcal{X}}, \epsilon_{\Omega}$ be non-negative error parameters that represent the errors in these approximations, respectively. We consider the case that the following four error inequalities hold for any $t \geq 0$, $m \in [M_f]$, $l \in [L_m]$, $\mathbf{x}, \mathbf{x}_{t+1} \in \mathcal{X}$, $\mathbf{w}_{t+1} \in \Omega_{\mathbf{x}_{t+1}}$ and $g(\mathbf{x}, \mathbf{w}) \in \tilde{G}_t^{(m)}(\mathbf{x})$:

$$\text{lcb}_t^{(m,l)}(\mathbf{x}) - \epsilon_{\text{lcb}} \leq \rho^{(m,l)}(g(\mathbf{x}, \mathbf{w})) \leq \text{ucb}_t^{(m,l)}(\mathbf{x}) + \epsilon_{\text{ucb}}, \quad (6)$$

$$\max_{\mathbf{y} \in \text{Par}(\text{LCB}_t(\hat{\Pi}_t))} \text{dist}(\mathbf{y}, \text{Par}(\text{LCB}_t(\mathcal{X}))) \leq \epsilon_{\text{PF}}, \quad (7)$$

$$\max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}) - a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq \epsilon_{\mathcal{X}}, \quad (8)$$

$$\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} a_t^{(\Omega_{\mathbf{x}_{t+1}})}(\mathbf{w}) - a_t^{(\Omega_{\mathbf{x}_{t+1}})}(\mathbf{w}_{t+1}) \leq \epsilon_{\Omega}. \quad (9)$$

Theorem A.4 (Simulator setting). Suppose that the assumption in Lemma A.2 holds. Let $t \geq 0$, $m \in [M_f]$, $l \in [L_m]$, $\delta \in (0, 1)$, and let $\tilde{\beta}_{m,t+1}^{1/2}$ be defined as in Lemma A.2. In addition, let $\epsilon > 0$ be a predetermined desired accuracy parameter. Moreover, let $\epsilon_{\text{lcb}}, \epsilon_{\text{ucb}}, \epsilon_{\text{PF}}, \epsilon_{\mathcal{X}}, \epsilon_{\Omega}$ be non-negative error parameters satisfying (6)–(9). Then, with probability at least $1 - \delta$, the inequality $I_t \leq a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) + \epsilon_{\text{lcb}} + \epsilon_{\text{ucb}} + \epsilon_{\mathcal{X}}$ holds for any $t \geq 0$ and \mathbf{x}_{t+1} . Therefore, if the stopping condition satisfies at T iterations, the inference discrepancy I_T satisfies $I_T \leq \epsilon + \epsilon_{\text{lcb}} + \epsilon_{\text{ucb}} + \epsilon_{\mathcal{X}}$ with probability at least $1 - \delta$.

Theorem A.5 (Simulator setting). Suppose that the assumption in Theorem A.4 holds. Let $q : [0, \infty) \rightarrow [0, \infty)$ be a strictly increasing function satisfying $q(0) = 0$ and (5). Then, the inequality $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1}) \leq \epsilon_{\text{PF}} + q(\epsilon_{\Omega} + s_t)$ holds for any $t \geq 0$ and some $\hat{t} \leq t$, where s_t is given by Theorem A.2. Therefore, Algorithm 2 terminates after at most T iterations, where T is the smallest positive integer satisfying $\epsilon_{\text{PF}} + q(\epsilon_{\Omega} + s_T) \leq \epsilon$.

Note that for Theorem A.5, the integer T satisfying the theorem's last inequality does not always exist. However, the left hand side in this inequality is merely an upper bound of $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1})$. Thus, in some cases the actual value of $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1})$ satisfies $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1}) \leq \epsilon$ and the stopping condition is satisfied.

Uncontrollable Setting We provide theoretical results for the uncontrollable setting. First, we define the following two additional conditions:

Condition A.1. Let $\text{Nei}(\mathbf{a}; r)$ be an open ball with center \mathbf{a} and radius $r > 0$, where the distance is taken with respect to L_1 -distance. Then, for any $\mathbf{x} \in \mathcal{X}$, $\hat{\mathbf{w}} \in \Omega_{\mathbf{x}}$ and $\zeta > 0$, $P_{\mathbf{w}}(\mathbf{x})$ satisfies

$$\mathbb{P}_{P_{\mathbf{w}}(\mathbf{x})}[\mathbf{w} \in \text{Nei}(\hat{\mathbf{w}}; \zeta)] > 0,$$

where $\mathbb{P}_{P_{\mathbf{w}}(\mathbf{x})}[\cdot]$ is the probability measure with respect to $P_{\mathbf{w}}(\mathbf{x})$.

Condition A.2. Let L_σ be a positive number. Then, $\tilde{\sigma}_t^{(m)}(\mathbf{x}, \mathbf{w})$ is an L_σ -data-independent-Lipschitz continuous, that is, the following inequality holds for any $t \geq 1$, $m \in [M_f]$ and $\{(\mathbf{x}_i, \mathbf{w}_i)\}_{i=1}^t$:

$$\forall (\mathbf{x}, \mathbf{w}), (\tilde{\mathbf{x}}, \tilde{\mathbf{w}}) \in \mathcal{X} \times \Omega, |\tilde{\sigma}_t^{(m)}(\mathbf{x}, \mathbf{w}) - \tilde{\sigma}_t^{(m)}(\tilde{\mathbf{x}}, \tilde{\mathbf{w}})| \leq L_\sigma \|(\mathbf{x}^\top, \mathbf{w}^\top)^\top - (\tilde{\mathbf{x}}^\top, \tilde{\mathbf{w}}^\top)^\top\|_1$$

Condition A.1 implies that the support of $P_{\mathbf{w}}(\mathbf{x})$ is equal to $\Omega_{\mathbf{w}}$. The assumption that the support of the distribution of \mathbf{w} and the set of \mathbf{w} are the same is also used in existing studies that conduct theoretical analysis of BOs for risk measures under IU (Nguyen et al., 2021b; Inatsu et al., 2022). Similarly, Condition A.2 is introduced by Kusakawa et al. (2022), and they proved that Condition A.2 holds if the linear, Gaussian or Matérn (with parameter $\nu > 1$) is used. Their proof is given under constant variance of the normal error distribution for GP models, but similar arguments can be derived in the setting considered in this section. We also define a maximal ζ -separated subset of $\Omega_{\mathbf{x}}$:

Definition A.2. Let ζ be a positive number. Then, a subset $S \subset \Omega_{\mathbf{x}}$ is called the maximal ζ -separated subset of $\Omega_{\mathbf{x}}$, if the following holds:

1. For any $\mathbf{w}, \mathbf{w}' \in S$, $\mathbf{w} \neq \mathbf{w}' \Rightarrow \|\mathbf{w} - \mathbf{w}'\|_1 > \zeta$.
2. For any $\mathbf{w} \in \Omega_{\mathbf{x}}$, there exists $\mathbf{w}' \in S$ such that $\|\mathbf{w} - \mathbf{w}'\|_1 \leq \zeta$.

Note that a compact set A has a maximal ζ -separated subset. Let $\mathcal{S}(\Omega_{\mathbf{x}}; \zeta)$ be a maximal ζ -separated subset of $\Omega_{\mathbf{x}}$. From Condition A.1 and compactness of $\Omega_{\mathbf{x}}$, for any $\zeta > 0$ and $\mathbf{x} \in \mathcal{X}$, the following holds:

$$\min_{\hat{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}}; \zeta)} \mathbb{P}_{P_{\mathbf{w}}(\mathbf{x})}[\mathbf{w} \in \text{Nei}(\hat{\mathbf{w}}; \zeta/2)] \equiv \underline{p}_{\mathbf{x}, \zeta} > 0. \quad (10)$$

In contrast, (10) does not necessarily guarantee $\inf_{\mathbf{x} \in \mathcal{X}} \underline{p}_{\mathbf{x}, \zeta} > 0$. However, $\inf_{\mathbf{x} \in \mathcal{X}} \underline{p}_{\mathbf{x}, \zeta} = 0$ implies that given $\zeta > 0$ and for any $\nu > 0$, there exist an open ball $\text{Nei}(\hat{\mathbf{w}}; \zeta/2)$ and $\hat{\mathbf{x}} \in \mathcal{X}$ such that $\mathcal{P}_{P_{\mathbf{w}}(\hat{\mathbf{x}})}[\mathbf{w} \in \text{Nei}(\hat{\mathbf{w}}; \zeta/2)] < \nu$. This means that the probability of \mathbf{w} realizes to the open ball with radius ζ can be as small as desired. Thus, to avoid this extreme case, we assume

$$\inf_{\mathbf{x} \in \mathcal{X}} \underline{p}_{\mathbf{x}, \zeta} \equiv \underline{p}_\zeta > 0. \quad (11)$$

Then, the following theorems hold:

Theorem A.6 (Uncontrollable setting). Suppose that the assumption in Theorem A.1 holds. Let $q : [0, \infty) \rightarrow [0, \infty)$ be a strictly increasing function satisfying $q(0) = 0$ and (5). Assume that Condition A.1 and A.2 hold. Let ζ_1, \dots, ζ_t be positive numbers and $\underline{p}_{\zeta_1}, \dots, \underline{p}_{\zeta_t}$ be numbers defined by (11). Let $\tilde{p}_{\zeta_i} = \min_{1 \leq j \leq i} \underline{p}_{\zeta_j}$, $\tilde{\beta}_t^{1/2} = \max_{1 \leq m \leq M_f} \tilde{\beta}_{m,t}^{1/2}$, $\tilde{\kappa}_t = \max_{1 \leq m \leq M_f} \tilde{\kappa}_t^{(m)}$ and define

$$\hat{s}_t = \frac{2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} (1 + \tilde{p}_{\zeta_{t+1}}^{-1})}{t+1} \sum_{i=1}^{t+1} \zeta_i + \frac{16J \log(8J/\delta) \tilde{\beta}_{t+1}^{1/2} \tilde{p}_{\zeta_{t+1}}^{-1}}{t+1} + \sqrt{\frac{\hat{C} \tilde{p}_{\zeta_{t+1}}^{-2} \tilde{\beta}_{t+1} \tilde{\kappa}_{t+1}}{t+1}},$$

where $J = M_f \max\{1, \lambda_1^{-1}, \dots, \lambda_{M_f}^{-1}\}$, $\hat{C} = M_f \max_{1 \leq m \leq M_f} \hat{C}_m$ and $\hat{C}_m = \frac{32M_f}{\log(1 + \lambda_m^{-1} \underline{p}_\zeta^{-2})}$. Then, with probability at least $1 - \delta$, the inequality $a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq q(\hat{s}_t)$ holds for any $t \geq 0$ and some $\hat{t} \leq t$. Therefore, with probability at least $1 - \delta$, Algorithm 2 terminates after at most T iterations, where T is the smallest positive integer satisfying $q(\hat{s}_T) \leq \epsilon$.

In Theorem A.6, the choice of ζ_1, \dots, ζ_t is important and must be chosen that \hat{s}_t converges to 0. The simplest example is the case where $\Omega_{\mathbf{x}}$ is a finite set and equal to Ω for all $\mathbf{x} \in \mathcal{X}$. In this case, noting that $\lim_{\zeta \rightarrow 0} \underline{p}_\zeta > 0$ and $\sum_{t=1}^{\infty} \zeta_t = t^{-2} = \pi^2/6$, \hat{s}_t converges to 0 when $\tilde{\beta}_t^{1/2}$ and $\tilde{\beta}_t \tilde{\kappa}_t$ are sublinear. Inatsu et al. (2022) used the finiteness assumption for set of the environmental variables in theoretical analysis for uncontrollable settings under IU. On the other hand, Iwazaki et al. (2021b) considered the Bayes risk and standard deviation risk under the uncontrollable setting, and they derived the similar theoretical result without the finiteness assumption. Their approach can be used for moment-based risk measures such as Bayes risk, but not for quantile-based methods such as the worst-case risk. As another example, when $\Omega_{\mathbf{x}} = \Omega = [0, 1]$ and $P_{\mathbf{w}}$ follows the uniform distribution on Ω , the orders of $\tilde{p}_{\zeta_t}^{-1}$ and $\sum_{i=1}^t \zeta_i$ are respectively $\log t$ and $t/\log t$ if $\zeta_i = 1/(\log i)$. Then, the

Table 4: Specific forms of $q^{(m)}(a)$ for commonly used risk measures

Risk measure	Definition	$q^{(m)}(a)$
Bayes risk	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}]$	a
Worst-case	$\inf_{\mathbf{w} \in \Omega} f_{\mathbf{x},\mathbf{w}}^{(m)}$	a
Best-case	$\sup_{\mathbf{w} \in \Omega} f_{\mathbf{x},\mathbf{w}}^{(m)}$	a
α -value-at-risk	$\inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(f_{\mathbf{x},\mathbf{w}}^{(m)} \leq b)\}$	a
α -conditional value-at-risk	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} \mid f_{\mathbf{x},\mathbf{w}}^{(m)} \leq v_{f^{(m)}}(\mathbf{x}; \alpha)]$	a
Mean absolute deviation	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}]]$	$2a$
Standard deviation	$\sqrt{\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}] ^2]}$	$\sqrt{8B_m a + 5a^2}$
Variance	$\mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)} - \mathbb{E}[f_{\mathbf{x},\mathbf{w}}^{(m)}] ^2]$	$8B_m a + 5a^2$
Distributionally robust	$\inf_{P \in \mathcal{A}} F^{(m)}(\mathbf{x}; P)$	$q^{(m)}(a; F^{(m)})$
Monotonic Lipschitz map	$\mathcal{M}(F^{(m)}(\mathbf{x}))$	$Kq^{(m)}(a)$
Weighted sum	$\alpha_1 F^{(m_1)}(\mathbf{x}) + \alpha_2 F^{(m_2)}(\mathbf{x})$	$\alpha_1 q^{(m_1)}(a) + \alpha_2 q^{(m_2)}(a)$
Probabilistic threshold	$\mathbb{P}(f_{\mathbf{x},\mathbf{w}}^{(m)} \geq \theta)$	-

$f_{\mathbf{x},\mathbf{w}}^{(m)} \equiv f^{(m)}(\mathbf{x}, \mathbf{w})$, $v_{f^{(m)}}(\mathbf{x}; \alpha) \equiv \inf\{b \in \mathbb{R} \mid \mathbb{P}(f_{\mathbf{x},\mathbf{w}}^{(m)} \leq b) \geq \alpha\}$, $\alpha \in (0, 1)$
 $F^{(m)}(\mathbf{x}; P)$: Risk measure $F^{(m)}(\mathbf{x})$ defined based on the distribution P
 $q^{(m)}(a; F^{(m)})$: a function $q^{(m)}(a)$ for $F^{(m)}(\mathbf{x})$, does not depend on P
 $\mathcal{M}(\cdot)$: Monotonic Lipschitz continuous map with a Lipschitz constant K
 $\alpha_1, \alpha_2 \geq 0$

α -value-at-risk is the same meaning as α -quantile

dominant term of \hat{s}_t is the first term and its order is $\tilde{\beta}_t^{1/2}$. Recently, Takeno et al. (2023) has proposed a method in which $\tilde{\beta}_t$ does not diverge to infinity by stochastically sampling $\tilde{\beta}_t$ under the assumption that the true black-box function follows GP. Since their method is not an RKHS setting, nor is it a multi-objective optimization setting, it is not clear whether it is applicable to our setting, but it is one direction to consider.

Theorem A.7 (Uncontrollable setting). Suppose that the assumption in Lemma A.2 holds. Let $t \geq 0$, $m \in [M_f]$, $l \in [L_m]$, $\delta \in (0, 1)$, and let $\tilde{\beta}_{m,t+1}^{1/2}$ be defined as in Lemma A.2. In addition, let $\epsilon > 0$ be a predetermined desired accuracy parameter. Moreover, let $\epsilon_{\text{lcb}}, \epsilon_{\text{ucb}}, \epsilon_{\text{PF}}, \epsilon_{\mathcal{X}}$ be non-negative error parameters satisfying (6)–(8). Then, with probability at least $1 - \delta$, the inequality $I_t \leq a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) + \epsilon_{\text{lcb}} + \epsilon_{\text{ucb}} + \epsilon_{\mathcal{X}}$ holds for any $t \geq 0$ and \mathbf{x}_{t+1} . Therefore, if the stopping condition satisfies at T iterations, the inference discrepancy I_T satisfies $I_T \leq \epsilon + \epsilon_{\text{lcb}} + \epsilon_{\text{ucb}} + \epsilon_{\mathcal{X}}$ with probability at least $1 - \delta$.

Theorem A.8 (Uncontrollable setting). Suppose that the assumptions in Theorem A.6 and Theorem A.7 holds. Let $q : [0, \infty) \rightarrow [0, \infty)$ be a strictly increasing function satisfying $q(0) = 0$ and (5). Then, with probability at least $1 - \delta$, the inequality $a_{\hat{t}}^{(\mathcal{X})}(\mathbf{x}_{\hat{t}+1}) \leq \epsilon_{\text{PF}} + q(\hat{s}_t)$ holds for any $t \geq 0$ and some $\hat{t} \leq t$, where \hat{s}_t is given by Theorem A.6. Therefore, with probability at least $1 - \delta$, Algorithm 2 terminates after at most T iterations, where T is the smallest positive integer satisfying $\epsilon_{\text{PF}} + q(\hat{s}_T) \leq \epsilon$.

B PROOFS

In this section, we prove all theorems, lemmas and the results in Table 3 and 4.

B.1 Proof of Table 3 and 4

In this proof, we omit the notation $\tilde{\cdot}$ and (m) for simplicity. Let $\mathbf{x} \in \mathcal{X}$, $\mathbf{w} \in \Omega_{\mathbf{x}}$, $t \geq 0$ and $\beta_{t+1}^{1/2} \geq 0$. Assume that $l_{t,\mathbf{x},\mathbf{w}} \leq f(\mathbf{x}, \mathbf{w}) \leq u_{t,\mathbf{x},\mathbf{w}}$, where $l_{t,\mathbf{x},\mathbf{w}} = \mu_t(\mathbf{x}, \mathbf{w}) - \beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w})$ and $u_{t,\mathbf{x},\mathbf{w}} = \mu_t(\mathbf{x}, \mathbf{w}) + \beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w})$.

Bayes Risk Since \mathbf{w} is a random variable, $l_{t,\mathbf{x},\mathbf{w}}$, $u_{t,\mathbf{x},\mathbf{w}}$ and $f(\mathbf{x}, \mathbf{w})$ are also random variables. Hence, from the monotonicity of expectation and $l_{t,\mathbf{x},\mathbf{w}} \leq f(\mathbf{x}, \mathbf{w}) \leq u_{t,\mathbf{x},\mathbf{w}}$, we have

$$\text{lcb}_t(\mathbf{x}) \equiv \mathbb{E}[l_{t,\mathbf{x},\mathbf{w}}] \leq \mathbb{E}[f(\mathbf{x}, \mathbf{w})] \leq \mathbb{E}[u_{t,\mathbf{x},\mathbf{w}}] \equiv \text{ucb}_t(\mathbf{x}).$$

In addition, from the definition of $l_{t,\mathbf{x},\mathbf{w}}$ and $u_{t,\mathbf{x},\mathbf{w}}$, we get

$$0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) = \mathbb{E}[u_{t,\mathbf{x},\mathbf{w}} - l_{t,\mathbf{x},\mathbf{w}}] = \mathbb{E}[2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w})] \leq \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}).$$

Worst-case From the definition of infimum, noting that $l_{t,\mathbf{x},\mathbf{w}} \leq f(\mathbf{x}, \mathbf{w}) \leq u_{t,\mathbf{x},\mathbf{w}}$, we obtain

$$\text{lcb}_t(\mathbf{x}) \equiv \inf_{\mathbf{w} \in \Omega_{\mathbf{x}}} l_{t,\mathbf{x},\mathbf{w}} \leq \inf_{\mathbf{w} \in \Omega_{\mathbf{x}}} f(\mathbf{x}, \mathbf{w}) \leq \inf_{\mathbf{w} \in \Omega_{\mathbf{x}}} u_{t,\mathbf{x},\mathbf{w}} \equiv \text{ucb}_t(\mathbf{x}).$$

Moreover, from the property of infimum, for any $\epsilon > 0$, there exists $\hat{\mathbf{w}} \in \Omega_{\mathbf{x}}$ such that $l_{t,\mathbf{x},\hat{\mathbf{w}}} \leq \text{lcb}_t(\mathbf{x}) + \epsilon$. Therefore, noting that $\text{ucb}_t(\mathbf{x}) \leq u_{t,\mathbf{x},\hat{\mathbf{w}}}$, we get

$$\text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq u_{t,\mathbf{x},\hat{\mathbf{w}}} - l_{t,\mathbf{x},\hat{\mathbf{w}}} + \epsilon = 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \hat{\mathbf{w}}) + \epsilon \leq \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) + \epsilon.$$

Since ϵ is an arbitrary positive number, we have

$$0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}).$$

Best-case From the definition of supremum, noting that $l_{t,\mathbf{x},\mathbf{w}} \leq f(\mathbf{x}, \mathbf{w}) \leq u_{t,\mathbf{x},\mathbf{w}}$, we obtain

$$\text{lcb}_t(\mathbf{x}) \equiv \sup_{\mathbf{w} \in \Omega_{\mathbf{x}}} l_{t,\mathbf{x},\mathbf{w}} \leq \sup_{\mathbf{w} \in \Omega_{\mathbf{x}}} f(\mathbf{x}, \mathbf{w}) \leq \sup_{\mathbf{w} \in \Omega_{\mathbf{x}}} u_{t,\mathbf{x},\mathbf{w}} \equiv \text{ucb}_t(\mathbf{x}).$$

Moreover, from the property of supremum, for any $\epsilon > 0$, there exists $\hat{\mathbf{w}} \in \Omega_{\mathbf{x}}$ such that $\text{ucb}_t(\mathbf{x}) - \epsilon \leq u_{t,\mathbf{x},\hat{\mathbf{w}}}$. Therefore, noting that $\text{lcb}_t(\mathbf{x}) \geq l_{t,\mathbf{x},\hat{\mathbf{w}}}$, we get

$$\text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq u_{t,\mathbf{x},\hat{\mathbf{w}}} - l_{t,\mathbf{x},\hat{\mathbf{w}}} + \epsilon = 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \hat{\mathbf{w}}) + \epsilon \leq \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) + \epsilon.$$

Since ϵ is an arbitrary positive number, we have

$$0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}).$$

α -value-at-risk Let $\alpha \in (0, 1)$. For any $b \in \mathbb{R}$, $f(\mathbf{x}, \mathbf{w}) \leq u_{t,\mathbf{x},\mathbf{w}}$ implies that $\mathbb{P}(u_{t,\mathbf{x},\mathbf{w}} \leq b) \leq \mathbb{P}(f(\mathbf{x}, \mathbf{w}) \leq b)$. Thus, letting $\text{ucb}_t(\mathbf{x}) \equiv \inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(u_{t,\mathbf{x},\mathbf{w}} \leq b)\}$, we obtain

$$\alpha \leq \mathbb{P}(u_{t,\mathbf{x},\mathbf{w}} \leq \text{ucb}_t(\mathbf{x})) \leq \mathbb{P}(f(\mathbf{x}, \mathbf{w}) \leq \text{ucb}_t(\mathbf{x})).$$

This implies that

$$\inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(f(\mathbf{x}, \mathbf{w}) \leq b)\} \leq \text{ucb}_t(\mathbf{x}).$$

By using the same argument, we get

$$\text{lcb}_t(\mathbf{x}) \equiv \inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(l_{t,\mathbf{x},\mathbf{w}} \leq b)\} \leq \inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(f(\mathbf{x}, \mathbf{w}) \leq b)\}.$$

Furthermore, noting that the definition of $l_{t,\mathbf{x},\mathbf{w}}$ and $u_{t,\mathbf{x},\mathbf{w}}$, we get

$$u_{t,\mathbf{x},\mathbf{w}} \leq l_{t,\mathbf{x},\mathbf{w}} + \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}).$$

Therefore, we have

$$\begin{aligned} 0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) &= \inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(u_{t,\mathbf{x},\mathbf{w}} \leq b)\} - \inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(l_{t,\mathbf{x},\mathbf{w}} \leq b)\} \\ &\leq \inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(l_{t,\mathbf{x},\mathbf{w}} + \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) \leq b)\} - \inf\{b \in \mathbb{R} \mid \alpha \leq \mathbb{P}(l_{t,\mathbf{x},\mathbf{w}} \leq b)\} \\ &= \text{lcb}_t(\mathbf{x}) + \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) - \text{lcb}_t(\mathbf{x}) = \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}). \end{aligned}$$

α -conditional value-at-risk Let $\alpha \in (0, 1)$. From Nguyen et al. (2021a), α -conditional value-at-risk for $f(\mathbf{x}, \mathbf{w})$ can be written as follows:

$$\frac{1}{\alpha} \int_0^\alpha v_f(\mathbf{x}; \alpha') d\alpha'.$$

Thus, we have

$$\text{lcb}_t(\mathbf{x}) \equiv \frac{1}{\alpha} \int_0^\alpha v_{l_t}(\mathbf{x}; \alpha') d\alpha' \leq \frac{1}{\alpha} \int_0^\alpha v_f(\mathbf{x}; \alpha') d\alpha' \leq \frac{1}{\alpha} \int_0^\alpha v_{u_t}(\mathbf{x}; \alpha') d\alpha' \equiv \text{ucb}_t(\mathbf{x}).$$

In addition, noting that $v_{u_t}(\mathbf{x}; \alpha') - v_{l_t}(\mathbf{x}; \alpha') \leq \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w})$, we get

$$0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) = \frac{1}{\alpha} \int_0^\alpha (v_{u_t}(\mathbf{x}; \alpha') - v_{l_t}(\mathbf{x}; \alpha')) d\alpha' \leq \frac{1}{\alpha} \int_0^\alpha \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}) d\alpha' = \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}).$$

Mean Absolute Deviation, Standard Deviation and Variance From $l_{t,\mathbf{x},\mathbf{w}} \leq f(\mathbf{x}, \mathbf{w}) \leq u_{t,\mathbf{x},\mathbf{w}}$, we get

$$-\mathbb{E}[u_{t,\mathbf{x},\mathbf{w}}] \leq -\mathbb{E}[f(\mathbf{x}, \mathbf{w})] \leq -\mathbb{E}[l_{t,\mathbf{x},\mathbf{w}}].$$

Hence, we have

$$\check{l}_{t,\mathbf{x},\mathbf{w}} \equiv l_{t,\mathbf{x},\mathbf{w}} - \mathbb{E}[u_{t,\mathbf{x},\mathbf{w}}] \leq f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})] \leq u_{t,\mathbf{x},\mathbf{w}} - \mathbb{E}[l_{t,\mathbf{x},\mathbf{w}}] \equiv \check{u}_{t,\mathbf{x},\mathbf{w}}.$$

Therefore, we obtain

$$|f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]| \leq \max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\}.$$

Similarly, if $\check{l}_{t,\mathbf{x},\mathbf{w}} < 0$ and $\check{u}_{t,\mathbf{x},\mathbf{w}} > 0$, then we have

$$|f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]| \geq 0.$$

On the other hand, if $\check{l}_{t,\mathbf{x},\mathbf{w}} \geq 0$ or $\check{u}_{t,\mathbf{x},\mathbf{w}} \leq 0$, then we get

$$|f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]| \geq \min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\}.$$

Thus, by combining these, for any $\check{l}_{t,\mathbf{x},\mathbf{w}}$ and $\check{u}_{t,\mathbf{x},\mathbf{w}}$, we obtain

$$\begin{aligned} |f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]| &\geq \min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\} - \max\{\min\{-\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}}\}, 0\} \\ &\equiv \min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\} - \text{STR}(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}}). \end{aligned}$$

Hence, we have

$$\begin{aligned} \text{lcb}_t(\mathbf{x}) \equiv \mathbb{E}[\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\} - \text{STR}(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}})] &\leq \mathbb{E}[|f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]|] \\ &\leq \mathbb{E}[\max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\}] \equiv \text{ucb}_t(\mathbf{x}). \end{aligned}$$

Moreover, noting that

$$\begin{aligned} &\max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\} - (\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|, |\check{u}_{t,\mathbf{x},\mathbf{w}}|\} - \text{STR}(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}})) \\ &\leq \check{u}_{t,\mathbf{x},\mathbf{w}} - \check{l}_{t,\mathbf{x},\mathbf{w}} = (u_{t,\mathbf{x},\mathbf{w}} - l_{t,\mathbf{x},\mathbf{w}}) + \mathbb{E}[u_{t,\mathbf{x},\mathbf{w}} - l_{t,\mathbf{x},\mathbf{w}}] \\ &= 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}) + \mathbb{E}[2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w})] \leq 2 \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}), \end{aligned}$$

we obtain

$$0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq \mathbb{E}[2 \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w})] = 2 \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}).$$

Next, we prove the case of the standard deviation. By using the same argument as in the mean absolute deviation, we get

$$\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\} - \text{STR}^2(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}}) \leq |f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]|^2 \leq \max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\}.$$

Therefore, we have

$$\begin{aligned} \text{lcb}_t(\mathbf{x}) &\equiv \sqrt{\mathbb{E}[\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\} - \text{STR}^2(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}})]} \\ &\leq \sqrt{\mathbb{E}[|f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]|^2]} \leq \sqrt{\mathbb{E}[\max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\}]} \equiv \text{ucb}_t(\mathbf{x}). \end{aligned}$$

In addition, noting that $\sqrt{u} - \sqrt{v} \leq \sqrt{u-v}$ for any $u \geq v \geq 0$, we obtain

$$0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq \sqrt{\mathbb{E}[\max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\}] - \mathbb{E}[\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\} - \text{STR}^2(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}})]}.$$

From Equation (17) of Appendix A.2 in Iwazaki et al. (2021b), we have

$$\begin{aligned} &\mathbb{E}[\max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\}] - \mathbb{E}[\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\} - \text{STR}^2(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}})] \\ &\leq 16B\beta_{t+1}^{1/2}\mathbb{E}[\sigma_t(\mathbf{x}, \mathbf{w})] + 20\beta_{t+1}\mathbb{E}[\sigma_t^2(\mathbf{x}, \mathbf{w})] \\ &\leq 16B\beta_{t+1}^{1/2} \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} \sigma_t(\mathbf{x}, \mathbf{w}) + 20\beta_{t+1} \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} \sigma_t^2(\mathbf{x}, \mathbf{w}) = 8B \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) + 5 \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) \right)^2. \end{aligned}$$

Hence, we get

$$0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq \sqrt{8B \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) + 5 \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) \right)^2}.$$

Finally, we prove the case of the variance. By using the same argument as in the standard deviation, we get

$$\begin{aligned} \text{lcb}_t(\mathbf{x}) &\equiv \mathbb{E}[\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\} - \text{STR}^2(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}})] \\ &\leq \mathbb{E}[|f(\mathbf{x}, \mathbf{w}) - \mathbb{E}[f(\mathbf{x}, \mathbf{w})]|^2] \leq \mathbb{E}[\max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\}] \equiv \text{ucb}_t(\mathbf{x}). \end{aligned}$$

Furthermore, we obtain

$$\begin{aligned} 0 \leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) &= \mathbb{E}[\max\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\}] - \mathbb{E}[\min\{|\check{l}_{t,\mathbf{x},\mathbf{w}}|^2, |\check{u}_{t,\mathbf{x},\mathbf{w}}|^2\} - \text{STR}^2(\check{l}_{t,\mathbf{x},\mathbf{w}}, \check{u}_{t,\mathbf{x},\mathbf{w}})] \\ &\leq 16B\beta_{t+1}^{1/2}\mathbb{E}[\sigma_t(\mathbf{x}, \mathbf{w})] + 20\beta_{t+1}\mathbb{E}[\sigma_t^2(\mathbf{x}, \mathbf{w})] \\ &\leq 16B\beta_{t+1}^{1/2} \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} \sigma_t(\mathbf{x}, \mathbf{w}) + 20\beta_{t+1} \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} \sigma_t^2(\mathbf{x}, \mathbf{w}) \\ &= 8B \max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) + 5 \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2}\sigma_t(\mathbf{x}, \mathbf{w}) \right)^2. \end{aligned}$$

Distributionally Robust Let P be a candidate distribution of $P_{\mathbf{w}}(\mathbf{x})$, and let \mathcal{A} be a family of candidate distributions. Also let $F(\mathbf{x}; P)$, $\text{lcb}_t(\mathbf{x}; P)$ and $\text{ucb}_t(\mathbf{x}; P)$ be respectively risk measure, and its lower and upper with respect to P . Define

$$F(\mathbf{x}) \equiv \inf_{P \in \mathcal{A}} F(\mathbf{x}; P), \quad \text{lcb}_t(\mathbf{x}) \equiv \inf_{P \in \mathcal{A}} \text{lcb}_t(\mathbf{x}; P), \quad \text{ucb}_t(\mathbf{x}) \equiv \inf_{P \in \mathcal{A}} \text{ucb}_t(\mathbf{x}; P).$$

From the property of infimum, for any $\epsilon > 0$, there exists a distribution \hat{P} such that

$$\text{ucb}_t(\mathbf{x}; \hat{P}) \leq \text{ucb}_t(\mathbf{x}) + \epsilon.$$

Hence, we get

$$F(\mathbf{x}) \leq F(\mathbf{x}; \hat{P}) \leq \text{ucb}_t(\mathbf{x}; \hat{P}) \leq \text{ucb}_t(\mathbf{x}) + \epsilon.$$

Since ϵ is an arbitrary positive number, we obtain

$$F(\mathbf{x}) \leq \text{ucb}_t(\mathbf{x}).$$

Similarly, we also get

$$\text{lcb}_t(\mathbf{x}) \leq F(\mathbf{x}).$$

Furthermore, for any $\eta > 0$, there exists a distribution \tilde{P} such that

$$\text{lcb}_t(\mathbf{x}; \tilde{P}) \leq \text{lcb}_t(\mathbf{x}) + \eta.$$

Thus, we get

$$\text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq \text{ucb}_t(\mathbf{x}; \tilde{P}) - \text{lcb}_t(\mathbf{x}; \tilde{P}) + \eta \leq q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}); F \right) + \eta.$$

Since η is an arbitrary positive number, we have

$$\text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) \leq q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}); F \right).$$

Monotone Lipschitz Map Let \mathcal{M} be a K -Lipschitz map, and let $F(\mathbf{x})$, $\text{lcb}_t(\mathbf{x})$ and $\text{ucb}_t(\mathbf{x})$ be respectively risk measure, and its lower and upper. Then, from the monotonicity of \mathcal{M} , we have

$$\min\{\mathcal{M}(\text{lcb}_t(\mathbf{x})), \mathcal{M}(\text{ucb}_t(\mathbf{x}))\} \leq \mathcal{M}(F(\mathbf{x})) \leq \max\{\mathcal{M}(\text{lcb}_t(\mathbf{x})), \mathcal{M}(\text{ucb}_t(\mathbf{x}))\}.$$

In addition, using the Lipschitz continuity of \mathcal{M} we get

$$\begin{aligned} 0 &\leq \max\{\mathcal{M}(\text{lcb}_t(\mathbf{x})), \mathcal{M}(\text{ucb}_t(\mathbf{x}))\} - \min\{\mathcal{M}(\text{lcb}_t(\mathbf{x})), \mathcal{M}(\text{ucb}_t(\mathbf{x}))\} \\ &\leq |\mathcal{M}(\text{lcb}_t(\mathbf{x})) - \mathcal{M}(\text{ucb}_t(\mathbf{x}))| \leq K|\text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x})| \leq Kq \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}) \right). \end{aligned}$$

Weighted Sum Let $\alpha_1, \alpha_2 \geq 0$, and let $F_i(\mathbf{x})$, $\text{lcb}_{t,i}(\mathbf{x})$ and $\text{ucb}_{t,i}(\mathbf{x})$ be respectively risk measure, and its lower and upper with $i = 1, 2$. Then, noting that $\alpha_1, \alpha_2 \geq 0$, we obtain

$$\text{lcb}_t(\mathbf{x}) \equiv \alpha_1 \text{lcb}_{t,1}(\mathbf{x}) + \alpha_2 \text{lcb}_{t,2}(\mathbf{x}) \leq \alpha_1 F_1(\mathbf{x}) + \alpha_2 F_2(\mathbf{x}) \leq \alpha_1 \text{ucb}_{t,1}(\mathbf{x}) + \alpha_2 \text{ucb}_{t,2}(\mathbf{x}) \equiv \text{ucb}_t(\mathbf{x}).$$

Moreover, we get

$$\begin{aligned} 0 &\leq \text{ucb}_t(\mathbf{x}) - \text{lcb}_t(\mathbf{x}) = \alpha_1 (\text{ucb}_{t,1}(\mathbf{x}) - \text{lcb}_{t,1}(\mathbf{x})) + \alpha_2 (\text{ucb}_{t,2}(\mathbf{x}) - \text{lcb}_{t,2}(\mathbf{x})) \\ &\leq \alpha_1 q_1 \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}) \right) + \alpha_2 q_2 \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}}} 2\beta_{t+1}^{1/2} \sigma_t(\mathbf{x}, \mathbf{w}) \right). \end{aligned}$$

Probabilistic Threshold Let $\theta \in \mathbb{R}$ be a threshold. Then, $l_{t,\mathbf{x},\mathbf{w}} \leq f(\mathbf{x}, \mathbf{w}) \leq u_{t,\mathbf{x},\mathbf{w}}$ implies that

$$\text{lcb}_t(\mathbf{x}) \equiv \mathbb{P}(l_{t,\mathbf{x},\mathbf{w}} \geq \theta) \leq \mathbb{P}(f(\mathbf{x}, \mathbf{w}) \geq \theta) \leq \mathbb{P}(u_{t,\mathbf{x},\mathbf{w}} \geq \theta) \equiv \text{ucb}_t(\mathbf{x}).$$

B.2 Extension of Theorem E.4 in Kusakawa et al. (2022)

We show the extension of Theorem E.4 in Kusakawa et al. (2022). In this subsection, we use \mathbf{x} and \mathcal{X} as the input variable and set of all input variables, respectively. In Theorem E.4 in Kusakawa et al. (2022), they proved that if linear, Gaussian or Matérn (with parameter $\nu > 1$) kernel is used, then the posterior standard deviation satisfies the L_σ -data-independent-Lipschitz continuity. They have assumed that the variance of an error distribution for GP models is $\sigma^2 > 0$ for any input \mathbf{x} . We show that this assumption can be relaxed to the assumption that the noise variance is positive and depends on \mathbf{x} . Since the relaxation of noise variance to the heteroscedastic setting does not affect any essential part of their proof, only the sketch of the proof is given here. Let \mathbf{X}_t be a $t \times d$ matrix. Then, in their proof, σ^2 appears only within the formula given below:

$$\mathbf{I}_d - \mathbf{X}_t^\top (\mathbf{X}_t \mathbf{X}_t^\top + a^{-2} \sigma^2 \mathbf{I}_t)^{-1} \mathbf{X}_t,$$

where a is some positive constant. They considered the singular value decomposition $\mathbf{X}_t = \mathbf{H}' \mathbf{\Lambda} \mathbf{H}$ and calculated

$$\mathbf{I}_d - \mathbf{X}_t^\top (\mathbf{X}_t \mathbf{X}_t^\top + a^{-2} \sigma^2 \mathbf{I}_t)^{-1} \mathbf{X}_t = \mathbf{H} \mathbf{\Theta} \mathbf{H}^\top,$$

where Θ is the diagonal matrix whose (j, j) -th element θ_j satisfies $0 \leq \theta_j \leq 1$. In their proof, only the fact that \mathbf{H} is an orthogonal matrix and $0 \leq \theta_j \leq 1$. On the other hand, when the noise variance is heteroscedastic, that is, the variance is expressed as s_t^2 at iteration t , we have to consider the following:

$$\mathbf{I}_d - \mathbf{X}_t^\top (\mathbf{X}_t \mathbf{X}_t^\top + a^{-2} \mathbf{S}_t)^{-1} \mathbf{X}_t,$$

where \mathbf{S}_t is the diagonal matrix whose (j, j) -th element is s_t^2 . Also in this case, noting that

$$\begin{aligned} \mathbf{I}_d - \mathbf{X}_t^\top (\mathbf{X}_t \mathbf{X}_t^\top + a^{-2} \mathbf{S}_t)^{-1} \mathbf{X}_t &= \mathbf{I}_d - \mathbf{X}_t^\top (\mathbf{S}_t^{1/2} \{ \mathbf{S}_t^{-1/2} \mathbf{X}_t \mathbf{X}_t^\top \mathbf{S}_t^{-1/2} + a^{-2} \mathbf{I}_t \} \mathbf{S}_t^{1/2})^{-1} \mathbf{X}_t \\ &= \mathbf{I}_d - \mathbf{X}_t^\top \mathbf{S}_t^{-1/2} (\mathbf{S}_t^{-1/2} \mathbf{X}_t \mathbf{X}_t^\top \mathbf{S}_t^{-1/2} + a^{-2} \mathbf{I}_t)^{-1} \mathbf{S}_t^{-1/2} \mathbf{X}_t \\ &= \mathbf{I}_d - \tilde{\mathbf{X}}_t^\top (\tilde{\mathbf{X}}_t \tilde{\mathbf{X}}_t^\top + a^{-2} \mathbf{I}_t)^{-1} \tilde{\mathbf{X}}_t, \end{aligned}$$

using the singular value decomposition $\tilde{\mathbf{X}}_t = \tilde{\mathbf{H}}' \tilde{\Lambda} \tilde{\mathbf{H}}$ we have

$$\mathbf{I}_d - \tilde{\mathbf{X}}_t^\top (\tilde{\mathbf{X}}_t \tilde{\mathbf{X}}_t^\top + a^{-2} \mathbf{I}_t)^{-1} \tilde{\mathbf{X}}_t = \tilde{\mathbf{H}} \tilde{\Theta} \tilde{\mathbf{H}}^\top,$$

where $\tilde{\mathbf{H}}$ is an orthogonal matrix and the (j, j) -th element $\tilde{\theta}_j$ of the diagonal matrix $\tilde{\Theta}$ satisfies $0 \leq \tilde{\theta}_j \leq 1$. Therefore, also in the heteroscedastic setting, L_σ -data-independent-Lipschitz continuity holds.

B.3 Proof of Lemma A.1

Let $\mathbf{UCB}_t(\mathbf{x}) = (u_1, \dots, u_L) \equiv \mathbf{u}$ and $\mathbf{LCB}_t(\hat{\Pi}_t) = \{(l_1^{(i)}, \dots, l_L^{(i)}) \mid 1 \leq i \leq k\} \equiv \mathcal{L}$. Here, if $\mathbf{u} \in \text{Dom}(\mathcal{L})$, then the following holds from the definition of $\text{dist}(\mathbf{a}, B)$:

$$a_t^{(\mathcal{X})}(\mathbf{x}) = \text{dist}(\mathbf{u}, \text{Dom}(\mathcal{L})) = \inf_{\mathbf{b} \in \text{Dom}(\mathcal{L})} d_\infty(\mathbf{u}, \mathbf{b}) = d_\infty(\mathbf{u}, \mathbf{u}) = 0.$$

In addition, since $\mathbf{u} \in \text{Dom}(\mathcal{L})$, there exists $(l_1^{(i)}, \dots, l_L^{(i)})$ such that $u_j \leq l_j^{(i)}$ for any $j \in [L]$. Thus, we have $\max\{u_1 - l_1^{(i)}, \dots, u_L - l_L^{(i)}\} \leq 0$. This implies that

$$\tilde{a}_t(\mathbf{x}) = \min_{1 \leq i \leq k} \max\{u_1 - l_1^{(i)}, \dots, u_L - l_L^{(i)}\} \leq 0$$

and $\max\{\tilde{a}_t(\mathbf{x}), 0\} = 0$. Therefore, we get $a_t^{(\mathcal{X})}(\mathbf{x}) = \max\{\tilde{a}_t(\mathbf{x}), 0\}$. Next, we consider the case where $\mathbf{u} \notin \text{Dom}(\mathcal{L})$. Let $a_t^{(\mathcal{X})}(\mathbf{x}) = \eta$. Then, noting that $\mathbf{u} \notin \text{Dom}(\mathcal{L})$, for any $i \in \{1, \dots, k\}$, there exists $j \in [L]$ such that $u_j > l_j^{(i)}$. This implies that

$$\tilde{a}_t(\mathbf{x}) = \min_{1 \leq i \leq k} \max\{u_1 - l_1^{(i)}, \dots, u_L - l_L^{(i)}\} \equiv \tilde{\eta} > 0$$

and $\max\{\tilde{a}_t(\mathbf{x}), 0\} = \tilde{a}_t(\mathbf{x}) = \tilde{\eta}$. For this $\tilde{\eta}$, there exists i such that

$$u_j - l_j^{(i)} \leq \tilde{\eta} \quad \forall j \in [L].$$

Hence, we have $\tilde{\mathbf{u}} \equiv (u_1 - \tilde{\eta}, \dots, u_L - \tilde{\eta}) \in \text{Dom}(\mathcal{L})$ because $u_j - \tilde{\eta} \leq l_j^{(i)}$ for any $j \in [L]$. Thus, from the definition of $a_t^{(\mathcal{X})}(\mathbf{x})$, the following holds:

$$\eta = a_t^{(\mathcal{X})}(\mathbf{x}) = \text{dist}(\mathbf{u}, \text{Dom}(\mathcal{L})) = \inf_{\mathbf{b} \in \text{Dom}(\mathcal{L})} d_\infty(\mathbf{u}, \mathbf{b}) \leq d_\infty(\mathbf{u}, \tilde{\mathbf{u}}) = \tilde{\eta}.$$

Here, we assume $\eta < \tilde{\eta}$. Then, noting that $\text{Dom}(\mathcal{L})$ is the closed set, there exists $\tilde{\mathbf{l}} = (\tilde{l}_1, \dots, \tilde{l}_L) \in \text{Dom}(\mathcal{L})$ such that $d_\infty(\mathbf{u}, \tilde{\mathbf{l}}) = \eta$. Therefore, $\tilde{\mathbf{l}}$ can be expressed as $\tilde{\mathbf{l}} = (u_1 - s_1, \dots, u_L - s_L)$, where $0 \leq |s_j| \leq \eta$ and at least one of s_1, \dots, s_L is η . Thus, since $(u_1 - \eta, \dots, u_L - \eta) \leq \tilde{\mathbf{l}}$, noting that $(u_1 - \eta, \dots, u_L - \eta) \in \text{Dom}(\mathcal{L})$ there exists i such that

$$u_j - \eta \leq l_j^{(i)} \quad \forall j \in [L].$$

This implies that $\max\{u_1 - l_1^{(i)}, \dots, u_L - l_L^{(i)}\} \leq \eta$. Hence, it follows that

$$\tilde{\eta} = \min_{1 \leq i \leq k} \max\{u_1 - l_1^{(i)}, \dots, u_L - l_L^{(i)}\} \leq \eta.$$

However, this is a contradiction with $\eta < \tilde{\eta}$. Consequently, we obtain $a_t^{(\mathcal{X})}(\mathbf{x}) = \max\{\tilde{a}_t(\mathbf{x}), 0\}$.

B.4 Proof of Theorem A.1

From the theorem's assumption, the bounding box $\tilde{B}_t(\mathbf{x})$ is HPBB. Therefore, with probability at least $1 - \delta$, the following holds for any $t \geq 0$:

$$\text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t)) \subset \text{Dom}(\mathbf{F}(\hat{\Pi}_t)) \subset \text{Dom}(Z^*) \subset \text{Dom}(\mathbf{UCB}_t(\mathcal{X})).$$

Hence, using this, noting that the definition of $d_\infty(\cdot, \cdot)$, we get

$$\begin{aligned} I_t^{(i)} &= \max_{\mathbf{y} \in Z^*} \min_{\mathbf{y}' \in \text{Par}(\mathbf{F}(\hat{\Pi}_t))} d_\infty(\mathbf{y}, \mathbf{y}') \leq \max_{\mathbf{y}' \in \text{Par}(\mathbf{UCB}_t(\mathcal{X}))} \min_{\mathbf{y}' \in \text{Par}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{y}, \mathbf{y}') \\ &= \max_{\mathbf{y}' \in \text{Par}(\mathbf{UCB}_t(\mathcal{X}))} \min_{\mathbf{y}' \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{y}, \mathbf{y}') \\ &= \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y}' \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{UCB}_t(\mathbf{x}), \mathbf{y}') = \max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}). \end{aligned}$$

Similarly, we get

$$\begin{aligned} I_t^{(ii)} &= \max_{\mathbf{y}' \in \mathbf{F}(\hat{\Pi}_t)} \min_{\mathbf{y} \in Z^*} d_\infty(\mathbf{y}, \mathbf{y}') \leq \max_{\mathbf{y}' \in \text{Par}(\mathbf{UCB}_t(\mathcal{X}))} \min_{\mathbf{y} \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{y}', \mathbf{y}) \\ &= \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y} \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{UCB}_t(\mathbf{x}), \mathbf{y}) = \max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}). \end{aligned}$$

Thus, we have $I_t = \max\{I_t^{(i)}, I_t^{(ii)}\} \leq \max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}) = a_t^{(\mathcal{X})}(\mathbf{x}_{t+1})$. Hence, if $a_T^{(\mathcal{X})}(\mathbf{x}_{T+1}) \leq \epsilon$, then $I_T \leq \epsilon$.

B.5 Proof of Theorem A.2

From the definition of $a_t^{(\mathcal{X})}(\mathbf{x})$, \mathbf{x}_{t+1} and \mathbf{w}_{t+1} , noting that $\mathbf{LCB}_t(\mathbf{x}_{t+1}) \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))$ we get

$$\begin{aligned} a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) &\leq \|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty \leq q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right) \\ &= q \left(\sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}) \right). \end{aligned}$$

Let $\hat{t} = \text{argmin}_{0 \leq i \leq t} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1})$. Then, the following inequality holds:

$$\begin{aligned} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,\hat{t}+1}^{1/2} \tilde{\sigma}_{\hat{t}}^{(m)}(\mathbf{x}_{\hat{t}+1}, \mathbf{w}_{\hat{t}+1}) &\leq \frac{1}{t+1} \sum_{i=1}^{t+1} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,i}^{1/2} \tilde{\sigma}_{i-1}^{(m)}(\mathbf{x}_i, \mathbf{w}_i) \\ &\leq \frac{1}{t+1} \sqrt{(t+1) \sum_{i=1}^{t+1} \sum_{m=1}^{M_f} 4M_f \tilde{\beta}_{m,i} \tilde{\sigma}_{i-1}^{(m)2}(\mathbf{x}_i, \mathbf{w}_i)} \\ &\leq \frac{1}{t+1} \sqrt{(t+1) \sum_{m=1}^{M_f} 4M_f \tilde{\beta}_{m,t+1} \sum_{i=1}^{t+1} \tilde{\sigma}_{i-1}^{(m)2}(\mathbf{x}_i, \mathbf{w}_i)} \\ &\leq \frac{1}{t+1} \sqrt{(t+1) \sum_{m=1}^{M_f} 4M_f \tilde{\beta}_{m,t+1} \frac{2}{\log(1 + \lambda_m^{-1} \underline{\tau}^{-2})} \tilde{\kappa}_{t+1}^{(m)}} = \sqrt{\frac{\sum_{m=1}^{M_f} \tilde{C}_m \tilde{\beta}_{m,t+1} \tilde{\kappa}_{t+1}^{(m)}}{t+1}}, \end{aligned}$$

where the second inequality is derived by Cauchy-Schwarz inequality and $(a_1 + \dots + a_{M_f})^2 \leq M_f(a_1^2 + \dots + a_{M_f}^2)$, the third inequality is derived by monotonicity of $\tilde{\beta}_{m,t}$, and the fourth inequality is derived by the definition of the maximum information gain, $s^2 \leq (\varsigma^{-2} / \log(1 + \varsigma^{-2})) \log(1 + s^2)$ for $s^2 \in [0, \varsigma^{-2}]$, and $\tau_{\mathbf{x}_i, \mathbf{w}_i, m}^{-2} \tilde{\sigma}_{i-1}^{(m)2}(\mathbf{x}_i, \mathbf{w}_i) \leq \tau_{\mathbf{x}_i, \mathbf{w}_i, m}^{-2} \lambda_m^{-1}$. Therefore, we obtain

$$\max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}) = a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \leq q \left(\sum_{m=1}^{M_f} 2\tilde{\beta}_{m,\hat{t}+1}^{1/2} \tilde{\sigma}_{\hat{t}}^{(m)}(\mathbf{x}_{\hat{t}+1}, \mathbf{w}_{\hat{t}+1}) \right) \leq q \left(\sqrt{\frac{\sum_{m=1}^{M_f} \tilde{C}_m \tilde{\beta}_{m,t+1} \tilde{\kappa}_{t+1}^{(m)}}{t+1}} \right) = q(s_t).$$

Thus, for some $T \geq 0$ satisfying $q(s_T) \leq \epsilon$, there exists $\hat{T} \leq T$ such that $a_{\hat{T}}^{(\mathcal{X})}(\mathbf{x}_{\hat{T}+1}) \leq q(s_T) \leq \epsilon$. Noting that $0 \leq \hat{T} \leq T$, the algorithm terminates after at most T iterations.

B.6 Proof of Theorem A.3

From the definition of $q(a)$, since $q^{(m,l)}(a)$ is a strictly increasing function satisfying $q^{(m,l)}(0) = 0$, $q(a)$ is a strictly increasing function and satisfies $q(0) = 0$. Furthermore, noting that $\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \leq \max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w})$, since $q^{(m,l)}(a)$ is a strictly increasing function, we get

$$\begin{aligned} \|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty &= \max_{m \in [M_f], l \in [L_m]} |\text{ucb}_t^{(m,l)}(\mathbf{x}_{t+1}) - \text{lcb}_t^{(m,l)}(\mathbf{x}_{t+1})| \\ &\leq \max_{m \in [M_f], l \in [L_m]} q^{(m,l)} \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right) \\ &\leq \max_{m \in [M_f], l \in [L_m]} q^{(m,l)} \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right) \\ &= q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right). \end{aligned}$$

B.7 Proof of Theorem A.4

Let r be a number. For any vector $\mathbf{a} = (a_1, \dots, a_s)$ and subset $B \subset \mathbb{R}^s$, we define $r + \mathbf{a} \equiv (r + a_1, \dots, r + a_s)$ and $r + B \equiv \{r + \mathbf{b} \mid \mathbf{b} \in B\}$. Then, from the theorem's assumption, with probability at least $1 - \delta$, the following holds for any $t \geq 0$:

$$\text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t) - \epsilon_{\text{lcb}}) \subset \text{Dom}(\mathbf{F}(\hat{\Pi}_t)) \subset \text{Dom}(Z^*) \subset \text{Dom}(\mathbf{UCB}_t(\mathcal{X}) + \epsilon_{\text{ucb}}).$$

Hence, using this, noting that the definition of $d_\infty(\cdot, \cdot)$, we get

$$\begin{aligned} I_t^{(i)} &= \max_{\mathbf{y} \in Z^*} \min_{\mathbf{y}' \in \text{Par}(\mathbf{F}(\hat{\Pi}_t))} d_\infty(\mathbf{y}, \mathbf{y}') \leq \max_{\mathbf{y} \in \text{Par}(\mathbf{UCB}_t(\mathcal{X}) + \epsilon_{\text{ucb}})} \min_{\mathbf{y}' \in \text{Par}(\mathbf{LCB}_t(\hat{\Pi}_t) - \epsilon_{\text{lcb}})} d_\infty(\mathbf{y}, \mathbf{y}') \\ &= \max_{\mathbf{y} \in \text{Par}(\mathbf{UCB}_t(\mathcal{X}) + \epsilon_{\text{ucb}})} \min_{\mathbf{y}' \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t) - \epsilon_{\text{lcb}})} d_\infty(\mathbf{y}, \mathbf{y}') \\ &= \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y}' \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t) - \epsilon_{\text{lcb}})} d_\infty(\mathbf{UCB}_t(\mathbf{x}) + \epsilon_{\text{ucb}}, \mathbf{y}') \\ &\leq \epsilon_{\text{ucb}} + \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y}' \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t) - \epsilon_{\text{lcb}})} d_\infty(\mathbf{UCB}_t(\mathbf{x}), \mathbf{y}') \\ &\leq \epsilon_{\text{ucb}} + \epsilon_{\text{lcb}} + \max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{y}' \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))} d_\infty(\mathbf{UCB}_t(\mathbf{x}), \mathbf{y}') \\ &= \epsilon_{\text{ucb}} + \epsilon_{\text{lcb}} + \max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}) \\ &= \epsilon_{\text{ucb}} + \epsilon_{\text{lcb}} + a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) + \max_{\mathbf{x} \in \mathcal{X}} a_t^{(\mathcal{X})}(\mathbf{x}) - a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) \\ &\leq \epsilon_{\text{ucb}} + \epsilon_{\text{lcb}} + \epsilon_{\mathcal{X}} + a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}). \end{aligned}$$

Similarly, we get

$$\begin{aligned} I_t^{(ii)} &= \max_{\mathbf{y} \in \mathbf{F}(\hat{\Pi}_t)} \min_{\mathbf{y}' \in Z^*} d_\infty(\mathbf{y}, \mathbf{y}') \leq \max_{\mathbf{y}' \in \text{Par}(\mathbf{UCB}_t(\mathcal{X}) + \epsilon_{\text{ucb}})} \min_{\mathbf{y} \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t) - \epsilon_{\text{lcb}})} d_\infty(\mathbf{y}', \mathbf{y}) \\ &\leq \epsilon_{\text{ucb}} + \epsilon_{\text{lcb}} + \epsilon_{\mathcal{X}} + a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}). \end{aligned}$$

Thus, we have $I_t = \max\{I_t^{(i)}, I_t^{(ii)}\} \leq \epsilon_{\text{ucb}} + \epsilon_{\text{lcb}} + \epsilon_{\mathcal{X}} + a_t^{(\mathcal{X})}(\mathbf{x}_{t+1})$. Hence, if $a_T^{(\mathcal{X})}(\mathbf{x}_{T+1}) \leq \epsilon$, then $I_T \leq \epsilon + \epsilon_{\text{ucb}} + \epsilon_{\text{lcb}} + \epsilon_{\mathcal{X}}$.

B.8 Proof of Theorem A.5

From the definition of $a_t^{(\mathcal{X})}(\mathbf{x})$, \mathbf{x}_{t+1} and \mathbf{w}_{t+1} , noting that $-\epsilon_{\text{PF}} + \mathbf{LCB}_t(\mathbf{x}_{t+1}) \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))$ we get

$$\begin{aligned} a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) &\leq \|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - (\mathbf{LCB}_t(\mathbf{x}_{t+1}) - \epsilon_{\text{PF}})\|_\infty \leq \epsilon_{\text{PF}} + \|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty \\ &\leq \epsilon_{\text{PF}} + q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right) \\ &\leq \epsilon_{\text{PF}} + q \left(\epsilon_\Omega + \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}) \right). \end{aligned}$$

Thus, by letting $\hat{t} = \text{argmin}_{0 \leq i \leq t} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1})$, using the same argument as in the proof of Theorem A.2, we have the desired result.

B.9 Proof of Theorem A.6

From the definition of $a_t^{(\mathcal{X})}(\mathbf{x})$, \mathbf{x}_{t+1} and \mathbf{w}_{t+1} , noting that $\mathbf{LCB}_t(\mathbf{x}_{t+1}) \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))$ we get

$$\begin{aligned} a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) &\leq \|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty \leq q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right) \\ &= q \left(\sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}^*) \right). \end{aligned}$$

Let $\mathcal{S}(\Omega_{\mathbf{x}_{t+1}}; \zeta_{t+1})$ be a maximal ζ_{t+1} -separated subset of $\Omega_{\mathbf{x}_{t+1}}$. Then, from the definition of $\mathcal{S}(\Omega_{\mathbf{x}_{t+1}}; \zeta_{t+1})$, there exists a point $\tilde{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}_{t+1}}; \zeta_{t+1})$ such that $\|\mathbf{w}_{t+1}^* - \tilde{\mathbf{w}}\|_1 \leq \zeta_{t+1}$. Hence, from the L_σ -data-independent Lipschitz continuity, we obtain

$$\begin{aligned} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}^*) &= \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}}) + \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \{\tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}^*) - \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}})\} \\ &\leq \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} L_\sigma \zeta_{t+1} + \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}}) \\ &\leq 2M_f \tilde{\beta}_{t+1}^{1/2} L_\sigma \zeta_{t+1} + \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}}). \end{aligned}$$

In addition, we get

$$\begin{aligned} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}}) &\leq \sum_{\tilde{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}_{t+1}}; \zeta_{t+1})} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}}) \\ &\leq \underline{p}_{\zeta_{t+1}}^{-1} \sum_{\tilde{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}_{t+1}}; \zeta_{t+1})} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}}) \mathbb{P}_{P_{\mathbf{w}}(\mathbf{x}_{t+1})}[\mathbf{w} \in \text{Nei}(\tilde{\mathbf{w}}; \zeta_{t+1}/2)] \\ &\leq 2\tilde{\beta}_{t+1}^{1/2} \underline{p}_{\zeta_{t+1}}^{-1} \mathbb{E}_{P_{\mathbf{w}}(\mathbf{x}_{t+1})} \left[\sum_{\tilde{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}_{t+1}}; \zeta_{t+1})} \sum_{m=1}^{M_f} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \tilde{\mathbf{w}}) \mathbb{1}[\mathbf{w} \in \text{Nei}(\tilde{\mathbf{w}}; \zeta_{t+1}/2)] \right] \\ &\equiv 2\tilde{\beta}_{t+1}^{1/2} \underline{p}_{\zeta_{t+1}}^{-1} \mathbb{E}_{P_{\mathbf{w}}(\mathbf{x}_{t+1})} [S(\mathbf{x}_{t+1}, \mathbf{w})], \end{aligned}$$

where $\mathbb{1}[\cdot]$ represents the indicator function. Thus, we have

$$\sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{t+1}^*) \leq 2M_f \tilde{\beta}_{t+1}^{1/2} L_\sigma \zeta_{t+1} + 2\tilde{\beta}_{t+1}^{1/2} \underline{p}_{\zeta_{t+1}}^{-1} \mathbb{E}_{P_{\mathbf{w}}(\mathbf{x}_{t+1})} [S(\mathbf{x}_{t+1}, \mathbf{w})].$$

Therefore, we get

$$\begin{aligned} \sum_{i=0}^t \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,i+1}^{1/2} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}^*) &\leq \sum_{i=0}^t 2M_f \tilde{\beta}_{i+1}^{1/2} L_\sigma \zeta_{i+1} + \sum_{i=0}^t 2\tilde{\beta}_{i+1}^{1/2} \underline{p}_{\zeta_{i+1}}^{-1} \mathbb{E}_{P_{\mathbf{w}}(\mathbf{x}_{i+1})}[S(\mathbf{x}_{i+1}, \mathbf{w})] \\ &\leq 2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} \sum_{i=0}^t \zeta_{i+1} + 2\tilde{\beta}_{t+1}^{1/2} \underline{p}_{\zeta_{t+1}}^{-1} \sum_{i=0}^t \mathbb{E}_{P_{\mathbf{w}}(\mathbf{x}_{i+1})}[S(\mathbf{x}_{i+1}, \mathbf{w})]. \end{aligned}$$

Here, $S(\mathbf{x}_{i+1}, \mathbf{w})$ is the non-negative random variable satisfying $S(\mathbf{x}_{i+1}, \mathbf{w}) \leq M_f \max\{1, \lambda_1^{-1}, \dots, \lambda_m^{-1}\} = J$. Hence, from Lemma 3 in Kirschner and Krause (2018), with probability at least $1 - \delta$, the following holds for any $i \geq 0$:

$$\sum_{i=0}^t \mathbb{E}_{P_{\mathbf{w}}(\mathbf{x}_{i+1})}[S(\mathbf{x}_{i+1}, \mathbf{w})] \leq 4J \log \frac{1}{\delta} + 8J \log(4J) + 1 + 2 \sum_{i=0}^t S(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) \leq 8J \log \frac{8J}{\delta} + 2 \sum_{i=0}^t S(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}).$$

Furthermore, from the definition of $S(\mathbf{x}_{i+1}, \mathbf{w}_{i+1})$, we have

$$S(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) = \sum_{m=1}^{M_f} \sum_{\check{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}_{i+1}}; \zeta_{i+1})} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \check{\mathbf{w}}) \mathbb{1}[\mathbf{w}_{i+1} \in \text{Nei}(\check{\mathbf{w}}; \zeta_{i+1}/2)].$$

Noting that $\check{\mathbf{w}}_1 \neq \check{\mathbf{w}}_2 \Rightarrow \text{Nei}(\check{\mathbf{w}}_1; \zeta_{i+1}/2) \cap \text{Nei}(\check{\mathbf{w}}_2; \zeta_{i+1}/2) = \emptyset$, if there exists $\check{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}_{i+1}}; \zeta_{i+1})$ such that $\mathbf{w}_{i+1} \in \text{Nei}(\check{\mathbf{w}}; \zeta_{i+1}/2)$, then we obtain

$$\tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \check{\mathbf{w}}) = \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) + \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \check{\mathbf{w}}) - \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) \leq \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) + L_\sigma \zeta_{i+1}/2.$$

Similarly, if $\mathbf{w}_{i+1} \notin \text{Nei}(\check{\mathbf{w}}; \zeta_{i+1}/2)$ for any $\check{\mathbf{w}} \in \mathcal{S}(\Omega_{\mathbf{x}_{i+1}}; \zeta_{i+1})$, then we get

$$\tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \check{\mathbf{w}}) = 0 \leq \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) + L_\sigma \zeta_{i+1}/2.$$

Therefore, we have

$$2 \sum_{i=0}^t S(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) \leq M_f L_\sigma \sum_{i=0}^t \zeta_{i+1} + 2 \sum_{i=0}^t \sum_{m=1}^{M_f} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}).$$

By combining previous results, we obtain

$$\begin{aligned} &\sum_{i=0}^t \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,i+1}^{1/2} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}^*) \\ &\leq 2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} \sum_{i=0}^t \zeta_{i+1} + 2\tilde{\beta}_{t+1}^{1/2} \underline{p}_{\zeta_{t+1}}^{-1} \left(8J \log \frac{8J}{\delta} + 2 \sum_{i=0}^t S(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) \right) \\ &\leq 2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} \sum_{i=0}^t \zeta_{i+1} + 2\tilde{\beta}_{t+1}^{1/2} \underline{p}_{\zeta_{t+1}}^{-1} \left(8J \log \frac{8J}{\delta} + M_f L_\sigma \sum_{i=0}^t \zeta_{i+1} + 2 \sum_{i=0}^t \sum_{m=1}^{M_f} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}) \right) \\ &= 2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} (1 + \underline{p}_{\zeta_{t+1}}^{-1}) \sum_{i=0}^t \zeta_{i+1} + 16J \log \frac{8J}{\delta} \tilde{\beta}_{t+1}^{1/2} \underline{p}_{\zeta_{t+1}}^{-1} + 2\underline{p}_{\zeta_{t+1}}^{-1} \sum_{i=0}^t \sum_{m=1}^{M_f} 2\tilde{\beta}_{t+1}^{1/2} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}). \end{aligned}$$

Finally, let $\hat{t} = \operatorname{argmin}_{0 \leq i \leq t} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,i+1}^{1/2} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}^*)$. Then, the following inequality holds:

$$\begin{aligned} &\sum_{m=1}^{M_f} 2\tilde{\beta}_{m,\hat{t}+1}^{1/2} \tilde{\sigma}_{\hat{t}}^{(m)}(\mathbf{x}_{\hat{t}+1}, \mathbf{w}_{\hat{t}+1}^*) \\ &\leq \frac{2M_f L_\sigma \tilde{\beta}_{\hat{t}+1}^{1/2} (1 + \underline{p}_{\zeta_{\hat{t}+1}}^{-1})}{\hat{t} + 1} \sum_{i=0}^{\hat{t}} \zeta_{i+1} + \frac{16J \log \frac{8J}{\delta} \tilde{\beta}_{\hat{t}+1}^{1/2} \underline{p}_{\zeta_{\hat{t}+1}}^{-1}}{\hat{t} + 1} + \frac{2\underline{p}_{\zeta_{\hat{t}+1}}^{-1}}{\hat{t} + 1} \sum_{i=0}^{\hat{t}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{i+1}^{1/2} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}). \end{aligned}$$

In addition, by using the same argument as in the proof of Theorem A.2, we get

$$\begin{aligned}
 & \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}^*) \\
 & \leq \frac{2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} (1 + \tilde{p}_{\zeta_{t+1}}^{-1})}{t+1} \sum_{i=0}^t \zeta_{i+1} + \frac{16J \log \frac{8J}{\delta} \tilde{\beta}_{t+1}^{1/2} \tilde{p}_{\zeta_{t+1}}^{-1}}{t+1} + \sqrt{\frac{4\tilde{p}_{\zeta_{t+1}}^{-2} \sum_{m=1}^{M_f} \tilde{C}_m \tilde{\beta}_{t+1} \tilde{\kappa}_{t+1}}{t+1}} \\
 & = \frac{2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} (1 + \tilde{p}_{\zeta_{t+1}}^{-1})}{t+1} \sum_{i=0}^t \zeta_{i+1} + \frac{16J \log \frac{8J}{\delta} \tilde{\beta}_{t+1}^{1/2} \tilde{p}_{\zeta_{t+1}}^{-1}}{t+1} + \sqrt{\frac{\tilde{p}_{\zeta_{t+1}}^{-2} \tilde{\beta}_{t+1} \tilde{\kappa}_{t+1} \sum_{m=1}^{M_f} \hat{C}_m}{t+1}} \\
 & \leq \frac{2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} (1 + \tilde{p}_{\zeta_{t+1}}^{-1})}{t+1} \sum_{i=0}^t \zeta_{i+1} + \frac{16J \log \frac{8J}{\delta} \tilde{\beta}_{t+1}^{1/2} \tilde{p}_{\zeta_{t+1}}^{-1}}{t+1} + \sqrt{\frac{\tilde{p}_{\zeta_{t+1}}^{-2} \tilde{\beta}_{t+1} \tilde{\kappa}_{t+1} M_f \max\{\hat{C}_1, \dots, \hat{C}_{M_f}\}}{t+1}} \\
 & = \frac{2M_f L_\sigma \tilde{\beta}_{t+1}^{1/2} (1 + \tilde{p}_{\zeta_{t+1}}^{-1})}{t+1} \sum_{i=0}^t \zeta_{i+1} + \frac{16J \log \frac{8J}{\delta} \tilde{\beta}_{t+1}^{1/2} \tilde{p}_{\zeta_{t+1}}^{-1}}{t+1} + \sqrt{\frac{\hat{C} \tilde{p}_{\zeta_{t+1}}^{-2} \tilde{\beta}_{t+1} \tilde{\kappa}_{t+1}}{t+1}} \\
 & \equiv \hat{s}_t.
 \end{aligned}$$

Therefore, we obtain

$$a_i^{(\mathcal{X})}(\mathbf{x}_{i+1}) \leq q(\hat{s}_t).$$

Thus, for some $T \geq 0$ satisfying $q(\hat{s}_T) \leq \epsilon$, there exists $\hat{T} \leq T$ such that $a_{\hat{T}}^{(\mathcal{X})}(\mathbf{x}_{\hat{T}+1}) \leq q(\hat{s}_T) \leq \epsilon$. Noting that $0 \leq \hat{T} \leq T$, the algorithm terminates after at most T iterations.

B.10 Proof of Theorem A.7

The proof of Theorem A.7 is same as in the proof of Theorem A.4.

B.11 Proof of Theorem A.8

From the definition of $a_t^{(\mathcal{X})}(\mathbf{x})$ and \mathbf{x}_{t+1} , noting that $-\epsilon_{\text{PF}} + \mathbf{LCB}_t(\mathbf{x}_{t+1}) \in \text{Dom}(\mathbf{LCB}_t(\hat{\Pi}_t))$ we get

$$\begin{aligned}
 a_t^{(\mathcal{X})}(\mathbf{x}_{t+1}) & \leq \|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - (\mathbf{LCB}_t(\mathbf{x}_{t+1}) - \epsilon_{\text{PF}})\|_\infty \leq \epsilon_{\text{PF}} + \|\mathbf{UCB}_t(\mathbf{x}_{t+1}) - \mathbf{LCB}_t(\mathbf{x}_{t+1})\|_\infty \\
 & \leq \epsilon_{\text{PF}} + q \left(\max_{\mathbf{w} \in \Omega_{\mathbf{x}_{t+1}}} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}) \right) \\
 & = \epsilon_{\text{PF}} + q \left(\sum_{m=1}^{M_f} 2\tilde{\beta}_{m,t+1}^{1/2} \tilde{\sigma}_t^{(m)}(\mathbf{x}_{t+1}, \mathbf{w}_{i+1}^*) \right).
 \end{aligned}$$

Thus, by letting $\hat{t} = \text{argmin}_{0 \leq i \leq t} \sum_{m=1}^{M_f} 2\tilde{\beta}_{m,i+1}^{1/2} \tilde{\sigma}_i^{(m)}(\mathbf{x}_{i+1}, \mathbf{w}_{i+1}^*)$, using the same argument as in the proof of Theorem A.6, we have the desired result.

C EXPERIMENTAL DETAILS AND ADDITIONAL EXPERIMENTS

In this section, we give experimental details and additional experiments. All experiments were performed using R software version 3.6.3. For all experiments except for additional experiments, we set the tradeoff parameter $\beta_{m,t}^{1/2}$ to 3.

C.1 Details of Synthetic Function Experiments without Input Uncertainty

In the synthetic function experiments without IU, the input space \mathcal{X} was a set of grid points divided into $[-5, 5] \times [-5, 5]$ equally spaced at 50×50 . For black-box functions, we used Booth, Matyas, Himmelblau's and McCormic benchmark functions. We standardized these functions and further multiplied by minus one. The functional forms we actually used in our experiments are given as follows:

- Booth function:

$$f(x_1, x_2) = \frac{-(x_1 + 2x_2 - 7)^2 - (2x_1 + x_2 - 5)^2 + 157.35}{\sqrt{28896.11}}.$$

- Matyas function:

$$f(x_1, x_2) = \frac{-0.26(x_1^2 + x_2^2) + 0.48x_1x_2 + 4.3342}{\sqrt{23.52052}}.$$

- Himmelblau's function:

$$f(x_1, x_2) = \frac{-(x_1^2 + x_2 - 11)^2 + (x_1 + x_2^2 - 7)^2 + 136.71}{\sqrt{12503.63}}.$$

- McCormic function:

$$f(x_1, x_2) = \frac{-\sin(x_1 + x_2) - (x_1 - x_2)^2 + 1.5x_1 - 2.5x_2 - 117.67}{\sqrt{460.573}}.$$

We performed the following two cases: (i) Two-objective Pareto optimization problem using first two benchmark functions, (ii) four-objective Pareto optimization problem using all benchmark functions. For each black-box function, we used the independent GP model $\mathcal{GP}(0, k(\mathbf{x}, \mathbf{x}'))$, where the kernel function $k(\mathbf{x}, \mathbf{x}')$ is given by

$$k(\mathbf{x}, \mathbf{x}') = 2 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2}\right).$$

We used the zero-mean independent Gaussian noise with variance 10^{-6} for all black-box functions. As evaluation indicators, we used the simple Pareto hypervolume (PHV) regret, which is a commonly used indicator in the context of MOBOs, and inference discrepancy. Let $\mathcal{X}_t = \{\mathbf{x}_1, \dots, \mathbf{x}_t\}$ and $\mathcal{Y}_t = \{\mathbf{y}_1, \dots, \mathbf{y}_t\}$ be the set of input variables and observed values, respectively. Also let \mathbf{r} be a reference point of a multi-objective black-box function $\mathbf{f}(\mathbf{x}) = (f^{(1)}(\mathbf{x}), \dots, f^{(m)}(\mathbf{x}))$. Then, the simple PHV that we used in the experiments is given by

$$\text{Vol}(\mathbf{f}(\mathcal{X}); \mathbf{r}) - \text{Vol}(\mathbf{f}(\mathcal{X}_t); \mathbf{r}),$$

where $\mathbf{f}(A) \equiv \{\mathbf{f}(\mathbf{a}) \mid \mathbf{a} \in A\}$ and $\text{Vol}(\mathbf{f}(A); \mathbf{r})$ is the Lebesgue measure for $\{\mathbf{b} \mid \mathbf{r} \leq \mathbf{b} \text{ and } \mathbf{b} \leq \mathbf{f}(\mathbf{a}), \mathbf{a} \in A\}$. For a multi-objective black-box function $\mathbf{f}(\mathbf{x})$, we used $r_j = \min_{\mathbf{x} \in \mathcal{X}} f^{(j)}(\mathbf{x})$ as the j -th reference point. As AFs, we considered the random sampling (Random), uncertainty sampling (US), EHVI (Emmerich and Klinkenberg, 2008), EMmI (Svenson and Santner, 2010), ePAL (Zuluaga et al., 2016), ParEGO (Knowles, 2006), PFES (Suzuki et al., 2020) and proposed AF (Proposed). The next evaluation point was selected at random in Random. We used the AF $a_t(\mathbf{x}) = \sigma_t^{(1)2}(\mathbf{x}) + \dots + \sigma_t^{(m)2}(\mathbf{x})$ for US. In EHVI, we calculated sampling-based expected hypervolume improvement given by

$$\frac{1}{S} \sum_{s=1}^S \{\text{Vol}(\mathcal{Y}_t \cup \{\mathbf{y}_s(\mathbf{x})\}; \mathbf{r}) - \text{Vol}(\mathcal{Y}_t; \mathbf{r})\},$$

where $\mathbf{y}_s(\mathbf{x})$ is generated from the posterior distribution of $\mathbf{f}(\mathbf{x})$ and we set $S = 20$. In EMmI, we calculated sampling-based expected maximin distance improvement given by

$$\frac{1}{S} \sum_{s=1}^S \text{dist}(\mathbf{y}_s(\mathbf{x}), \text{Dom}(\mathcal{Y}_t)),$$

where S and $\mathbf{y}_s(\mathbf{x})$ are the same definition in EHVI. In ePAL, we performed the ϵ -PAL algorithm with parameter $\epsilon = (\epsilon_1, \dots, \epsilon_m) = (0, \dots, 0)$. In ParEGO, for each iteration t , we uniformly generated the vector of coefficients $\boldsymbol{\lambda}_t = (\lambda_t^{(1)}, \dots, \lambda_t^{(m)})^\top$ with $0 \leq \lambda_t^{(i)} \leq 1$ and $\sum_{i=1}^m \lambda_t^{(i)} = 1$, and calculated the scalarization $\tilde{y}_{t, \tilde{t}} = 0.05 \boldsymbol{\lambda}_t^\top \mathbf{y}_t + \max_{1 \leq i \leq m} \lambda_t^{(i)} y_t^{(i)}$ for all $\tilde{t} \leq t$. We constructed the GP model for $\tilde{y}_{t,1}, \dots, \tilde{y}_{t,t}$ using \mathcal{X}_t , where the kernel function was used the same kernel for $\mathbf{f}(\mathbf{x})$ but the noise variance was set to 10^{-8} . We calculated the expected improvement (EI) (Moćkus, 1975) and the next point was selected by maximizing EI. In PFES, we used the random feature map (Rahimi and Recht, 2007) to obtain posterior sample path. We generated a 500-dimensional random feature vector before BO, and used it for all iterations. The posterior sample path was generated 10

times for each iteration, and we calculated the PFES AF. In the four-objective Pareto optimization setting, the maximum number of Pareto optimal input points defined based on the sample path was restricted to 50 due to the computational cost. We also compared the commonly used evolutionary computation-based method NSGA-II (Deb et al., 2002). The NSGA-II method was performed using `nsga2R` version 1.1 in R. In `nsga2R`, we set the tournament size, crossover probability, crossover distribution index, mutation probability and mutation distribution index to 2, 0.9, 20, 0.1 and 3, respectively. We considered the population size p to 5, 10, 15, 20, 30, 50, 100, 150 and 300. For each p , we set the number of generations to $300/p$. In NSGA-II, we used $\hat{\Pi}_t$ as the set of input variables reported by the algorithm. Under this setup, one initial point was taken at random and the algorithm was run until the number of iterations reached 300. This simulation repeated 100 times and the average simple PHV regret and inference discrepancy at each iteration were calculated. In NSGA-II, only results with the highest average performance at the end of the 300 iterations are shown ($p = 30, 150$ in the two and four-objective settings, respectively).

C.2 Details of Synthetic Function Experiments with Input Uncertainty

Here, the input space $\mathcal{X} \times \Omega$ was a compact subset. For $\mathcal{X} \times \Omega$, we considered infinite and finite set settings. We set $\mathcal{X} \times \Omega = [0.25, 0.75]^2 \times [-0.25, 0.25]^2$ in the infinite set setting. In the finite setting, $\mathcal{X} \times \Omega$ was a set of grid points divided into $[-1, 1]^3 \times [-1, 1]^3$ equally spaced at $7^3 \times 7^3 = 117649$.

ZDT1 Function The black-box function in the infinite setting was used the ZDT1 benchmark function $\mathbf{ZDT1}(\mathbf{a}) \in \mathbb{R}^2$ with a two-dimensional input \mathbf{a} , and the environmental variable \mathbf{w} was used as the input noise for \mathbf{x} . We standardized the ZDT1 function and further multiplied by minus one. The functional form we actually used is given as follows:

$$\begin{aligned} g^{(1)}(a_1, a_2) &= a_1, \\ h(a_1, a_2) &= 1 + 9a_2, \\ g^{(2)}(a_1, a_2) &= h(a_1, a_2) - \sqrt{g^{(1)}(a_1, a_2)h(a_1, a_2)}, \\ \mathbf{ZDT1}(\mathbf{a}) &= (f^{(1)}(a_1, a_2), f^{(2)}(a_1, a_2)) = \left(-\frac{g^{(1)}(a_1, a_2) - 0.5}{\sqrt{0.042}}, -\frac{g^{(2)}(a_1, a_2) - 3.9085}{\sqrt{2.5615}} \right). \end{aligned}$$

Thus, our considered black-box function was defined by $\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})$. We assumed \mathbf{w} was the uniform distribution on Ω and used the Bayes risk $\mathbb{E}[\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})]$. We used the independent zero-mean Gaussian noise distribution with variance 10^{-6} for $f^{(i)}(a_1, a_2)$. We constructed the independent GP model $\mathcal{GP}(0, k(\boldsymbol{\theta}, \boldsymbol{\theta}'))$ for $f^{(i)}$, where $\boldsymbol{\theta} = (x_1, x_2, w_1, w_2)$ and

$$k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \exp\left(-\frac{\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2}{0.2}\right).$$

In order to calculate the true PF Z^* , we performed `nsga2R` with population size 500 and the number of generations is 100. As comparison methods, we considered the Bayesian quadrature-based method (BQ) (Qing et al., 2023) and surrogate-assisted bounding box approach (SABBa) (Rivier and Congedo, 2022). Furthermore, four naive methods, Naive-random, Naive-US, Naive-EMmI and Naive-ePAL, were used for comparison. In BQ, Bayes risk $\mathbb{E}[\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})]$ was modeled by the Bayesian quadrature, and its posterior distribution is again a GP. In this experiment, we can compute the exact posterior mean and variance, and we used them. Let $\boldsymbol{\mu}_t^{(\text{BQ})}(\mathbf{x})$ be a posterior mean for $\mathbb{E}[\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})]$. Then, we used $\hat{\Pi}_t$ to the set of Pareto optimal inputs calculated by $\boldsymbol{\mu}_t^{(\text{BQ})}(\mathbf{x}_1), \dots, \boldsymbol{\mu}_t^{(\text{BQ})}(\mathbf{x}_t)$. The AF for \mathbf{x} , we used sampling-based approximation with sample size 20. In Proposed, we computed the sample average for $\boldsymbol{\mu}_t^{(m)}(\mathbf{x}, \mathbf{w}) - 3\sigma_t^{(m)}(\mathbf{x}, \mathbf{w})$ and $\boldsymbol{\mu}_t^{(m)}(\mathbf{x}, \mathbf{w}) + 3\sigma_t^{(m)}(\mathbf{x}, \mathbf{w})$ by generating only two sample \mathbf{w}_1 and \mathbf{w}_2 , and used them to $\text{lcb}_t^{(m)}(\mathbf{x})$ and $\text{ucb}_t^{(m)}(\mathbf{x})$. In order to calculate $\hat{\Pi}_t$, we used `nsga2R` with population size 50 and the number of generations is 50. In SABBa, we set the number of new design set \mathcal{X}_{new} to be read to 10. The elements in \mathcal{X}_{new} were selected uniformly at random. We omitted the first approximation and then set $N_{first} = 0$. The number of initial design set was set to 1, and for each iteration. Similarly, the number of function evaluation was also set to 1. In the GP model for $\mathbb{E}[\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})]$, we used

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{0.1}\right).$$

In the AF for \mathbf{x} , we calculated the sampling-based AF calculation with sample size 20. For the threshold \mathbf{s}_1 and \mathbf{s}_2 , we set $\mathbf{s}_1 = \mathbf{s}_2 = (h_1 r_t, h_2 r_t)$, where $h_1 = \max_{\mathbf{x} \in \mathcal{X}} F^{(1)}(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} F^{(1)}(\mathbf{x})$ and $h_2 = \max_{\mathbf{x} \in \mathcal{X}} F^{(2)}(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} F^{(2)}(\mathbf{x})$. Here, $F^{(i)}(\mathbf{x})$ is the i -th element of $\mathbb{E}[\mathbf{ZDT1}(\mathbf{x} + \mathbf{w})]$. Furthermore, the initial value of r_t was set to 0.5 and multiplied by 0.9 each time a new \mathcal{X}_{new} was read, and if $r_t < 0.001$, then we fixed $r_t = 0.001$. The $\hat{\Pi}_t$ was set to the Pareto-optimal points defined based on $\rho_{SA}(\mathbf{x})$ and $\tilde{\rho}(\mathbf{x})$ (see, Rivier and Congedo (2022) for details) with respect to the inputs read so far. In the naive methods, \mathbf{w} was generated five times from the same \mathbf{x} in one iteration t , and the sample mean of the black-box function values were calculated. By using \mathbf{x} and these values, the experiments in naive four methods were performed as a usual MOBO. We used the following kernel function:

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{0.1}\right).$$

The same calculation (approximation) method as in the without IU setting was used for calculating AFs. For all methods, the maximization of AFs was performed using *optim* function with the L-BFGS-B method in R.

6D-Rosenbrock Function The black-box function in the finite setting was used the six-dimensional Rosenbrock function $f(w_1, w_2, x_1, x_2, x_3, w_3) \in \mathbb{R}$. The functional form that we used is given as follows:

$$f(a_1, a_2, a_3, a_4, a_5, a_6) = \frac{273.45 - \sum_{i=1}^5 \{100(a_{i+1} - a_i^2) + (1 - a_i)^2\}}{\sqrt{28153.22}}.$$

We assume that \mathbf{w} was a discretized normal distribution on $\Omega = \Omega_1 \times \Omega_2 \times \Omega_3$. For each $\mathbf{w} \in \Omega_i$, the probability math function of \mathbf{w} is given by

$$p(\mathbf{w}) = \frac{\phi(\mathbf{w})}{\sum_{\hat{\mathbf{w}} \in \Omega_i} \phi(\hat{\mathbf{w}})},$$

where $\phi(x)$ is the probability density function of standard normal distribution. As risk measures, we used the expectation $\mathbb{E}[f(w_1, w_2, x_1, x_2, x_3, w_3)]$ and negative standard deviation $-\sqrt{\mathbb{V}[f(w_1, w_2, x_1, x_2, x_3, w_3)]}$. As comparison methods, we considered the Mean-variance-based method (MVA) (Iwazaki et al., 2021b), SABBa, Naive-random, Naive-US, Naive-EMMI and Naive-ePAL. We used the independent zero-mean Gaussian noise distribution with variance 10^{-6} for $f(w_1, w_2, x_1, x_2, x_3, w_3)$. We constructed the GP model $\mathcal{GP}(0, k(\boldsymbol{\theta}, \boldsymbol{\theta}'))$ for f , where $\boldsymbol{\theta} = (x_1, x_2, x_3, w_1, w_2, w_3)$ and

$$k(\boldsymbol{\theta}, \boldsymbol{\theta}') = \exp\left(-\frac{\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2}{4}\right).$$

This experiment is the setting that the number of black-box functions and risk measures are different. Thus, in Proposed, the algorithm was performed using Algorithm 2. In SABBa, we set the number of new design set \mathcal{X}_{new} to be read to 10. The elements in \mathcal{X}_{new} were selected uniformly at random. We omitted the first approximation and then set $N_{first} = 0$. The number of initial design set was set to 1, and for each iteration. Similarly, the number of function evaluation was also set to 1. In the GP model for Bayes risk and negative standard deviation, we used the following kernel:

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{4}\right).$$

In the AF for \mathbf{x} , we calculated the sampling-based AF calculation with sample size 20. For the threshold \mathbf{s}_1 and \mathbf{s}_2 , we set $\mathbf{s}_1 = \mathbf{s}_2 = (h_1 r_t, h_2 r_t)$, where $h_1 = \max_{\mathbf{x} \in \mathcal{X}} F^{(1)}(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} F^{(1)}(\mathbf{x})$ and $h_2 = \max_{\mathbf{x} \in \mathcal{X}} F^{(2)}(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} F^{(2)}(\mathbf{x})$. Here, $F^{(1)}(\mathbf{x})$ and $F^{(2)}(\mathbf{x})$ are Bayes risk and negative standard deviation, respectively. Furthermore, the initial value of r_t was set to 2 and multiplied by 0.9 each time a new \mathcal{X}_{new} was read, and if $r_t < 0.001$, then we fixed $r_t = 0.001$. The $\hat{\Pi}_t$ was set to the Pareto-optimal points defined based on $\rho_{SA}(\mathbf{x})$ and $\tilde{\rho}(\mathbf{x})$ with respect to the inputs read so far. In the naive methods, \mathbf{w} was generated five times from the same \mathbf{x} in one iteration t , and the sample mean and the negative square root of the sample variance of the black-box function values were calculated. By using \mathbf{x} and these values, the experiments in naive four methods were performed as a usual MOBO. We used the following kernel function:

$$k(\mathbf{x}, \mathbf{x}') = \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{2}\right).$$

The same calculation (approximation) method as in the without IU setting was used for calculating AFs.

C.3 Details of Real-world Simulation Model

We applied the proposed method to docking simulations for real-world chemical compounds. As a dataset for compounds, we used the software suite *Schrödinger* (Schrödinger LLC, 2021). Given a set of compounds, we applied the software “QikProp” in the suite, a software to compute various chemical properties, for explanatory variables. We also applied the software “Glide” in the suite, a software for calculating docking scores. We took the black-box function as the original docking score plus 5 and then multiplied by -1. When performing docking simulations, it is necessary to specify both the protein of interest and the specific site on the protein where compounds are expected to dock. We used the protein “KAT1”, whose structure is available at <https://pdbe.org/mine/summary/6v1x>, and the 16th and 18th sites computed by the software “SiteMap” in the suite. Each chemical compound C_i may have an isomer S_{ij} , and in this experiment the maximum number of isomers was limited to 10. For each i , we computed a 51-dimensional isomer-independent vector of explanatory variables \mathbf{x}_i and a 51-dimensional environment variable \mathbf{w}_{ij} that can vary with isomers, using physicochemical features of (C_i, S_{ij}) computed using QikProp. Specifically, the 51-dimensional physicochemical features of (C_i, S_{ij}) calculated by QikProp were used as \mathbf{w}_{ij} . In addition, we defined \mathbf{x}_i as $\mathbf{x}_i = \frac{1}{N_i} \sum_{j=1}^{N_i} \mathbf{w}_{ij}$. Thus, the black-box functions, the docking scores in the 16th and 18th sites, can be expressed as $f^{(1)}(\mathbf{x}_i, \mathbf{w}_{ij})$ and $f^{(2)}(\mathbf{x}_i, \mathbf{w}_{ij})$, respectively. As risk measures for C_i , we considered the following measures:

Worst-case (WC): $F^{(m)}(\mathbf{x}_i) = \min_{j \in [N_i]} f^{(m)}(\mathbf{x}_i, \mathbf{w}_{ij})$.

Worst-case Bayes risk (WCBR): Define the Bayes risk under the worst-case candidate distribution, that is,

$$F^{(m)}(\mathbf{x}_i) = \min_{\boldsymbol{\alpha}_i \in \mathcal{A}_i} \sum_{j=1}^{N_i} \alpha_{ij} f^{(m)}(\mathbf{x}_i, \mathbf{w}_{ij}).$$

The \mathcal{A}_i is the set of $\boldsymbol{\alpha}_i \in \mathbb{R}^{N_i}$ satisfying

$$0 \leq \alpha_{ij} \leq 1, \sum_{j=1}^{N_i} \alpha_{ij} = 1, \|\boldsymbol{\alpha}_i - \tilde{\boldsymbol{\alpha}}_i\|_1 \leq 0.25,$$

where $\tilde{\alpha}_{ij} = N_i^{-1}$. The total number of compounds was 429, and the total number of data including isomers was 920. We compared Proposed, SABba and naive four methods. In this experiment, we used the independent GP model for $f^{(m)}$, where the kernel function is given by

$$k(\boldsymbol{\theta}, \boldsymbol{\theta}') = 25 \exp\left(-\frac{\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2}{l}\right).$$

The length scale parameter was computed using the median heuristic $l = 0.5 \text{Median}\{\|\boldsymbol{\theta}_i - \boldsymbol{\theta}_j\|^2 \mid 1 \leq i < j \leq 920\}$. In this experiment, there was no observation noise. Nevertheless, we added $10^{-3} \mathbf{I}_t$ to the kernel matrix \mathbf{K}_t to stabilize the inverse matrix calculation. In SABba, we set the number of new design set \mathcal{X}_{new} to be read to 10. The elements in \mathcal{X}_{new} were selected uniformly at random. We omitted the first approximation and then set $N_{first} = 0$. The number of initial design set was set to 1, and for each iteration. Similarly, the number of function evaluation was also set to 1. In the GP model for risk measures, we used the following kernel:

$$k(\mathbf{x}, \mathbf{x}') = 25 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{l}\right),$$

where the length scale parameter was computed using the median heuristic $l = 0.5 \text{Median}\{\|\mathbf{x}_i - \mathbf{x}_j\|^2 \mid 1 \leq i < j \leq 429\}$. In the AF for \mathbf{x} , we calculated the sampling-based AF calculation with sample size 20. For the threshold \mathbf{s}_1 and \mathbf{s}_2 , we set $\mathbf{s}_1 = \mathbf{s}_2 = (h_1 r_t, h_2 r_t)$, where $h_1 = \max_{\mathbf{x} \in \mathcal{X}} F^{(1)}(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} F^{(1)}(\mathbf{x})$ and $h_2 = \max_{\mathbf{x} \in \mathcal{X}} F^{(2)}(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} F^{(2)}(\mathbf{x})$. Furthermore, the initial value of r_t was set to 0.5 and multiplied by 0.9 each time a new \mathcal{X}_{new} was read, and if $r_t < 0.01$, then we fixed $r_t = 0.01$. We regarded this as a high accurate setting. As a low accurate setting, we considered that the initial value of r_t was set to 2 and multiplied by 0.99 each time a new \mathcal{X}_{new} was read, and if $r_t < 0.01$, then we fixed $r_t = 0.01$. In the original SABba method, Rivier and Congedo (2022) does not provide the worst-case Bayes risk setting. Hence, we modified

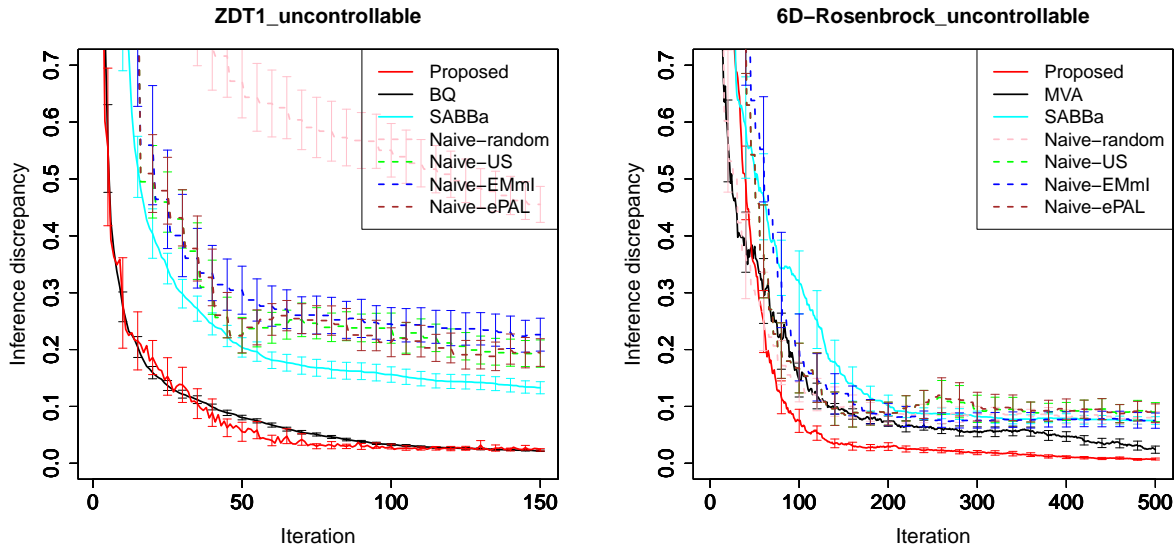


Figure 3: Comparison with MOBO methods. Solid (and dashed) lines are averages of the inference discrepancy for each iteration in 100 trials. Each error bar length represents the six times the standard error. The left and right columns respectively represent the ZDT1 and six-dimensional Rosenbrock setups under the uncontrollable setting.

the first formula of Equation (6) in Rivier and Congedo (2022) to $\inf_{\xi \in \mathcal{A}} \mathbb{E}_{\xi}[\bar{\varepsilon}_{q_{\alpha}}(\xi)]$. The $\hat{\Pi}_t$ was set to the Pareto-optimal points defined based on $\rho_{SA}(\mathbf{x})$ and $\hat{\rho}(\mathbf{x})$ with respect to the inputs read so far. In the naive four methods, we calculated docking scores for all isomers in the compound C_i at iteration t and determined the exact risk values. By using \mathbf{x} and these values, the experiments in naive four methods were performed as a usual MOBO. We used the following kernel function:

$$k(\mathbf{x}, \mathbf{x}') = 25 \exp\left(-\frac{\|\mathbf{x} - \mathbf{x}'\|_2^2}{l}\right),$$

where the length scale parameter was computed using the median heuristic $l = 0.5 \text{Median}\{\|\mathbf{x}_i - \mathbf{x}_j\|^2 \mid 1 \leq i < j \leq 429\}$. The same calculation (approximation) method as in the without IU setting was used for calculating AFs.

C.4 Additional Experiments

Uncontrollable Setting for Synthetic Experiments Here, we give the results of synthetic experiments (ZDT1 and 6D-Rosenbrock) under the uncontrollable setting. We performed the same experiments except for the selection of \mathbf{w} . Figure 3 shows the similar results as in the simulator setting.

Docking Simulation Experiments Using Bayes Risk In the docking simulation experiments, we also considered Bayes risk (BR) $F^{(m)}(\mathbf{x}_i) = \frac{1}{N_i} f^{(m)}(\mathbf{x}_i, \mathbf{w}_{ij})$. In this experiment, we also considered the BQ method. In BQ, $f^{(m)}(\mathbf{x}_i, \mathbf{w}_{ij})$ was modeled in the same way as in Proposed. The AF for \mathbf{x} was calculated using sampling-based approximation with sample size 20. Figure 4 shows the similar results as in the WC and WCBR settings. Also in the BR setting, only the proposed method correctly identifies the true PF at the end of 500 iterations at all 920 different initial points. Specifically, after 481 iterations, the true PF is identified for all 920 different initial points.

Hyperparameter Sensitivity In this section, we confirm the sensitivity for hyperparameters. In this experiment, the input space $\mathcal{X} \times \Omega \subset \mathbb{R}^2 \times \mathbb{R}$ was a set of grid points divided into $[-2, 2]^3$ equally spaced at $16 \times 16 \times 16 = 4096$. The true black-box functions $f^{(1)}(x_1, x_2, w_1)$ and $f^{(2)}(x_1, x_2, w_1)$ were defined as the

independent sample path from the GP $\mathcal{GP}(0, k^*(\cdot, \cdot))$, where $k^*(\cdot, \cdot)$ is given by

$$k^*((x_1, x_2, w_1), (x'_1, x'_2, w'_1)) = \exp\left(-\frac{(x_1 - x'_1)^2 + (x_2 - x'_2)^2 + (w_1 - w'_1)^2}{1}\right).$$

We used the zero-mean independent Gaussian noise with variance 10^{-6} . As the distribution of $w \in \Omega$, we used the discretized normal distribution $p(w)$ given by

$$p(w) = \frac{\phi(w)}{\sum_{w' \in \Omega} \phi(w')}.$$

We considered Bayes risk in this experiment. As the GP surrogate model, we used independent GP model for $f^{(1)}$ and $f^{(2)}$, and the kernel function that we used is given by

$$k((x_1, x_2, w_1), (x'_1, x'_2, w'_1)) = \sigma^2 \exp\left(-\frac{(x_1 - x'_1)^2 + (x_2 - x'_2)^2 + (w_1 - w'_1)^2}{L}\right).$$

We considered six cases for (L, σ^2) ,

$$(L, \sigma^2) = (1, 1), (L, \sigma^2) = (1, 2), (L, \sigma^2) = (0.5, 1), (L, \sigma^2) = (1, 0.5), (L, \sigma^2) = (2, 1), (L, \sigma^2) = (1, 0.1).$$

Similarly, we considered seven cases for $\beta_t^{1/2}$,

$$\begin{aligned} \beta_t^{1/2} &= 1, \beta_t^{1/2} = 2, \beta_t^{1/2} = 3, \beta_t^{1/2} = 4, \beta_t^{1/2} = 5, \\ \beta_t^{1/2} &= \sqrt{2 \log(2 \times 4096/2) + r_t}, \beta_t^{1/2} = \sqrt{2 \log(2 \times 4096\pi^2 t^2 / (6 \times 0.1))}, \end{aligned}$$

where r_t is a realized value from the exponential distribution with mean 0.5. The last two definitions of $\beta_t^{1/2}$ are proposed by Takeno et al. (2023) and Srinivas et al. (2010), respectively. We regarded them as Sampled and Theoretical values, respectively. Under this setup, one initial point was taken at random and the algorithm was run until the number of iterations reached 500. This simulation repeated 100 times and the average inference discrepancy at each iteration was calculated. From the top of Fig. 5, it can be confirmed that $\beta_t^{1/2} = 2$ is sufficient if the correct kernel is used, and $\beta_t^{1/2} = 1$ is sufficient for the right two columns of the top row that the posterior variance is predicted larger. On the other hand, if the cases the posterior variance is predicted smaller, $\beta_t^{1/2} = 3$ is still insufficient in the case of $L = 1, \sigma^2 = 0.1$.

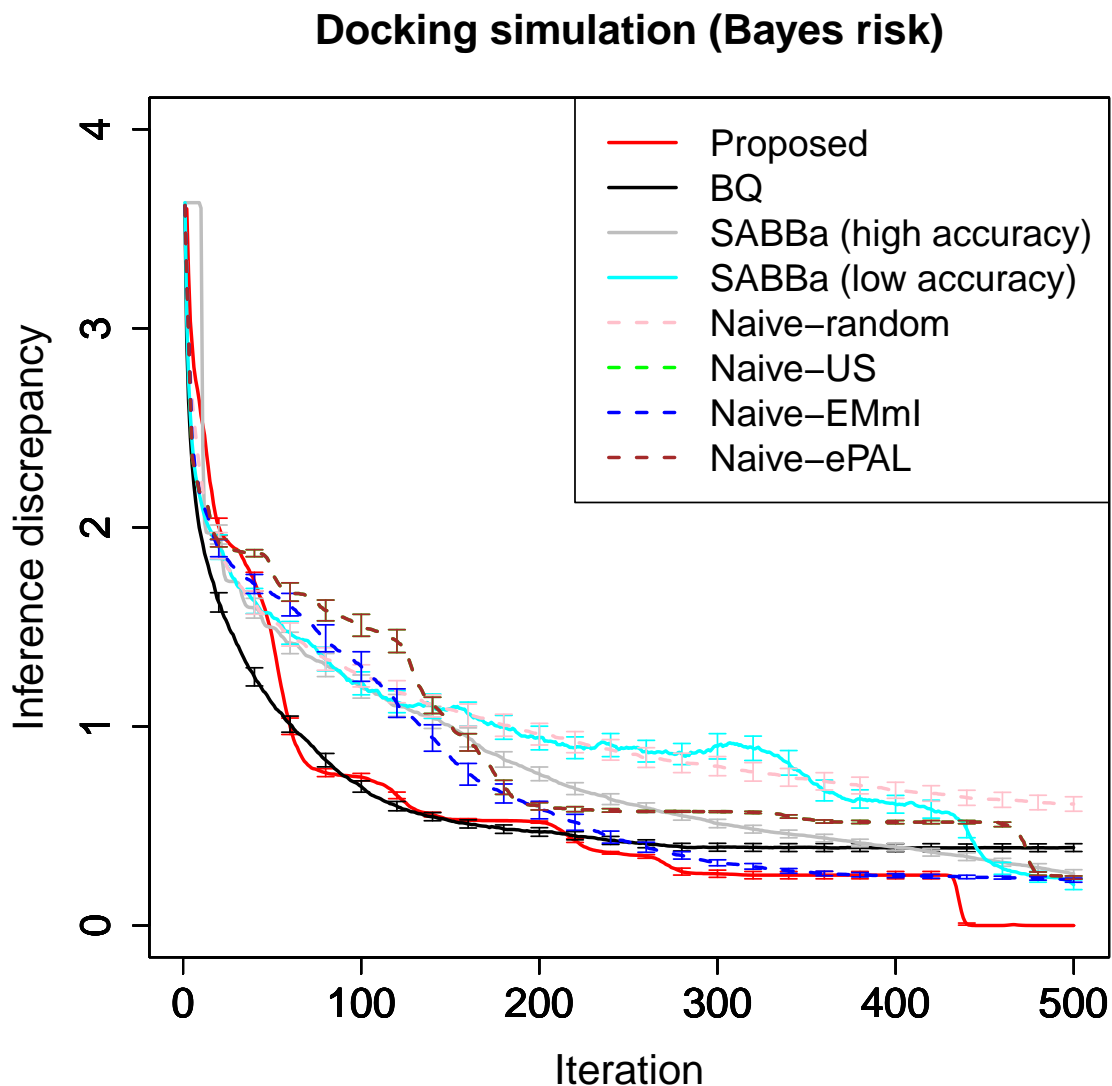


Figure 4: Comparison with MOBO methods. Solid (and dashed) lines are averages of the inference discrepancy of Bayes risk setting for each iteration in 920 or 429 trials. Each error bar length represents the six times the standard error.

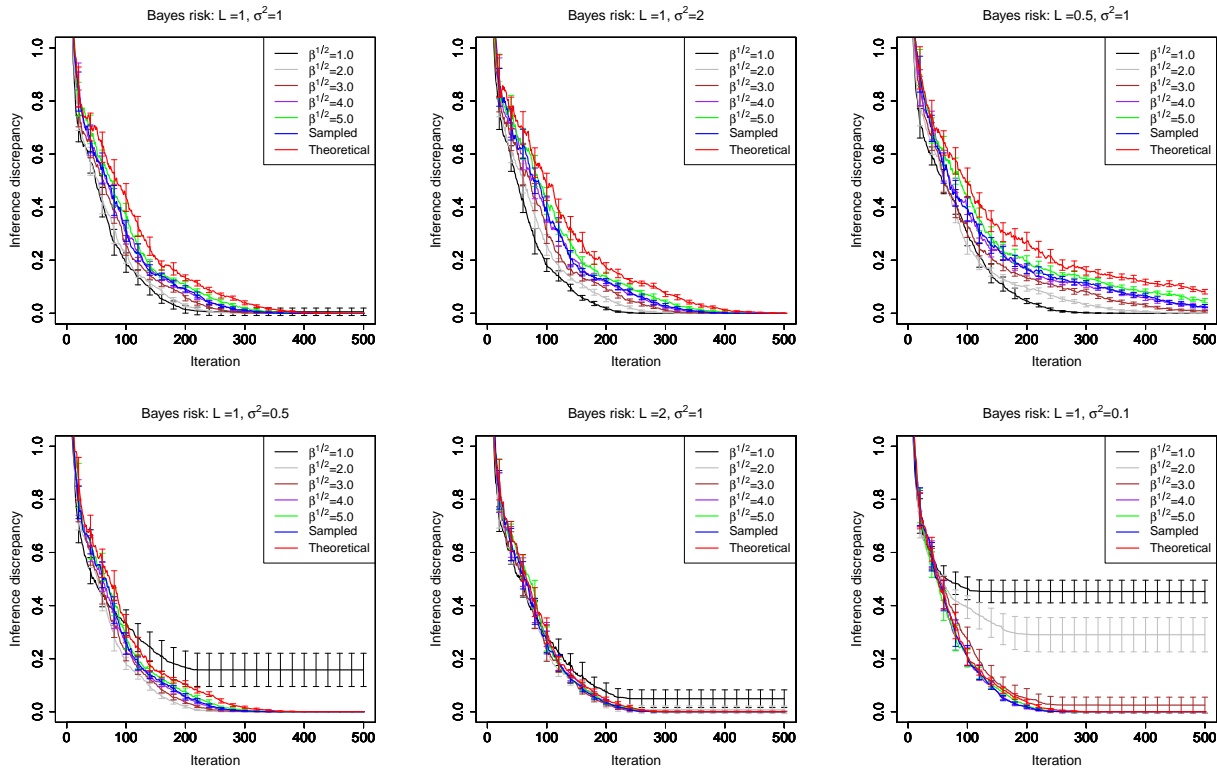


Figure 5: Comparison with different hyperparameters. Solid lines are averages of the inference discrepancy for each iteration in 100 trials. Each error bar length represents the six times the standard error. In the top row, the left column represents the case that the kernel of the surrogate model is equal to the true kernel. The right two columns represent the cases that the posterior variance is predicted larger. In the bottom row, the left, center and right columns represent the cases that the posterior variance is predicted smaller.