Randomized Kriging Believer For Parallel Bayesian Optimization With Regret Bounds

Shuhei Sugiura
RIKEN AIP
shuhei.sugiura@riken.jp

Ichiro Takeuchi
Nagoya University and RIKEN AIP
takeuchi.ichiro.n6@f.mail.nagoya-u.ac.jp

Shion Takeno
Nagoya University
takeno.shion.m6@f.mail.nagoya-u.ac.jp

Abstract

We consider the optimization problem of an expensive-to-evaluate black-box function, in which we can obtain noisy function values in parallel. For this problem, parallel Bayesian optimization (PBO) is a promising approach, which aims to optimize with fewer function evaluations by selecting a diverse input set for parallel evaluation. However, existing PBO methods suffer from poor practical performance or lack theoretical guarantees. In this study, we propose a PBO method, called randomized kriging believer (KB), based on a well-known KB heuristic and inheriting the advantages of the original KB: low computational complexity, a simple implementation, versatility across various BO methods, and applicability to asynchronous parallelization. Furthermore, we show that our randomized KB achieves Bayesian expected regret guarantees. We demonstrate the effectiveness of the proposed method through experiments, including those on real-data emulators.

1 Introduction

Bayesian optimization (BO) [25, 28] is a promising approach for optimization of expensive-to-evaluate black-box functions with a smaller number of function evaluations. For this purpose, BO sequentially queries the input that maximizes the acquisition function (AF) based on a Bayesian model. BO has been leveraged to a wide range of applications, such as AutoML [37], robotics [4], and materials informatics [48]. However, in many real-world applications, observations can be obtained in parallel. For example, if the objective function involves computer simulation and multiple computational resources are available, parallel execution of simulations is important to minimize wall-clock time for optimization. Naively applying vanilla BO methods to such problems can waste query budget, since the input points often cluster and provide redundant information.

Parallel BO (PBO) [36] aims to improve optimization efficiency by selecting diverse input points that effectively leverage parallel evaluations. Simple heuristics to extend BO methods to PBO by promoting diversity [2, 12, 13] are widely used for their advantages, such as simple implementation, low computational complexity, and applicability to asynchronous parallelization. A representative heuristic is kriging believer (KB) [12], which promotes diversity by sequentially imputing predictive values as fictitious observations at input points currently under evaluation. While the effectiveness of these heuristics has been shown empirically, they generally lack theoretical guarantees.

In contrast to heuristic approaches, several PBO methods with theoretical guarantees have been proposed, including parallel Thompson sampling (PTS) [22, 29] and batched upper confidence bound (BUCB) [10]. Under standard regularity assumptions, these methods have guarantees regarding regret

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

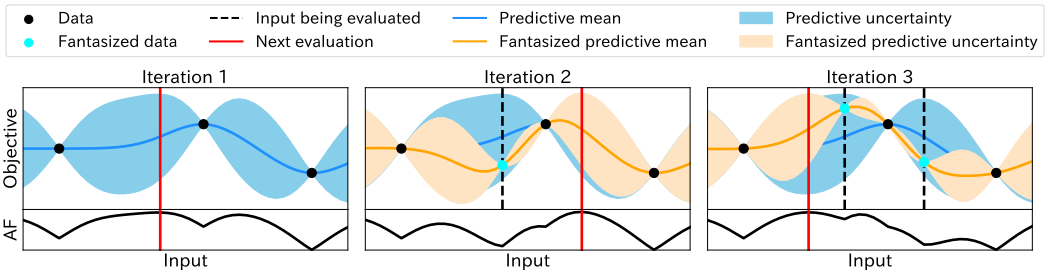


Figure 1: Schematic illustration of the proposed method for three consecutive iterations. For efficient parallel optimization, it is necessary to evaluate diverse input points while avoiding redundant evaluation. The left figure shows iteration 1, where no inputs are under evaluation, and thus, the next input is chosen as in standard BO. The middle and right figures show iterations 2 and 3, respectively. There are inputs currently being evaluated, indicated by dashed lines. In these cases, RKB generates fantasized data from the predictive distribution at the inputs under evaluation. Then, the next input is selected by maximizing an AF based on a model trained on both observed (real) and fantasized data. As a result, BO avoids redundant evaluation and ensures diversity among the evaluated inputs.

[38]. However, such methods typically suffer from poor practical performance, lack theoretical guarantees for practical tuning of the confidence width parameter, or involve complex implementations. This gap between strong theoretical guarantees and practical effectiveness motivates the development of PBO methods that are both theoretically principled and empirically competitive.

This paper proposes and analyzes a randomized variant of the KB heuristic, referred to as randomized KB (RKB), shown in Fig. 1. Our contributions are summarized as follows:

1. We propose a PBO method, RKB, that selects a diverse input set by conditioning on one random posterior realization for the ongoing evaluation points, inheriting the practical advantages of the original KB: ease of implementation, low computational complexity, applicability to asynchronous parallelization, and versatility across a wide range of BO methods.
2. We establish upper bounds on both Bayesian cumulative regret (BCR) and Bayesian simple regret (BSR) for RKB when combined with several BO methods that admit regret guarantees [38, 43, 45, 46, 47] in Theorems 4.1, 4.2, and 4.3, which especially provide the BSR upper bound independent of the number of parallel workers.
3. We demonstrate the effectiveness of the proposed method via a wide range of experiments, including emulators of real-world data.

1.1 Related work

Joint selection methods. Several studies have proposed PBO methods that select a set of inputs either by jointly optimizing a utility-based AF that explicitly quantifies a batch's utility, or by sampling from a joint distribution designed to encourage diversity. Representative examples include PBO methods based on expected improvement (EI) [6, 12, 27, 52], predictive entropy search (PES) [35], and knowledge gradient (KG) [56], all of which jointly optimize AFs over a batch of $O(Q)$ inputs, where Q is the number of parallel workers. Alternatively, determinantal point process (DPP)-based approaches, such as DPP sampling [23] and DPP-TS [29], generate input sets by sampling from a DPP distribution that explicitly incorporates diversity. These DPP-based methods admit regret guarantees [23, 29]. However, both approaches incur substantial computational costs that grow rapidly with Q, due to high-dimensional AF optimization or sampling from high-dimensional distributions using Markov chain Monte Carlo (MCMC) methods.

Greedy selection methods. Greedy PBO methods, including our RKB, have also been actively studied. Importantly, these greedy approaches have low computational complexity in Q, comparable to sequential BO, and naturally support asynchronous parallelization. BUCB [10], UCB with pure exploration (UCB-PE) [9], and DPP-based greedy approximations [23] have regret guarantees but often suffer from over-exploration. In contrast, KB [12], constant liar [12], simulation matching

[2], and local penalization (LP) [13] often show superior performance. Furthermore, Monte Carlo (MC)-estimation-based PBO for utility-based AFs, such as EI [3, 37, 52, 53] and max-value entropy search (MES) [40, 41, 42] has also been widely used. The MC estimation-based methods greedily select inputs by maximizing the AF conditioned on ongoing evaluations and averaging over multiple posterior samples. Thus, our RKB can be interpreted as a special case that uses a single MC sample, although our regret analysis applies only to this single-sample setting. The empirical effectiveness of these heuristics has been demonstrated in multiple studies above, but lacks theoretical guarantees.

Distributed selection methods. A third line of work is fully distributed PBO, especially for the case that even linear computational dependence on Q is prohibitive. Distributed PBO methods, including PTS [15, 22, 29, 50] and Boltzmann policies derived from vanilla BOs [11], select input sets in a fully distributed manner, such as independent posterior sampling, by which they avoid computational dependence on Q in terms of wall-clock time. However, by design, these methods do not incorporate other selected inputs, which can result in batches with insufficient diversity.

Regret analysis for PBO. Regarding PBO, several works have established regret guarantees. High-probability regret bounds have been derived for parallel UCB variants [9, 10, 23]. Desautels et al. [10] showed that the deterioration of regret with respect to the batch size Q can be avoided by performing uncertainty sampling $O(Q^c)$ times at an initial phase, where $c \geq 1$ is a constant. Regret bounds for PTS and DPP-TS have been established [8, 22, 29]. Notably, Nava et al. [29] derived a BSR bound that is independent of Q, for which we obtain the comparable upper bound for RKB. However, the above algorithms with regret guarantees suffer from practical disadvantages discussed already. Finally, the analysis of regret lower bounds in our problem setup is limited to vanilla BO for one-dimensional objectives [34].

2 Preliminaries

This section provides background knowledge.

2.1 Problem statement

We consider the optimization of a black-box and expensive-to-evaluate function $f : \mathcal{X} \to \mathbb{R}$:

$$\boldsymbol{x}^* = \operatorname*{argmax}_{\boldsymbol{x} \in \mathcal{X}} f(\boldsymbol{x}). \tag{1}$$

For this problem, we sequentially query $\boldsymbol{x}_t$ for all iterations $t \in \mathbb{N}$ and obtain an observation y_t, which can be contaminated, aiming for sample-efficient optimization. We consider the general case that up to $Q-1$ queries can be unobserved. This setting recovers vanilla BO with $Q = 1$ and includes synchronous and asynchronous parallelization with Q workers.

2.2 Gaussian process regression

We assume the following regularity assumption for the Gaussian process (GP) regression model [32]:

Assumption 2.1. *For $d \in \mathbb{N}$ and $r > 0$, let k be a kernel on $\mathcal{X} \subset [0, r]^d$ that satisfies $\max_{\boldsymbol{x} \in \mathcal{X}} k(\boldsymbol{x}, \boldsymbol{x}) \leq 1$. Then, function f follows $\mathcal{GP}(0, k)$, and its input-output data is defined as*

$$(\boldsymbol{x}_t, y_t) \in \mathcal{X} \times \mathbb{R}, \quad y_t = f(\boldsymbol{x}_t) + \varepsilon_t \qquad (t \in \mathbb{N}), \tag{2}$$

where $\varepsilon_t \sim \mathcal{N}\left(0, \sigma^2_{\mathrm{noise}}\right)$ with $\sigma_{\mathrm{noise}} > 0$.

For any index set $\mathcal{S} \subset \mathbb{N}$, we define data set $\mathcal{D}_{\mathcal{S}}$ as $\{(\boldsymbol{x}_i, y_i)\}_{i \in \mathcal{S}}$. For any $t \in \mathbb{N}$, we denote the set $\{1, \ldots, t\}$ as $[t]$ and the data set $\mathcal{D}_{[t-1]}$ as $\mathcal{D}_{t-1}$. Let $\mathcal{N}_{t-1} \subset [t-1]$ be the index set of data available at the beginning of the t-th iteration. Then, the set $\mathcal{N}_{t-1}$ satisfies $|\mathcal{N}_{t-1}| \geq t - Q$ and $\mathcal{N}_t \supset \mathcal{N}_{t-1}$. Under Assumption 2.1, the posterior distribution conditioned on $\mathcal{D}_{\mathcal{N}_t}$ is also a GP. The posterior covariance between $f(\boldsymbol{x})$ and $f(\boldsymbol{x}')$ is written as

$$k(\boldsymbol{x}, \boldsymbol{x}'; \mathcal{D}_{\mathcal{N}_t}) = k(\boldsymbol{x}, \boldsymbol{x}') - \boldsymbol{k}_{\mathcal{N}_t}(\boldsymbol{x})^\top \left(\boldsymbol{K}_{\mathcal{N}_t} + \sigma^2_{\mathrm{noise}} \boldsymbol{I}_{n_t}\right)^{-1} \boldsymbol{k}_{\mathcal{N}_t}(\boldsymbol{x}') \qquad (\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{X}), \tag{3}$$

where $I_{n_t} \in \mathbb{R}^{n_t \times n_t}$ with $n_t = |\mathcal{N}_t|$ denotes the identity matrix, $\boldsymbol{K}_{\mathcal{N}_t} \in \mathbb{R}^{n_t \times n_t}$ is a kernel matrix whose (i,j)-th element is $k\big(\boldsymbol{x}_{[\mathcal{N}_t]_i}, \boldsymbol{x}_{[\mathcal{N}_t]_j}\big)$, where $[\mathcal{N}_t]_i$ is the i-th smallest element of $\mathcal{N}_t$. Let $k(\boldsymbol{x}, \boldsymbol{x}'; \emptyset) = k(\boldsymbol{x}, \boldsymbol{x}')$. Then, the posterior mean and variance of $f(\boldsymbol{x})$ are given by

$$\mu(\boldsymbol{x}; \mathcal{D}_{\mathcal{N}_t}) = \boldsymbol{k}_{\mathcal{N}_t}(\boldsymbol{x})^\top \big(\boldsymbol{K}_{\mathcal{N}_t} + \sigma^2_{\mathrm{noise}} \boldsymbol{I}_{n_t}\big)^{-1} \boldsymbol{y}_{\mathcal{N}_t}, \quad \sigma^2(\boldsymbol{x}; \mathcal{D}_{\mathcal{N}_t}) = k(\boldsymbol{x}, \boldsymbol{x}; \mathcal{D}_{\mathcal{N}_t}) \quad (\boldsymbol{x} \in \mathcal{X}), \tag{4}$$

where $\boldsymbol{y}_{\mathcal{N}_t} = [y_i]^\top_{i \in \mathcal{N}_t} \in \mathbb{R}^{n_t}$, and $\boldsymbol{k}_{\mathcal{N}_t}(\boldsymbol{x}) = [k(\boldsymbol{x}, \boldsymbol{x}_i)]^\top_{i \in \mathcal{N}_t} \in \mathbb{R}^{n_t}$.

For continuous input domains, we assume the following regularity condition as in [22, 29, 38], which is satisfied by Gaussian kernels and Matérn-ν kernels with $\nu > 2$ [38]:

Assumption 2.2. *Assume that $\mathcal{X}$ is compact and convex, and $f \sim \mathcal{GP}(0, k)$ satisfies*

$$\exists a, b > 0, \forall L \geq 0, \forall j \in [d], \Pr\left(\sup_{\boldsymbol{x} \in \mathcal{X}} \left|\frac{\partial f(\boldsymbol{x})}{\partial x_j}\right| > L\right) \leq a \exp\left(-\frac{L^2}{b^2}\right), \tag{5}$$

where $[x_1 \ \cdots \ x_d]^\top = \boldsymbol{x}$.

2.3 Acquisition functions for BO

BO algorithm $\mathcal{A}$ selects an input to evaluate by maximizing the AF $\alpha : \mathcal{X} \to \mathbb{R}$. That is, $\mathcal{A}(\mathcal{D}_{t-1}) = \mathrm{argmax}_{\boldsymbol{x} \in \mathcal{X}} \alpha(\boldsymbol{x}; \mathcal{D}_{t-1})$. Here, we describe three AFs used in our numerical experiment.

The first is UCB [38] defined as

$$\alpha_{\mathrm{UCB}}(\boldsymbol{x}; \mathcal{D}_{t-1}) = \mu(\boldsymbol{x}; \mathcal{D}_{t-1}) + \beta_t^{\frac{1}{2}} \sigma(\boldsymbol{x}; \mathcal{D}_{t-1}) \quad (\boldsymbol{x} \in \mathcal{X}), \tag{6}$$

where $\beta_t > 0$ is the confidence width parameter. The second is EI [28] defined as

$$\alpha_{\mathrm{EI}}(\boldsymbol{x}; \mathcal{D}_{t-1}) = \sigma(\boldsymbol{x}; \mathcal{D}_{t-1}) \left(s\left(\boldsymbol{x}; \mathcal{D}_{t-1}\right) \Phi(s\left(\boldsymbol{x}; \mathcal{D}_{t-1}\right)) + \phi(s\left(\boldsymbol{x}; \mathcal{D}_{t-1}\right))\right) \quad (\boldsymbol{x} \in \mathcal{X}), \tag{7}$$

where Φ and ϕ denote the cumulative distribution function and probability density function of the standard normal distribution, respectively, and $s\left(\boldsymbol{x}; \mathcal{D}_{t-1}\right)$ is defined as

$$s\left(\boldsymbol{x}; \mathcal{D}_{t-1}\right) = \frac{\mu(\boldsymbol{x}; \mathcal{D}_{t-1}) - \max_{\boldsymbol{x}' \in \mathcal{X}} \mu(\boldsymbol{x}'; \mathcal{D}_{t-1})}{\sigma(\boldsymbol{x}; \mathcal{D}_{t-1})}. \tag{8}$$

The third is PI from the maximum of a sample path (PIMS) [45] defined as

$$\alpha_{\mathrm{PIMS}}(\boldsymbol{x}; \mathcal{D}_{t-1}) = 1 - \Phi\left(\frac{g_t^* - \mu(\boldsymbol{x}; \mathcal{D}_{t-1})}{\sigma(\boldsymbol{x}; \mathcal{D}_{t-1})}\right) \quad (\boldsymbol{x} \in \mathcal{X}), \tag{9}$$

where $g_t^* = \max_{\boldsymbol{x} \in \mathcal{X}} g_t(\boldsymbol{x})$ and $g_t \sim p(f \mid \mathcal{D}_{t-1})$. Hence, PIMS is random even if the data is fixed.

2.4 Kriging believer for PBO

Here, we describe KB [12], which serves as the basis of our approach. KB selects $\boldsymbol{x}_t$, incorporating the diversity of selected inputs, as

$$\boldsymbol{x}_t = \mathcal{A}\big(\mathcal{D}^{\mathrm{KB}}_{t-1}\big), \quad \mathcal{D}^{\mathrm{KB}}_{t-1} = \left\{\left(\boldsymbol{x}_i, y_i^{(t)}\right)\right\}_{i=1}^{t-1}, \quad y_i^{(t)} = \begin{cases} y_i & (i \in \mathcal{N}_{t-1}) \\ \mu\big(\boldsymbol{x}_i; \mathcal{D}_{\mathcal{N}_{t-1}}\big) & \big(i \in \overline{\mathcal{N}}_{t-1}\big) \end{cases}, \tag{10}$$

where $\mathcal{A}$ is an arbitrary vanilla BO algorithm, and $\overline{\mathcal{N}}_{t-1} = [t-1] \setminus \mathcal{N}_{t-1}$. As shown in Eq. (10), KB uses the posterior mean of f in place of real data that have not yet been obtained. Such conditioning on fantasized data prevents redundant evaluation.

2.5 Bayesian regret and maximum information gain

As a criterion of the performance of BO methods, we employ Bayesian regret [22, 33]. The BSR and BCR are defined as

$$\mathrm{BCR}_T = \mathbb{E}\left[\sum_{t=1}^{T} f^* - f(\boldsymbol{x}_t)\right], \quad \mathrm{BSR}_T = \mathbb{E}[f^* - f(\hat{\boldsymbol{x}}_T)] \qquad (T \in \mathbb{N}), \tag{11}$$

where $f^* = f(\boldsymbol{x}^*)$ and $\hat{\boldsymbol{x}}_T = \mathrm{argmax}_{\boldsymbol{x} \in \mathcal{X}} \mu(\boldsymbol{x}; \mathcal{D}_T)$. In Eq. (11), the expectation is taken with respect to all randomness, including f, ε_t, and the algorithm. The bounds of BSR and BCR are represented using maximum information gain (MIG) [38], defined as follows:

Algorithm 1 Randomized Kriging Believer

Input: input space $\mathcal{X}$, objective function $f : \mathcal{X} \to \mathbb{R}$, kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$, noise variance $\sigma^2_{\mathrm{noise}} > 0$, sequential BO algorithm $\mathcal{A}$
$\mathcal{D}_{\mathcal{N}_0} \leftarrow \emptyset$
for $t = 1, \ldots$ **do**
  Set $\mathcal{D}^{\mathrm{RKB}}_{t-1}$ by Eqs. (13) and (14)
  $\boldsymbol{x}_t \leftarrow \mathcal{A}\big(\mathcal{D}^{\mathrm{RKB}}_{t-1}\big)$
  Assign a worker to evaluate $\boldsymbol{x}_t$
  Wait for a free worker and set $\mathcal{N}_t$ and $\mathcal{D}_{\mathcal{N}_t}$ as the obtained indices and the dataset
end for

Table 1: Main symbols and their meanings.

Symbol	Meaning
$\mathcal{N}_{t-1}$	Index set of the completed evaluations at the start of iteration t
$\overline{\mathcal{N}}_{t-1}$	Index set of the pending evaluations at the start of iteration t
y_i	Observed value for evaluation $i \in \mathbb{N}$
$y_i^{(t)}$	Fantasized value for evaluation $i \in \overline{\mathcal{N}}_{t-1}$ used at iteration t
$\mathcal{D}_{\mathcal{N}_{t-1}}$	Observed data set at the start of iteration t
$\mathcal{D}^{\mathrm{KB}}_{t-1}$	Fantasized data set used at iteration t, where the fantasy $y_i^{(t)}$ is the predictive mean
$\mathcal{D}^{\mathrm{RKB}}_{t-1}$	Fantasized data set used at iteration t, where the fantasy $y_i^{(t)}$ is the posterior sample

Definition 2.1. *Let Assumption 2.1 hold. Then, for any $T \in \mathbb{N}$, MIG γ_T is defined as*

$$\gamma_T = \max_{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T \in \mathcal{X}} I(\boldsymbol{y}_T; \boldsymbol{f}_T), \tag{12}$$

where I denotes Shannon mutual information, $\boldsymbol{y}_T = [y_1 \ \cdots \ y_T]^\top$, and $\boldsymbol{f}_T = [f(\boldsymbol{x}_1) \ \cdots \ f(\boldsymbol{x}_T)]^\top$.

MIG γ_T has a sublinear and concave upper bound $\bar{\gamma}_T$ for commonly used kernels [18, 38, 49]: $\bar{\gamma}_T = O(d \log T)$ for linear kernels $k_{\mathrm{Lin}}(\boldsymbol{x}, \boldsymbol{x}') = \boldsymbol{x}^\top \boldsymbol{x}'$; $\bar{\gamma}_T = O\Big((\log T)^{d+1}\Big)$ for Gaussian kernels $k_{\mathrm{Gauss}}(\boldsymbol{x}, \boldsymbol{x}') = \exp\Big(-\frac{1}{2l^2}\left\|\boldsymbol{x} - \boldsymbol{x}'\right\|^2\Big)$; and $\bar{\gamma}_T = O\Big(T^{\frac{d}{2\nu+d}} (\log T)^{\frac{4\nu+d}{2\nu+d}}\Big)$ for Matérn-ν kernels $k_{\mathrm{Matern}}(\boldsymbol{x}, \boldsymbol{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}}{l}\left\|\boldsymbol{x} - \boldsymbol{x}'\right\|\right)^{\nu} J_\nu\Big(\frac{\sqrt{2\nu}}{l}\left\|\boldsymbol{x} - \boldsymbol{x}'\right\|\Big)$, where $l > 0$ and $\nu > 0$ are lengthscale and smoothness parameters, respectively, and Γ and J_ν are Gamma and modified Bessel functions of the second kind, respectively.

3 Randomized kriging believer

Algorithm 1 presents the pseudocode of RKB, and Table 1 summarizes the brief definitions of the symbols. The key difference from the original KB is that the fantasized data $\mathcal{D}^{\mathrm{RKB}}_{t-1}$ is defined as

$$\mathcal{D}^{\mathrm{RKB}}_{t-1} = \left\{\left(\boldsymbol{x}_i, y_i^{(t)}\right)\right\}_{i=1}^{t-1}, \quad y_i^{(t)} = \begin{cases} y_i & (i \in \mathcal{N}_{t-1}) \\ g_t(\boldsymbol{x}_i) + \varepsilon_i^{(t)} & \left(i \in \overline{\mathcal{N}}_{t-1}\right) \end{cases}, \tag{13}$$

where $g_t \sim p\big(f \mid \mathcal{D}_{\mathcal{N}_{t-1}}\big)$ and $\varepsilon_i^{(t)} \sim \mathcal{N}\big(0, \sigma^2_{\mathrm{noise}}\big)$. Thus, data set $\mathcal{D}^{\mathrm{RKB}}_{t-1}$ contains values of a posterior sample path g_t of f in place of real data that have not yet been obtained. Note that we do not need to generate the whole sample path g_t to generate $g_t(\boldsymbol{x}_i)$ because $g_t(\boldsymbol{x}_i)$ follows the normal distribution as

$$[g_t(\boldsymbol{x}_i)]_{i \in \overline{\mathcal{N}}_{t-1}} \sim \mathcal{N}(\boldsymbol{\mu}_g, \boldsymbol{K}_g), \tag{14}$$

where $\boldsymbol{\mu}_g = \big[\mu\big(\boldsymbol{x}_i; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\big]_{i \in \overline{\mathcal{N}}_{t-1}}$ and $\boldsymbol{K}_g \in \mathbb{R}^{|\overline{\mathcal{N}}_{t-1}| \times |\overline{\mathcal{N}}_{t-1}|}$ is a covariance matrix whose (i, j)-th element is $k\Big(\boldsymbol{x}_{[\overline{\mathcal{N}}_{t-1}]_i}, \boldsymbol{x}_{[\overline{\mathcal{N}}_{t-1}]_j}; \mathcal{D}_{\mathcal{N}_{t-1}}\Big)$, where $[\overline{\mathcal{N}}_{t-1}]_i$ is the i-th smallest element of $\overline{\mathcal{N}}_{t-1}$.

As a result of using a posterior realization as fantasized data, we obtain the following theoretical and practical benefits. First, from the theoretical perspective, $\mathcal{D}^{\mathrm{RKB}}_{t-1}$ has the same distribution as $\mathcal{D}_{t-1}$

conditioned on $\mathcal{D}_{\mathcal{N}_{t-1}}$ by construction. This property is essential for the regret analysis in the next section. Second, KB can be overconfident in the sense that the algorithm believes the point estimate, the posterior mean. In contrast, RKB incorporates posterior uncertainty in a randomized manner, as in TS, thereby controlling the exploration-exploitation trade-off based on both posterior mean and uncertainty in practice.

Our RKB algorithm is similar to the hallucination believer (HB) [44] for preferential BO [5]. HB leverages a posterior sample defined over the training inputs, which remains random because only preference data are observed, by which HB avoids a time-consuming MCMC method for the GP preference model. The algorithms of HB and RKB are closely related, as both rely on posterior sampling and conditioning on hallucinated observations. However, the motivation for using posterior samples differs fundamentally between the two approaches.

4 Regret analysis

This section provides the regret analysis of RKB and the theoretical conditions for base BO algorithms. We study regret in the Bayesian setting [10, 18, 22, 33, 38] defined as in Section 2, although the frequentist setting [7, 17, 19, 38] has also been extensively investigated.

4.1 Theoretical conditions for base BO algorithms

Our RKB can parallelize any BO algorithm $\mathcal{A}$. To obtain regret bounds, we need the following condition of $\mathcal{A}$:

Condition 4.1. Let $t \in \mathbb{N}$ be an arbitrary number. If either (i) Assumption 2.1 holds and $|\mathcal{X}| < \infty$, or (ii) Assumptions 2.1 and 2.2 hold, then the following propositions hold. There exist (possibly random) variables $u_t\left(\mathcal{D}_{t-1}\right)$ and $v_t\left(\mathcal{D}_{t-1}\right) \in \mathbb{R}$ that satisfy the following for all $\mathcal{D}_{t-1} \in 2^{(\mathcal{X}\times\mathbb{R})}$:

$$\mathbb{E}[f^* - f(\boldsymbol{x}_t) \mid \mathcal{D}_{t-1}] \leq \mathbb{E}[u_t(\mathcal{D}_{t-1})\sigma(\boldsymbol{x}_t; \mathcal{D}_{t-1}) + v_t(\mathcal{D}_{t-1}) \mid \mathcal{D}_{t-1}], \tag{15}$$

where $\boldsymbol{x}_t = \mathcal{A}(\mathcal{D}_{t-1})$. Moreover, there exist ζ_t and $\xi_t \in \mathbb{R}$ that satisfy the following for all $\mathcal{D}_{t-1} \in 2^{(\mathcal{X}\times\mathbb{R})}$:

$$\mathbb{E}\left[u_t^2(\mathcal{D}_{t-1})\right] \leq \zeta_t, \quad \mathbb{E}[v_t(\mathcal{D}_{t-1})] \leq \xi_t, \quad \sum_{t=1}^{T} \zeta_t = \tilde{O}(T), \quad \sum_{t=1}^{T} \xi_t = O(1), \tag{16}$$

where $\tilde{O}$ hides polylogarithmic factors.

Condition 4.1 is satisfied at least by UCB [38], randomized UCB [43, 46], PIMS [45], and EI from the maximum of a sample path (EIMS) [47] for both finite and continuous input domains. Note that, although Condition 4.1 holds for TS [33, 45], RKB combined with TS results in PTS. For more details, see Appendix F.

4.2 Regret bounds

First, we present the BCR bounds for finite and continuous input domains:

Theorem 4.1 (BCR bound for finite input domains)**.** *Suppose that Assumption 2.1 and Condition 4.1 hold and that γ_T has an upper bound $\bar{\gamma}_T$ concave for T. Let $|\mathcal{X}| < \infty$ and $\boldsymbol{x}_t = \mathcal{A}\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)$. Then, the following holds:*

$$\mathrm{BCR}_T \leq B_T + \sqrt{C_1 C_2 Q T \bar{\gamma}_{T/Q}} \qquad (T \in \mathbb{N}), \tag{17}$$

where $C_1 = 2/\log\left(1 + \sigma_{\mathrm{noise}}^{-2}\right)$, $C_2 = 2 + 2\log\left(|\mathcal{X}|/2\right)$, $Q = \max_{t\in\mathbb{N}} t - |\mathcal{N}_{t-1}|$, and B_T is defined as

$$B_T = \sqrt{C_1 \gamma_T \sum_{t=1}^{T} \zeta_t} + \sum_{t=1}^{T} \xi_t = \tilde{O}(\sqrt{T\gamma_T}) \quad (T \in \mathbb{N}). \tag{18}$$

Theorem 4.2 (BCR bound for continuous input domains)**.** *Suppose that Assumptions 2.1 and 2.2 and Condition 4.1 hold. Assume that the kernel is a linear kernel, a Gaussian kernel, or a Matérn-ν kernel with $\nu > 1$. Let $\bar{\gamma}_T$ be a concave upper bound of γ_T. Define $L_\sigma = O(1)$ as in*

Lemma E.5, $L = \max\left\{L_\sigma, b\left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)\right\}$*, and* $s_t = 2 + 2d\log\lceil drLt^2\rceil - 2\log 2$*. Let* $\boldsymbol{x}_t = \mathcal{A}\big(\mathcal{D}_{t-1}^{\mathrm{RKB}}\big)$*. Then, the following holds:*

$$\mathrm{BCR}_T \leq B_T + \frac{\pi^2}{3} + \frac{\pi^2}{6}\sqrt{s_T} + \sqrt{C_1 Q s_T T \bar{\gamma}_{T/Q}} \quad (T \in \mathbb{N}), \tag{19}$$

where C_1 *and* Q *are the same as in Theorem 4.1 and* B_T *is the same as that of Theorem 4.1.*

See Appendices B and C for the proof of Theorems 4.1 and 4.2, respectively.

The upper bounds $\tilde{O}(\sqrt{QT\bar{\gamma}_{T/Q}})$ in Theorems 4.1 and 4.2 are tighter with respect to the dependence on Q than the known results on BCR $\tilde{O}(\sqrt{QT\bar{\gamma}_T})$ [10, 22, 29] by leveraging the proof technique modified from [50], shown in Lemma B.2.[1] Let us consider the batch setting in which $T = QB$ for the number of batches $B \in \mathbb{N}$ to interpret the upper bound. Then, $\sqrt{QT\bar{\gamma}_{T/Q}} = Q\sqrt{B\bar{\gamma}_B}$. Hence, our BCR bounds are $o(B)$ and $O(Q)$ if $\bar{\gamma}_T = o(T/\log T)$ in contrast to that the existing bound $\tilde{O}(\sqrt{QT\bar{\gamma}_T})$ implies $\omega(Q)$.

The key step in the proof is decomposing the regret into the regret incurred in the fantasized sample path g_t and the error between the fantasized sample path g_t and the objective function f as follows:

$$\mathrm{BCR}_T = \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} g_t^* - g_t(\boldsymbol{x}_t)\right]}_{A_1} + \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} g_t(\boldsymbol{x}_t) - f(\boldsymbol{x}_t)\right]}_{A_2}, \tag{20}$$

where we leverage the fact that $\mathcal{D}_{t-1}$ and $\mathcal{D}_{t-1}^{\mathrm{RKB}}$ are identically distributed. By this decomposition, we can obtain the upper bound $A_1 \leq B_T$, which can be seen as regret incurred by sequential optimization. Actually, vanilla BO methods satisfying Condition 4.1 achieve the regret bound B_T as shown in Lemma A.1. Furthermore, A_2 can be bounded from above by $\tilde{O}\big(\sqrt{QT\gamma_{T/Q}}\big)$, which contains the penalized term Q incurred by imputing the posterior samples. For the proof of Theorem 4.2, we further consider the discretization error that is bounded above by $\frac{\pi^2}{3} + \frac{\pi^2}{6}s_T$.

Next, we show the BSR bound in a consistent way for finite and continuous input domains:

Theorem 4.3 (BSR bound)**.** *Suppose that Assumption 2.1 and Condition 4.1 hold. Additionally, suppose that either* $|\mathcal{X}| < \infty$ *or Assumption 2.2 holds. Let* $\boldsymbol{x}_t = \mathcal{A}\big(\mathcal{D}_{t-1}^{\mathrm{RKB}}\big)$*. Then, the following holds:*

$$\mathrm{BSR}_T \leq \frac{B_T}{T} \qquad (T \in \mathbb{N}), \tag{21}$$

where B_T *is the same as in Theorem 4.1.*

See Appendix D for the proof.

Importantly, Theorem 4.3 shows the vanishing BSR upper bound independent of Q. Although Iwazaki [18], Takeno and Iwazaki [39] have shown a tighter high-probability cumulative regret bound for the sequential BO methods, BSR bounds for sequential BO methods [33, 43, 45, 46], PTS [29], DPP-TS [29], and UCB-PE [9] have the same rate as ours. Thus, although we conjecture that our BSR bound remains loose when Q is small, this looseness is a common limitation for the prior studies. In addition, the existing PBO methods listed above are fully distributed or joint selection methods, as discussed in Section 1.1. Thus, RKB is the first greedy PBO method to explicitly encourage diversity and to achieve the BSR guarantee without dependence on Q.

The key observation for the proof is the monotonically decreasing property of the BSR, that is,

$$\mathrm{BSR}_t \geq \mathrm{BSR}_T, \tag{22}$$

for all $t \leq T$. Thus, we have the BSR upper bound by the average of fantasized regret as follows:

$$\mathrm{BSR}_T \leq \frac{1}{T}\mathbb{E}\left[\sum_{t=1}^{T} g_t^* - g_t(\boldsymbol{x}_t)\right]. \tag{23}$$

Then, this term can be bounded from above by B_T without a penalty term with respect to Q.

[1] We conjecture that the similar upper bound immediately follows from Lemma B.2 for TS-based methods [22, 29]

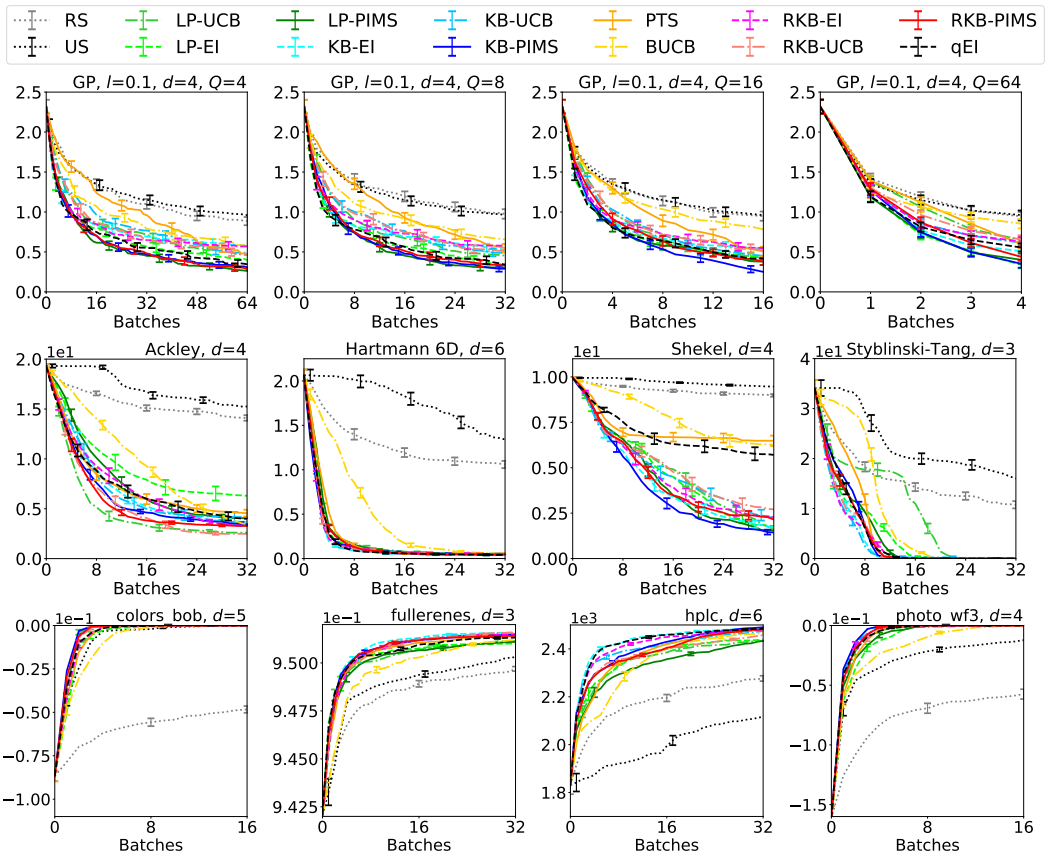


Figure 2: Mean and standard error of simple regret or best objective value across the 100 experiments on each condition. Rows correspond to the following objective functions: first, synthetic; second, benchmarks; third, emulators. One batch corresponds to Q iterations, and $Q = 8$ for benchmarks and emulators. Methods qLEI and qLNEI are denoted as qEI in the legend.

5 Experiments

We conducted experiments to demonstrate the efficiency of our RKB. As objective functions, we used synthetic and benchmark functions, as well as emulators trained on real-world data. As a performance measure, we report the simple regret $f(\boldsymbol{x}^*) - \max_{i \in [t]} f(\boldsymbol{x}_i)$ for synthetic and benchmark functions and the best objective value $\max_{i \in [t]} f(\boldsymbol{x}_i)$ for the emulators of real-world data since the optimum $f(\boldsymbol{x}^*)$ is unknown.

We compared RKB, KB, and LP in combination with UCB, EI, and PIMS. These combinations are denoted by connecting the names of the parallelization method and the AF with a hyphen, for example, KB-EI. Other methods included for comparison are BUCB, PTS, uncertainty sampling (US), and random search (RS). Variants of qEI implemented in BoTorch [3] are also included for comparison as follows: synthetic, q log noisy expected improvement (qLNEI); benchmark, q log expected improvement (qLEI); emulator, qEI. We chose the qEI variants based on the observation noise level (synthetic function experiments are noisy) and the package's compatibility with the emulator, which uses TensorFlow 1 [1]. Among the compared methods, RKB-PIMS, RKB-UCB, BUCB, PTS, and US have theoretical backing. For PIMS and PTS, which involve the posterior sampling, we consistently employ random Fourier features (RFF) [31] based approximation [54, 55]. Note that RKB itself does not need RFF-based approximation because the fantasized data are sampled from a multivariate normal distribution shown in Eq. (14).

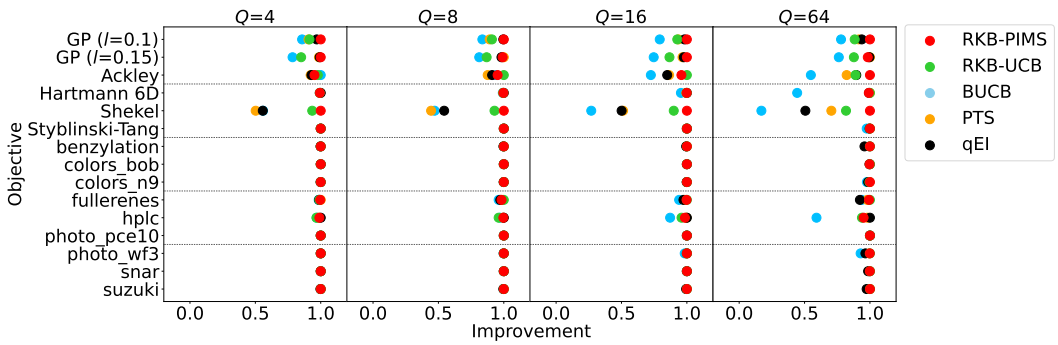


Figure 3: Mean improvement of the best objective value $\max_{i \in [t]} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition. The improvement amount has been normalized so that its maximum value among the five compared methods is one. Methods qLEI and qLNEI are denoted as qEI in the legend.

5.1 Setting

Common setting. We consider simple synchronous parallelization with $Q \in \{4, 8, 16, 64\}$ workers to facilitate the interpretation of the methods' effectiveness. In this setting, all ongoing evaluations are completed at the end of each iteration whose index is a multiple of Q. We also conducted experiments in the setting of asynchronous parallel optimization with $Q \in \{4, 8, 16, 64\}$, and the results are presented in Appendix G. We report the mean and standard error of the performance measure across 64 random trials for initial data generation, PBO algorithm's randomness, and synthetic objective function generation (in synthetic experiments only).

For synthetic function experiment. We generated the objective function by sampling from $\mathcal{GP}(0, k_{\mathrm{Gauss}})$, with input dimension $d = 4$ and lengthscale parameter $l \in \{0.1, 0.15\}$ in all dimensions. We show the results for $l = 0.1$ in this section, and the results for $l = 0.15$ are given in Appendix G. The search space is defined as $\mathcal{X} = \{0.1, 0.2, \ldots, 1.0\}^4$, which consists of 10^4 grid points. The variance of the Gaussian noise added to the observations is $\sigma^2_{\mathrm{noise}} = 10^{-3}$. As the prediction model, all algorithms used a GP with the true noise variance and the same kernel as that used to generate the objective function. Each experiment was initialized with 8 data points in $\mathcal{X}$ that were closest to 8 points in $[0, 1]^4$ chosen using Latin hypercube sampling (LHS) [26]. We set the confidence width parameter to $\beta_t = 2 \log \left(|\mathcal{X}| \, t^2 / \sqrt{2\pi} \right)$ for KB-UCB and RKB-UCB, following the theoretical value in [43]. For BUCB, we multiply β_t by $1 + ((t-1) \bmod Q) / \sigma^2_{\mathrm{noise}}$, where mod denotes the remainder operator, following the theoretical value in [10].

For benchmark and emulator experiment. We employed four benchmark objective functions, Ackley, Hartmann 6D, Shekel, and Styblinski–Tang on 4, 6, 4, and 3-dimensional search spaces, respectively, defined in `https://www.sfu.ca/~ssurjano/optimization.html`. As the objective function, we also used emulators provided by Olympus [14], a benchmarking framework for optimization. The search spaces of the benchmark and emulator functions are continuous hyperrectangular domains. The experiments were conducted on all emulators, except the "alkox" emulator, which returns only 0 and seems to have bugs. The prediction model is a GP using a Gaussian kernel with automatic relevance determination [32]. We selected the lengthscales and prior variance σ^2_{pri} of the kernel by marginal likelihood maximization [32] every Q iterations. For computational stability, the noise variance in the model was set to 10^{-8}. Each experiment was initialized with 16 data points chosen by LHS. We set the confidence width parameter to $\beta_t = 0.2d \log(2t)$, following the heuristics used in [20, 21]. As with the synthetic function experiments, we multiplied β_t by $1 + ((t-1) \bmod Q) \, \sigma^2_{\mathrm{pri}} / \sigma^2_{\mathrm{noise}}$ for BUCB. We show the results for $Q = 8$ in this section, and other results with $Q \in \{4, 16, 64\}$ are presented in Appendix G.

5.2 Result

Figure 2 shows the average and standard error of the performance measure across the 64 experiments on each condition. As shown in Fig. 2, RKB is comparable to KB and LP in each AF combination, except for KB- and LP-PIMS in the Shekel function experiment. In all cases, RKB-PIMS and RKB-UCB perform at least as well as other theoretically guaranteed methods, PTS, BUCB, and US. Moreover, PTS, BUCB, and US often underperform the proposed methods, likely due to over-exploration. Furthermore, in all cases, RKB-PIMS performs comparably to or better than qEI, a state-of-the-art method. While qEI underperformed RKB-PIMS significantly in the Shekel function experiment, we found no contrary example in our experiments. Figure 3 summarizes the results across all conditions, focusing on the improvement in the value of the objective function throughout the experiments. As shown in Fig. 3, RKB exhibits consistently high performance under various conditions.

6 Conclusion

We proposed the general-purpose PBO method, called randomized kriging believer (RKB), based on the widely used KB heuristic, inheriting the practical advantages of KB. Furthermore, we showed the BCR and BSR upper bounds for RKB comparable to TS-based PBO methods [29]. Finally, we demonstrated the effectiveness of the proposed method via extensive experiments.

There are several directions for future work. First, since our RKB is a general-purpose algorithm for parallelization, extensions to various problem settings, such as multi-fidelity [40, 41], multi-objective [16, 30], and constrained BO [42], are promising. Second, extending our regret analysis is of interest; for example, this includes analyses in the frequentist setting and its application to recent tighter regret bounds in [18, 39]. Finally, as in prior work, our BCR analysis requires uncertainty sampling in the initial phase to avoid an additional $O(\sqrt{Q})$ term. Addressing this common limitation in PBO analysis remains an important open problem.

Acknowledgments and Disclosure of Funding

This work was partially supported by JST CREST (Grant Numbers JPMJCR21D3 and JPMJCR22N2), JSPS KAKENHI (Grant Numbers JP26H02528 and JP24K20847), JST PRESTO (Grant Number JPMJPR24J6), and the RIKEN Center for Advanced Intelligence Project (AIP).

References

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL `https://www.tensorflow.org/`. Software available from tensorflow.org.

[2] J. Azimi, A. Fern, and X. Fern. Batch Bayesian optimization via simulation matching. In *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc., 2010.

[3] M. Balandat, B. Karrer, D. Jiang, S. Daulton, B. Letham, A. G. Wilson, and E. Bakshy. BoTorch: A framework for efficient Monte-Carlo Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 33, pages 21524–21538. Curran Associates, Inc., 2020.

[4] F. Berkenkamp, A. Krause, and A. P. Schoellig. Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *Machine learning*, 112(10):3713–3747, 2023.

[5] E. Brochu, V. M. Cora, and N. de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv.1012.2599*, 2010.

[6] C. Chevalier and D. Ginsbourger. Fast computation of the multi-points expected improvement with applications in batch selection. In *Learning and Intelligent Optimization*, pages 59–69. Springer Berlin Heidelberg, 2013.

[7] S. R. Chowdhury and A. Gopalan. On kernelized multi-armed bandits. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 844–853, 2017.

[8] S. R. Chowdhury and A. Gopalan. On batch bayesian optimization. *arXiv:1911.01032*, 2019.

[9] E. Contal, D. Buffoni, A. Robicquet, and N. Vayatis. Parallel Gaussian process optimization with upper confidence bound and pure exploration. In *Proceedings of the 2013th European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 225–240. Springer-Verlag, 2013.

[10] T. Desautels, A. Krause, and J. W. Burdick. Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization. *Journal of Machine Learning Research*, 15:4053–4103, 2014.

[11] J. Garcia-Barcos and R. Martinez-Cantin. Fully distributed Bayesian optimization with stochastic policies. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 2357–2363. International Joint Conferences on Artificial Intelligence Organization, 2019.

[12] D. Ginsbourger, R. Le Riche, and L. Carraro. *Kriging Is Well-Suited to Parallelize Optimization*, pages 131–162. Springer Berlin Heidelberg, 2010.

[13] J. Gonzalez, Z. Dai, P. Hennig, and N. Lawrence. Batch Bayesian optimization via local penalization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 648–657. PMLR, 2016.

[14] F. Häse, M. Aldeghi, R. J. Hickman, L. M. Roch, M. Christensen, E. Liles, J. E. Hein, and A. Aspuru-Guzik. Olympus: a benchmarking framework for noisy optimization and experiment planning. *Machine Learning: Science and Technology*, 2(3):035021, 2021.

[15] J. M. Hernández-Lobato, J. Requeima, E. O. Pyzer-Knapp, and A. Aspuru-Guzik. Parallel and distributed Thompson sampling for large-scale accelerated exploration of chemical space. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 1470–1479. PMLR, 2017.

[16] Y. Inatsu, S. Takeno, H. Hanada, K. Iwata, and I. Takeuchi. Bounding box-based multi-objective Bayesian optimization of risk measures under input uncertainty. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics*, volume 238 of *Proceedings of Machine Learning Research*, pages 4564–4572. PMLR, 2024.

[17] S. Iwazaki. Gaussian process upper confidence bound achieves nearly-optimal regret in noise-free Gaussian process bandits. In *Advances on Neural Information Processing Systems*, volume 38, pages 65863–65886. Curran Associates, Inc., 2025.

[18] S. Iwazaki. Improved regret bounds for Gaussian process upper confidence bound in Bayesian optimization. In *Advances on Neural Information Processing Systems*, volume 38, pages 96922–96964. Curran Associates, Inc., 2025.

[19] S. Iwazaki and S. Takeno. Improved regret analysis in Gaussian process bandits: Optimality for noiseless reward, RKHS norm, and non-stationary variance. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 26642–26672. PMLR, 2025.

[20] K. Kandasamy, J. Schneider, and B. Poczos. High dimensional Bayesian optimisation and bandits via additive models. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 295–304, 2015.

[21] K. Kandasamy, G. Dasarathy, J. Schneider, and B. Póczos. Multi-fidelity Bayesian optimisation with continuous approximations. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1799–1808, 2017.

[22] K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos. Parallelised Bayesian optimisation via Thompson sampling. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pages 133–142, 2018.

[23] T. Kathuria, A. Deshpande, and P. Kohli. Batched Gaussian process bandit optimization via determinantal point processes. In *Advances in Neural Information Processing Systems*, volume 29, pages 4206–4214. Curran Associates, Inc., 2016.

[24] S. Kusakawa, S. Takeno, Y. Inatsu, K. Kutsukake, S. Iwazaki, T. Nakano, T. Ujihara, M. Karasuyama, and I. Takeuchi. Bayesian optimization for cascade-type multistage processes. *Neural Computation*, 34(12):2408–2431, 2022.

[25] H. J. Kushner. A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *Journal of Basic Engineering*, 86(1):97–106, 1964.

[26] W.-L. Loh. On Latin hypercube sampling. *Ann. Statist.*, 24(5):2058–2080, 1996.

[27] S. Marmin, C. Chevalier, and D. Ginsbourger. Differentiating the multipoint expected improvement for optimal batch design. In *Revised Selected Papers of the First International Workshop on Machine Learning, Optimization, and Big Data*, pages 37–48. Springer-Verlag, 2015.

[28] J. Mockus, V. Tiesis, and A. Zilinskas. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2(117-129):2, 1978.

[29] E. Nava, M. Mutny, and A. Krause. Diversified sampling for batched Bayesian optimization with determinantal point processes. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 7031–7054, 2022.

[30] B. Paria, K. Kandasamy, and B. Póczos. A flexible framework for multi-objective Bayesian optimization using random scalarizations. In *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research*, pages 766–776, 2020.

[31] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems 20*, pages 1177–1184. Curran Associates, Inc., 2008.

[32] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[33] D. Russo and B. Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.

[34] J. Scarlett. Tight regret bounds for Bayesian optimization in one dimension. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4500–4508. PMLR, 2018.

[35] A. Shah and Z. Ghahramani. Parallel predictive entropy search for batch global optimization of expensive objective functions. In *Advances in Neural Information Processing Systems 28*, pages 3330–3338. Curran Associates, Inc., 2015.

[36] B. Shahriari, K. Swersky, Z. Wang, R. Adams, and N. De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.

[37] J. Snoek, H. Larochelle, and R. P. Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems 25*, pages 2951–2959. Curran Associates, Inc., 2012.

[38] N. Srinivas, A. Krause, S. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, pages 1015–1022. Omnipress, 2010.

[39] S. Takeno and S. Iwazaki. On regret bounds of Thompson sampling for Bayesian optimization. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306, pages 118354–118389. PMLR, 2026.

[40] S. Takeno, H. Fukuoka, Y. Tsukada, T. Koyama, M. Shiga, I. Takeuchi, and M. Karasuyama. Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 9334–9345. PMLR, 2020.

[41] S. Takeno, H. Fukuoka, Y. Tsukada, T. Koyama, M. Shiga, I. Takeuchi, and M. Karasuyama. A Generalized Framework of Multifidelity Max-Value Entropy Search Through Joint Entropy. *Neural Computation*, 34(10):2145–2203, 2022.

[42] S. Takeno, T. Tamura, K. Shitara, and M. Karasuyama. Sequential and parallel constrained max-value entropy search via information lower bound. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 20960–20986, 2022.

[43] S. Takeno, Y. Inatsu, and M. Karasuyama. Randomized Gaussian process upper confidence bound with tighter Bayesian regret bounds. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 33490–33515. PMLR, 2023.

[44] S. Takeno, M. Nomura, and M. Karasuyama. Towards practical preferential Bayesian optimization with skew Gaussian processes. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202, pages 33516–33533. PMLR, 2023.

[45] S. Takeno, Y. Inatsu, M. Karasuyama, and I. Takeuchi. Posterior sampling-based Bayesian optimization with tighter Bayesian regret bounds. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 47510–47534. PMLR, 2024.

[46] S. Takeno, Y. Inatsu, and M. Karasuyama. Regret analysis for randomized Gaussian process upper confidence bound. *Journal of Artificial Intelligence Research*, 84(18), 2025.

[47] S. Takeno, Y. Inatsu, M. Karasuyama, and I. Takeuchi. Regret analysis of posterior sampling-based expected improvement for Bayesian optimization. *Transactions on Machine Learning Research*, 2025. URL `https://openreview.net/forum?id=v0s9knY99c`.

[48] T. Ueno, T. D. Rhone, Z. Hou, T. Mizoguchi, and K. Tsuda. COMBO: An efficient Bayesian optimization library for materials science. *Materials discovery*, 4:18–21, 2016.

[49] S. Vakili, K. Khezeli, and V. Picheny. On information gain and regret bounds in Gaussian process bandits. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 82–90, 2021.

[50] S. Vakili, H. Moss, A. Artemev, V. Dutordoir, and V. Picheny. Scalable Thompson sampling using sparse Gaussian process models. In *Advances in Neural Information Processing Systems*, volume 34, pages 5631–5643. Curran Associates, Inc., 2021.

[51] R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.

[52] J. Wang, S. C. Clark, E. Liu, and P. I. Frazier. Parallel Bayesian global optimization of expensive functions. *Operations Research*, 68(6):1850–1865, 2020.

[53] J. Wilson, F. Hutter, and M. Deisenroth. Maximizing acquisition functions for Bayesian optimization. In *Advances in Neural Information Processing Systems 31*, pages 9906–9917. Curran Associates, Inc., 2018.

[54] J. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. Deisenroth. Efficiently sampling functions from Gaussian process posteriors. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pages 10292–10302. PMLR, 2020.

[55] J. T. Wilson, V. Borovitskiy, A. Terenin, P. Mostowsky, and M. P. Deisenroth. Pathwise conditioning of Gaussian processes. *Journal of Machine Learning Research*, 22(105):1–47, 2021.

[56] J. Wu and P. Frazier. The parallel knowledge gradient method for batch Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 29, pages 3126–3134. Curran Associates, Inc., 2016.

A Regret bound derived from Condition 4.1

In this section, we explain the meaning of the term B_T in Theorems 4.1, 4.2, and 4.3. For a sequencial BO method satisfying Condition 4.1, we have the following lemma:

Lemma A.1 (BCR bound for sequential optimization)**.** *Suppose that Condition 4.1 holds and that either (i) Assumption 2.1 holds and $|\mathcal{X}| < \infty$, or (ii) Assumptions 2.1 and 2.2 hold. Let $\boldsymbol{x}_t = \mathcal{A}(\mathcal{D}_{t-1})$. Then, the following holds:*

$$\mathrm{BCR}_T \leq B_T \qquad (T \in \mathbb{N}), \tag{24}$$

where B_T is defined using $C_1 = 2/\log\left(1+\sigma_{\mathrm{noise}}^{-2}\right)$ as

$$B_T = \sqrt{C_1 \gamma_T \sum_{t=1}^{T} \zeta_t} + \sum_{t=1}^{T} \xi_t = \tilde{O}(\sqrt{T\gamma_T}) \quad (T \in \mathbb{N}). \tag{25}$$

Lemma A.1 is a generalization of existing regret analyses. The term B_T is the derived regret bound defined using coefficients ζ_t and ξ_t depending on the BO method. For examples of ζ_t and ξ_t for a specific BO method, see Appendix F.

We use the following lemma, shown in Lemma 5.4 of [38], to prove Lemma A.1:

Lemma A.2. *Let Assumption 2.1 hold. Then, for any $T \in \mathbb{N}$ and $\mathcal{D}_T = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^T$, the following holds:*

$$\sum_{t=1}^{T} \sigma^2(\boldsymbol{x}_t; \mathcal{D}_{t-1}) \leq C_1 \gamma_T, \tag{26}$$

where $C_1 = 2/\log\left(1+\sigma_{\mathrm{noise}}^{-2}\right)$.

This lemma also plays an important role in the proof of our main theorem.

Proof of Lemma A.2. The lemma is proved as

$$\begin{aligned}
\sum_{t=1}^{T} \sigma^2(\boldsymbol{x}_t; \mathcal{D}_{t-1}) =& C_1 \sum_{t=1}^{T} \frac{1}{2} \log\left(1+\sigma_{\mathrm{noise}}^{-2}\right) \sigma^2(\boldsymbol{x}_t; \mathcal{D}_{t-1}) && \text{(27a)}\\
\leq& C_1 \sum_{t=1}^{T} \frac{1}{2} \log\left(1+\sigma_{\mathrm{noise}}^{-2} \sigma^2(\boldsymbol{x}_t; \mathcal{D}_{t-1})\right) && \text{(27b)}\\
=& C_1 I(\boldsymbol{y}_T; \boldsymbol{f}_T) && \text{(27c)}\\
\leq& C_1 \gamma_T, && \text{(27d)}
\end{aligned}$$

where the grounds of the equations are the following: Eq. (27b), Lemma E.3; Eq. (27c), Lemma E.1; Eq. (27d), the Definition 2.1 of γ_T. □

Using Lemma A.2, we prove Lemma A.1.

Proof of Lemma A.1. The proof is the following:

$$\begin{aligned}
\mathrm{BCR}_T = \mathbb{E}\left[\sum_{t=1}^{T} f^* - f(\boldsymbol{x}_t)\right] &= \sum_{t=1}^{T} \mathbb{E}_{\mathcal{D}_{t-1}}[\mathbb{E}[f^* - f(\boldsymbol{x}_t) \mid \mathcal{D}_{t-1}]] && \text{(28a)}\\
&\leq \sum_{t=1}^{T} \mathbb{E}_{\mathcal{D}_{t-1}}[\mathbb{E}[u_t \sigma(\boldsymbol{x}_t; \mathcal{D}_{t-1}) + v_t \mid \mathcal{D}_{t-1}]] && \text{(28b)}\\
&= \mathbb{E}\left[\sum\nolimits_{t=1}^{T} u_t \sigma(\boldsymbol{x}_t; \mathcal{D}_{t-1}) + v_t\right] && \text{(28c)}\\
&\leq \mathbb{E}\left[\sqrt{\sum\nolimits_{t=1}^{T} u_t^2}\sqrt{\sum\nolimits_{t=1}^{T} \sigma^2(\boldsymbol{x}_t; \mathcal{D}_{t-1})} + \sum\nolimits_{t=1}^{T} v_t\right] && \text{(28d)}\\
&\leq \mathbb{E}\left[\sqrt{\sum\nolimits_{t=1}^{T} u_t^2}\right]\sqrt{C_1\gamma_T} + \sum\nolimits_{t=1}^{T} \mathbb{E}[v_t] && \text{(28e)}\\
&\leq \sqrt{\mathbb{E}\left[\sum\nolimits_{t=1}^{T} u_t^2\right]}\sqrt{C_1\gamma_T} + \sum\nolimits_{t=1}^{T} \mathbb{E}[v_t] && \text{(28f)}\\
&\leq B_T, && \text{(28g)}
\end{aligned}$$

where the grounds of the equations are the following: Eq. (28b), Eq. (15) of Condition 4.1; Eq. (28d), Cauchy–Schwarz inequality; Eq. (28e), Lemma A.2; Eq. (28f), Jensen's inequality; Eq. (28g), Eq. (16) of Condition 4.1 and the definition of B_T. □

B Proof of Theorem 4.1

We use the following three lemmas to prove Theorem 1:

Lemma B.1. *Suppose that Condition 4.1 holds and that either (i) Assumption 2.1 holds and $|\mathcal{X}| < \infty$, or (ii) Assumptions 2.1 and 2.2 hold. Let $\boldsymbol{x}_t = \mathcal{A}\big(\mathcal{D}_{t-1}^{\mathrm{RKB}}\big)$. Then, the following holds:*

$$\mathbb{E}\left[\sum_{t=1}^{T} g_t^* - g_t(\boldsymbol{x}_t)\right] \leq B_T, \tag{29}$$

where $g_t \sim p\big(f \mid \mathcal{D}_{\mathcal{N}_{t-1}}\big)$ and $g_t^* = \max_{\boldsymbol{x}\in\mathcal{X}} g_t(\boldsymbol{x})$. Bound B_T is defined using $C_1 = 2/\log\left(1+\sigma_{\mathrm{noise}}^{-2}\right)$ as

$$B_T = \sqrt{C_1\gamma_T \sum\nolimits_{t=1}^{T} \zeta_t} + \sum_{t=1}^{T} \xi_t \qquad (T \in \mathbb{N})\,. \tag{30}$$

Proof of Lemma B.1. The proof is the following:

$$\mathbb{E}\left[\sum_{t=1}^{T} g_t^* - g_t(\boldsymbol{x}_t)\right] = \sum_{t=1}^{T} \mathbb{E}_{\mathcal{D}_{t-1}^{\mathrm{RKB}}}\left[\mathbb{E}\left[g_t^* - g_t(\boldsymbol{x}_t) \mid \mathcal{D}_{t-1}^{\mathrm{RKB}}\right]\right] \tag{31a}$$

$$= \sum_{t=1}^{T} \mathbb{E}_{\mathcal{D}_{t-1}^{\mathrm{RKB}}}\left[\mathbb{E}\left[f^* - f(\boldsymbol{x}_t) \mid \mathcal{D}_{t-1} = \mathcal{D}_{t-1}^{\mathrm{RKB}}\right]\right] \tag{31b}$$

$$\leq \sum_{t=1}^{T} \mathbb{E}_{\mathcal{D}_{t-1}^{\mathrm{RKB}}}\left[\mathbb{E}\left[u_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\sigma\left(\boldsymbol{x}_t; \mathcal{D}_{t-1}^{\mathrm{RKB}}\right) + v_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right) \mid \mathcal{D}_{t-1}^{\mathrm{RKB}}\right]\right] \tag{31c}$$

$$= \mathbb{E}\left[\sum\nolimits_{t=1}^{T} u_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\sigma\left(\boldsymbol{x}_t; \mathcal{D}_{t-1}^{\mathrm{RKB}}\right) + v_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\right] \tag{31d}$$

$$\leq \mathbb{E}\left[\sqrt{\sum\nolimits_{t=1}^{T} u_t^2\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)}\sqrt{\sum\nolimits_{t=1}^{T} \sigma^2\left(\boldsymbol{x}_t; \mathcal{D}_{t-1}^{\mathrm{RKB}}\right)} + \sum\nolimits_{t=1}^{T} v_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\right] \tag{31e}$$

$$\leq \mathbb{E}\left[\sqrt{\sum\nolimits_{t=1}^{T} u_t^2\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)}\sqrt{\sum\nolimits_{t=1}^{T} \sigma^2(\boldsymbol{x}_t; \mathcal{D}_{t-1})} + \sum\nolimits_{t=1}^{T} v_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\right] \tag{31f}$$

$$\leq \mathbb{E}\left[\sqrt{\sum\nolimits_{t=1}^{T} u_t^2\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)}\right]\sqrt{C_1\gamma_T} + \sum\nolimits_{t=1}^{T} \mathbb{E}\left[v_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\right] \tag{31g}$$

$$\leq \sqrt{\mathbb{E}\left[\sum\nolimits_{t=1}^{T} u_t^2\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\right]}\sqrt{C_1\gamma_T} + \sum\nolimits_{t=1}^{T} \mathbb{E}\left[v_t\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)\right] \tag{31h}$$

$$\leq B_T, \tag{31i}$$

where the grounds of the equations are the following: Eq. (31b), $p\left(g_t \mid \mathcal{D}_{t-1}^{\mathrm{RKB}}\right) = p\left(f \mid \mathcal{D}_{t-1} = \mathcal{D}_{t-1}^{\mathrm{RKB}}\right)$; Eq. (31c), Eq. (15) of Condition 4.1; Eq. (31e), Cauchy–Schwarz inequality; Eq. (31f), $\sigma\left(\boldsymbol{x}_t; \mathcal{D}_{t-1}^{\mathrm{RKB}}\right) = \sigma(\boldsymbol{x}_t; \mathcal{D}_{t-1})$; Eq. (31g), Lemma A.2; Eq. (31h), Jensen's inequality; Eq. (31i), Eq. (16) of Condition 4.1 and the definition of B_T.

□

Lemma B.2 (Modified from Lemma 3 of [50])**.** *Let Assumption 2.1 hold and MIG γ_T has an upper bound $\bar{\gamma}_T$ concave for T. Then, for any $t \in \mathbb{N}$, $\mathcal{D}_{t-1} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^{t-1} \in 2^{\mathcal{X}\times\mathbb{R}}$, and $\mathcal{N}_{t-1} \subset [t-1]$ satisfying $t - |\mathcal{N}_{t-1}| \leq Q$, the following holds:*

$$\sum_{t=1}^{T} \sigma^2\left(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\right) \leq C_1 Q \bar{\gamma}_{T/Q}. \tag{32}$$

Proof of Lemma B.2. Let $\mathcal{W}_{t-1} \subset [Q]$ be the set of indices of available workers for the t-th evaluation, and let the worker $w_t = \min \mathcal{W}_{t-1}$ be actually assigned to evaluation. Mathematically, the set $\mathcal{W}_{t-1}$ is defined as

$$\mathcal{W}_0 = [Q]\,, \qquad \mathcal{W}_t = (\mathcal{W}_{t-1} \setminus \{w_t\}) \cup \{w_{t'} \in [Q] \mid t' \in \mathcal{N}_t \setminus \mathcal{N}_{t-1}\} \qquad (t \in \mathbb{N})\,, \tag{33}$$

where $w_{t'}$ with $t' \in \mathcal{N}_t \setminus \mathcal{N}_{t-1}$ is a worker that has completed evaluation and become available again at iteration $t+1$. Let $\mathcal{T}_q \subset [T]$ be the set of evaluations that worker $q \in [Q]$ is assigned to, i.e.,

$$\mathcal{T}_q = \{t \leq T \mid w_t = q\}\,. \tag{34}$$

Obviously, the set $\mathcal{T}_q$ satisfies

$$\bigcup_{q=1}^{Q} \mathcal{T}_q = [T]\,, \qquad \forall q, q' \in [Q]\,, q \neq q' \Rightarrow \mathcal{T}_q \cap \mathcal{T}_{q'} = \emptyset. \tag{35}$$

Therefore, we have

$$\sum_{t=1}^{T}\sigma^2\big(\boldsymbol{x}_t;\mathcal{D}_{\mathcal{N}_{t-1}}\big)=\sum_{q=1}^{Q}\sum_{t\in\mathcal{T}_q}\sigma^2\big(\boldsymbol{x}_t;\mathcal{D}_{\mathcal{N}_{t-1}}\big)\leq\sum_{q=1}^{Q}\sum_{t\in\mathcal{T}_q}\sigma^2\big(\boldsymbol{x}_t;\mathcal{D}_{\mathcal{N}_{t-1}\cap\mathcal{T}_q}\big). \tag{36}$$

The inequality is obtained from $\mathcal{N}_{t-1}\supset\mathcal{N}_{t-1}\cap\mathcal{T}_q$. When worker q is assigned to evaluation t, worker q has already completed the past evaluations assigned to them, i.e.,

$$\forall t,t'\in\mathcal{T}_q, t>t'\Rightarrow t'\in\mathcal{N}_{t-1}. \tag{37}$$

Hence, the inner summation of the right-hand side of Eq. (36) is considered as the following operation: adding and learning each data in $\mathcal{D}_{\mathcal{T}_q}$ one by one and summing the prediction variances. Then, Lemma B.2 is proved as

$$\sum_{q=1}^{Q}\sum_{t\in\mathcal{T}_q}\sigma^2\big(\boldsymbol{x}_t;\mathcal{D}_{\mathcal{N}_{t-1}\cap\mathcal{T}_q}\big)\leq\sum_{q=1}^{Q}C_1\gamma_{|\mathcal{T}_q|}\leq C_1Q\sum_{q=1}^{Q}\frac{1}{Q}\bar{\gamma}_{|\mathcal{T}_q|}\leq C_1Q\bar{\gamma}_{T/Q}, \tag{38}$$

where the grounds are the following: the first inequality, Lemma A.2; the third inequality, Jensen's inequality and concavity of $\bar{\gamma}_T$. □

Lemma B.3. *Let Assumption 2.1 hold. For any $t\in\mathbb{N}$, $\mathcal{D}_{t-1}$, and finite $\mathcal{X}'\subset\mathcal{X}$, let $g_t\sim p(f\mid\mathcal{D}_{t-1})$ and define η_t, $\theta_t\in\mathbb{R}$ as*

$$\eta_t=\max_{\boldsymbol{x}\in\mathcal{X}'}\frac{g_t(\boldsymbol{x})-\mu(\boldsymbol{x};\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x};\mathcal{D}_{t-1})},\qquad\theta_t=\eta_t^2\mathbb{1}\{\eta_t\geq0\}, \tag{39}$$

where $\mathbb{1}\{\eta_t\geq0\}$ is 1 if $\eta_t\geq0$. Otherwise 0. Then, the following holds:

$$\mathbb{E}[\theta_t\mid\mathcal{D}_{t-1}]\leq2+2\log\left(|\mathcal{X}'|/2\right). \tag{40}$$

Proof of Lemma B.3. The proof mainly follows the proof of Lemma 3.2 of [45]. If $|\mathcal{X}'|=1$, η_t and η_t^2 follow the standard normal and Chi-squared distributions, respectively. Hence, we have $\mathbb{E}[\theta_t\mid\mathcal{D}_{t-1}]=1/2<0.61\simeq2+2\log\left(|\mathcal{X}'|/2\right)$. Next, we consider the case $|\mathcal{X}'|>1$. Let $\beta_\delta=2\log\left(|\mathcal{X}'|/(2\delta)\right)$ for any $\delta\in(0,1)$. We first show

$$\Pr(\theta_t\leq\beta_\delta\mid\mathcal{D}_{t-1})\geq1-\delta\qquad(\delta\in(0,1)). \tag{41}$$

The probability is transformed as

$$\begin{aligned}\Pr(\theta_t\leq\beta_\delta\mid\mathcal{D}_{t-1})=&\Pr\big(\eta_t^2\mathbb{1}\{\eta_t\geq0\}\leq\beta_\delta\mid\mathcal{D}_{t-1}\big)\\=&\Pr\Big(\eta_t\mathbb{1}\{\eta_t\geq0\}\leq\beta_\delta^{1/2}\mid\mathcal{D}_{t-1}\Big)\\=&\Pr\Big(\eta_t\leq\beta_\delta^{1/2}\mid\mathcal{D}_{t-1}\Big).\end{aligned} \tag{42}$$

We define $\boldsymbol{x}_{f,\eta}^*$ and $\boldsymbol{x}_{g,\eta}^*$ as

$$\boldsymbol{x}_{f,\eta}^*=\operatorname*{argmax}_{\boldsymbol{x}\in\mathcal{X}'}\frac{f(\boldsymbol{x})-\mu(\boldsymbol{x};\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x};\mathcal{D}_{t-1})},\qquad\boldsymbol{x}_{g,\eta}^*=\operatorname*{argmax}_{\boldsymbol{x}\in\mathcal{X}'}\frac{g_t(\boldsymbol{x})-\mu(\boldsymbol{x};\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x};\mathcal{D}_{t-1})}. \tag{43}$$

Because of $g_t\sim p(f\mid\mathcal{D}_{t-1})$, pairs $\Big(f,\boldsymbol{x}_{f,\eta}^*\Big)$ and $\big(g_t,\boldsymbol{x}_{g,\eta}^*\big)$ follow the same distribution under the condition $\mathcal{D}_{t-1}$. Hence, we have

$$\begin{aligned}\Pr\Big(\eta_t\leq\beta_\delta^{1/2}\mid\mathcal{D}_{t-1}\Big)=&\Pr\left(\frac{g_t\big(\boldsymbol{x}_{g,\eta}^*\big)-\mu\big(\boldsymbol{x}_{g,\eta}^*;\mathcal{D}_{t-1}\big)}{\sigma\big(\boldsymbol{x}_{g,\eta}^*;\mathcal{D}_{t-1}\big)}\leq\beta_\delta^{1/2}\mid\mathcal{D}_{t-1}\right)\\=&\Pr\left(\frac{f\Big(\boldsymbol{x}_{f,\eta}^*\Big)-\mu\Big(\boldsymbol{x}_{f,\eta}^*;\mathcal{D}_{t-1}\Big)}{\sigma\Big(\boldsymbol{x}_{f,\eta}^*;\mathcal{D}_{t-1}\Big)}\leq\beta_\delta^{1/2}\mid\mathcal{D}_{t-1}\right)\\=&\Pr\Big(f\big(\boldsymbol{x}_{f,\eta}^*\big)\leq\mu\big(\boldsymbol{x}_{f,\eta}^*;\mathcal{D}_{t-1}\big)+\beta_\delta^{1/2}\sigma\big(\boldsymbol{x}_{f,\eta}^*;\mathcal{D}_{t-1}\big)\mid\mathcal{D}_{t-1}\Big).\end{aligned} \tag{44}$$

Then, from Lemma E.2, we have

$$\begin{aligned}&\Pr\Big(f\big(\boldsymbol{x}_\eta^*\big) \le \mu\big(\boldsymbol{x}_\eta^*; \mathcal{D}_{t-1}\big) + \beta_\delta^{1/2}\sigma\big(\boldsymbol{x}_\eta^*; \mathcal{D}_{t-1}\big) \mid \mathcal{D}_{t-1}\Big)\\ \ge&\Pr\Big(\forall \boldsymbol{x} \in \mathcal{X}', f(\boldsymbol{x}) \le \mu(\boldsymbol{x}; \mathcal{D}_{t-1}) + \beta_\delta^{1/2}\sigma(\boldsymbol{x}; \mathcal{D}_{t-1}) \mid \mathcal{D}_{t-1}\Big)\\ \ge&1-\delta,\end{aligned} \tag{45}$$

which proves Eq. (41).

We define the cumulative distribution function F of θ_t conditioned by $\mathcal{D}_{t-1}$ and its generalized inverse function F^{-1} as

$$F(\beta) = \Pr(\theta_t \le \beta \mid \mathcal{D}_{t-1}), \quad F^{-1}(p) = \min\left\{\beta' \in \mathbb{R} \cup \{\infty\} \mid F(\beta') \ge p\right\} \quad (\beta \ge 0, p \in [0,1])\,. \tag{46}$$

Then, Eq. (41) is equivalent to

$$\beta_\delta \ge F^{-1}(1-\delta). \tag{47}$$

Hence, we have

$$2 + 2\log\left(|\mathcal{X}'|/2\right) = \int_0^1 \beta_\delta d\delta \ge \int_0^1 F^{-1}(1-\delta)d\delta = \int_0^1 F^{-1}(\delta)d\delta. \tag{48}$$

From the same principle of the inverse transform sampling, the following holds:

$$\int_0^1 F^{-1}(\delta)d\delta = \mathbb{E}_{\delta\sim\mathcal{U}(0,1)}\big[F^{-1}(\delta)\big] = \mathbb{E}[\theta_t \mid \mathcal{D}_{t-1}], \tag{49}$$

where $\mathcal{U}$ denotes the uniform distribution. This concludes the proof. □

Theorem 4.1 (BCR bound for finite input domains)**.** *Suppose that Assumption 2.1 and Condition 4.1 hold and that γ_T has an upper bound $\bar{\gamma}_T$ concave for T. Let $|\mathcal{X}| < \infty$ and $\boldsymbol{x}_t = \mathcal{A}\big(\mathcal{D}_{t-1}^{\mathrm{RKB}}\big)$. Then, the following holds:*

$$\mathrm{BCR}_T \le B_T + \sqrt{C_1 C_2 Q T \bar{\gamma}_{T/Q}} \qquad (T \in \mathbb{N})\,, \tag{17}$$

where $C_1 = 2/\log\big(1+\sigma_{\mathrm{noise}}^{-2}\big)$, $C_2 = 2 + 2\log\left(|\mathcal{X}|/2\right)$, $Q = \max_{t\in\mathbb{N}} t - |\mathcal{N}_{t-1}|$, and B_T is defined as

$$B_T = \sqrt{C_1\gamma_T \sum\nolimits_{t=1}^T \zeta_t} + \sum_{t=1}^T \xi_t = \tilde{O}(\sqrt{T\gamma_T}) \quad (T \in \mathbb{N})\,. \tag{18}$$

Proof of Theorem 4.1. Because of $g_t \sim p\big(f \mid \mathcal{D}_{\mathcal{N}_{t-1}}\big)$, we can transform BCR_T as

$$\begin{aligned}\mathrm{BCR}_T = \mathbb{E}\left[\sum_{t=1}^T f^* - f(\boldsymbol{x}_t)\right] =& \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{\mathcal{N}_{t-1}}}\left[\mathbb{E}\big[g_t^* - f(\boldsymbol{x}_t) \mid \mathcal{D}_{\mathcal{N}_{t-1}}\big]\right]\\ =& \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{\mathcal{N}_{t-1}}}\left[\mathbb{E}\big[g_t^* - g_t(\boldsymbol{x}_t) + g_t(\boldsymbol{x}_t) - f(\boldsymbol{x}_t) \mid \mathcal{D}_{\mathcal{N}_{t-1}}\big]\right]\\ =& \underbrace{\mathbb{E}\left[\sum_{t=1}^T g_t^* - g_t(\boldsymbol{x}_t)\right]}_{A_1} + \underbrace{\mathbb{E}\left[\sum_{t=1}^T g_t(\boldsymbol{x}_t) - f(\boldsymbol{x}_t)\right]}_{A_2}.\end{aligned} \tag{50}$$

From Lemma B.1, we have $A_1 \le B_T$. We transform A_2 as

$$\begin{aligned}A_2 = \mathbb{E}\left[\sum_{t=1}^T g_t(\boldsymbol{x}_t) - f(\boldsymbol{x}_t)\right] =& \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{\mathcal{N}_{t-1}}}\left[\mathbb{E}\big[g_t(\boldsymbol{x}_t) - f(\boldsymbol{x}_t) \mid \mathcal{D}_{\mathcal{N}_{t-1}}\big]\right]\\ =& \sum_{t=1}^T \mathbb{E}\big[g_t(\boldsymbol{x}_t) - \mu\big(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\big]\\ =& \sum_{t=1}^T \mathbb{E}\left[\frac{g_t(\boldsymbol{x}_t) - \mu\big(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\big)}{\sigma\big(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\big)}\sigma\big(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\right],\end{aligned} \tag{51}$$

Let $\eta_t = \max_{\boldsymbol{x}\in\mathcal{X}} \frac{g_t(\boldsymbol{x})-\mu\left(\boldsymbol{x};\mathcal{D}_{\mathcal{N}_{t-1}}\right)}{\sigma\left(\boldsymbol{x};\mathcal{D}_{\mathcal{N}_{t-1}}\right)}$ and $\theta_t = \eta_t^2 \mathbb{1}\{\eta_t \geq 0\}$. Then, A_2 is bounded above as

$$
\begin{aligned}
A_2 \leq \sum_{t=1}^{T} \mathbb{E}\left[\eta_t \sigma\left(\boldsymbol{x}_t;\mathcal{D}_{\mathcal{N}_{t-1}}\right)\right] \leq & \sum_{t=1}^{T} \mathbb{E}\left[\eta_t \mathbb{1}\{\eta_t \geq 0\} \sigma\left(\boldsymbol{x}_t;\mathcal{D}_{\mathcal{N}_{t-1}}\right)\right] && \text{(52a)}\\
\leq & \mathbb{E}\left[\sqrt{\textstyle\sum_{t=1}^{T} \eta_t^2 \mathbb{1}\{\eta_t \geq 0\}} \sqrt{\textstyle\sum_{t=1}^{T} \sigma^2\left(\boldsymbol{x}_t;\mathcal{D}_{\mathcal{N}_{t-1}}\right)}\right] && \text{(52b)}\\
\leq & \mathbb{E}\left[\sqrt{\textstyle\sum_{t=1}^{T} \theta_t}\right] \sqrt{C_1 Q \bar{\gamma}_{T/Q}} && \text{(52c)}\\
\leq & \sqrt{\textstyle\sum_{t=1}^{T} \mathbb{E}[\theta_t]} \sqrt{C_1 Q \bar{\gamma}_{T/Q}} && \text{(52d)}\\
= & \sqrt{\textstyle\sum_{t=1}^{T} \mathbb{E}_{\mathcal{D}_{\mathcal{N}_{t-1}}}\left[\mathbb{E}\left[\theta_t \mid \mathcal{D}_{\mathcal{N}_{t-1}}\right]\right]} \sqrt{C_1 Q \bar{\gamma}_{T/Q}} && \text{(52e)}\\
\leq & \sqrt{C_1 C_2 Q T \bar{\gamma}_{T/Q}}. && \text{(52f)}
\end{aligned}
$$

where $C_1 = 2/\log\left(1+\sigma_{\text{noise}}^{-2}\right)$ and $C_2 = 2 + 2\log\left(|\mathcal{X}|/2\right)$. The grounds of the equations are the following: Eq. (52b), Cauchy–Schwarz inequality; Eq. (52c), Lemma B.2; Eq. (52d), Jensen's inequality; Eq. (52f), Lemma B.3 with $\mathcal{X}' = \mathcal{X}$. As shown above, we have $\mathrm{BCR}_T = A_1 + A_2$, $A_1 \leq B_T$, and $A_2 \leq \sqrt{C_1 C_2 Q T \bar{\gamma}_{T/Q}}$, which proves Theorem 4.1. □

C Proof of Theorem 4.2

We first introduce the notation and lemmas needed to prove Theorem 4.2. Let $(l_t)_{t\in\mathbb{N}}$ be an arbitrary sequence of positive numbers. We equip $[0,r]^d$ with the metric induced by 1-norm $\|\cdot\|_1$. Then, the largest axis-aligned cube included in a ball of radius $l_t/2$ has edge length l_t/d. Hence, set $[0,r]^d$ can be covered by $\lceil dr/l_t \rceil^d$ balls of radius $l_t/2$. By this fact and Exercise 4.26 (b) of [51], set $\mathcal{X} \subset [0,r]^d$ can also be covered by the same number of balls with radius l_t. That is, there is some set of centers $\mathcal{X}_t \subset \mathcal{X}$ that satisfies $|\mathcal{X}_t| \leq \lceil dr/l_t \rceil^d$ and

$$
\forall \boldsymbol{x} \in \mathcal{X}, \exists \boldsymbol{x}' \in \mathcal{X}_t, \|\boldsymbol{x} - \boldsymbol{x}'\|_1 \leq l_t. \tag{53}
$$

For $\boldsymbol{x} \in \mathcal{X}$, we define $[\boldsymbol{x}]_t$ as one of $\boldsymbol{x}' \in \mathcal{X}_t$ that satisfies $\|\boldsymbol{x} - \boldsymbol{x}'\|_1 \leq l_t$.

We use the following lemmas to prove Theorem 2:

Lemma C.1 (Lemma H.2 of [43])**.** *Suppose that Assumptions 2.1, 2.2, and* $l_t \leq \frac{1}{bt^2} / \left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)$ *hold. Then, the following holds:*

$$
\sum_{t=1}^{T} \mathbb{E}\left[\sup_{\boldsymbol{x}\in\mathcal{X}} |f(\boldsymbol{x}) - f([\boldsymbol{x}]_t)|\right] \leq \frac{\pi^2}{6} \qquad (T \in \mathbb{N}). \tag{54}
$$

Lemma C.2 (Lemma D.4 of [45])**.** *Assume the same premise on k as Theorem E.5. Define L_σ by Eq. (78). Suppose that Assumption 2.1 and $l_t \leq \frac{1}{L_\sigma t^2}$ hold. Then, the following holds for any $t \in \mathbb{N}$ and $\mathcal{D}_{t-1}$:*

$$
\sup_{\boldsymbol{x}\in\mathcal{X}} |\sigma(\boldsymbol{x};\mathcal{D}_{t-1}) - \sigma([\boldsymbol{x}]_t;\mathcal{D}_{t-1})| \leq \frac{1}{t^2}. \tag{55}
$$

Theorem 4.2 (BCR bound for continuous input domains)**.** *Suppose that Assumptions 2.1 and 2.2 and Condition 4.1 hold. Assume that the kernel is a linear kernel, a Gaussian kernel, or a Matérn-ν kernel with $\nu > 1$. Let $\bar{\gamma}_T$ be a concave upper bound of γ_T. Define $L_\sigma = O(1)$ as in Lemma E.5, $L = \max\left\{L_\sigma, b\left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)\right\}$, and $s_t = 2 + 2d\log\lceil drLt^2\rceil - 2\log 2$. Let $\boldsymbol{x}_t = \mathcal{A}\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)$. Then, the following holds:*

$$
\mathrm{BCR}_T \leq B_T + \frac{\pi^2}{3} + \frac{\pi^2}{6}\sqrt{s_T} + \sqrt{C_1 Q s_T T \bar{\gamma}_{T/Q}} \quad (T \in \mathbb{N}), \tag{19}
$$

where C_1 and Q are the same as in Theorem 4.1 and B_T is the same as that of Theorem 4.1.

Proof of Theorem 4.2. In the same way as Eq. (50), we transform BCR_T as

$$\mathrm{BCR}_T = \mathbb{E}\left[\sum_{t=1}^{T} f^* - f(\boldsymbol{x}_t)\right] = \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} g_t^* - g_t(\boldsymbol{x}_t)\right]}_{A_1} + \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} g_t(\boldsymbol{x}_t) - f(\boldsymbol{x}_t)\right]}_{A_2}. \tag{56}$$

From Lemma B.1, we have $A_1 \leq B_T$. Let $l_t = \frac{1}{Lt^2}$ for $L = \max\left\{L_\sigma, b\left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)\right\}$. We transform A_2 as

$$\begin{aligned} A_2 =& \mathbb{E}\left[\sum_{t=1}^{T} g_t(\boldsymbol{x}_t) - f(\boldsymbol{x}_t)\right] \\ =& \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} g_t(\boldsymbol{x}_t) - g([\boldsymbol{x}_t]_t)\right]}_{A_3} + \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} g_t([\boldsymbol{x}_t]_t) - f([\boldsymbol{x}_t]_t)\right]}_{A_4} + \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} f([\boldsymbol{x}_t]_t) - f(\boldsymbol{x}_t)\right]}_{A_5} \end{aligned} \tag{57}$$

Then, from Lemma C.1, we have $A_3 \leq \pi^2/6$ and $A_5 \leq \pi^2/6$. We transform A_4 as

$$\begin{aligned} A_4 =& \mathbb{E}\left[\sum_{t=1}^{T} g_t([\boldsymbol{x}_t]_t) - f([\boldsymbol{x}_t]_t)\right] \\ =& \mathbb{E}_{\mathcal{D}_{\mathcal{N}_{t-1}}}\left[\mathbb{E}\left[\sum_{t=1}^{T} g_t([\boldsymbol{x}_t]_t) - f([\boldsymbol{x}_t]_t) \mid \mathcal{D}_{\mathcal{N}_{t-1}}\right]\right] \\ =& \mathbb{E}\left[\sum_{t=1}^{T} g_t([\boldsymbol{x}_t]_t) - \mu\big([\boldsymbol{x}_t]_t ; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\right] \\ =& \mathbb{E}\left[\sum_{t=1}^{T} \frac{g_t([\boldsymbol{x}_t]_t) - \mu\big([\boldsymbol{x}_t]_t ; \mathcal{D}_{\mathcal{N}_{t-1}}\big)}{\sigma\big([\boldsymbol{x}_t]_t ; \mathcal{D}_{\mathcal{N}_{t-1}}\big)} \sigma\big([\boldsymbol{x}_t]_t ; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\right] \end{aligned} \tag{58}$$

Let $\eta_t = \max_{\boldsymbol{x} \in \mathcal{X}_t} \frac{g_t(\boldsymbol{x}) - \mu\big(\boldsymbol{x}; \mathcal{D}_{\mathcal{N}_{t-1}}\big)}{\sigma\big(\boldsymbol{x}; \mathcal{D}_{\mathcal{N}_{t-1}}\big)}$. Then, A_4 is bounded above as

$$\begin{aligned} A_4 \leq& \mathbb{E}\left[\sum_{t=1}^{T} \eta_t \sigma\big([\boldsymbol{x}_t]_t ; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\right] \\ \leq& \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} \eta_t \sigma\big(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\right]}_{A_6} + \underbrace{\mathbb{E}\left[\sum_{t=1}^{T} \eta_t \left|\sigma\big(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\big) - \sigma\big([\boldsymbol{x}_t]_t ; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\right|\right]}_{A_7} \end{aligned} \tag{59}$$

In the same way as Eq. (52), A_6 is bounded above as

$$A_6 = \mathbb{E}\left[\sum_{t=1}^{T} \eta_t \sigma\big(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\big)\right] \leq \sqrt{C_1 Q s_T T \bar{\gamma}_{T/Q}}, \tag{60}$$

where $s_t \geq 2 + 2\log\left(|\mathcal{X}_t|/2\right)$ is defined as

$$s_t = 2 + 2d\log\lceil dr/l_t \rceil - 2\log 2 = 2 + 2d\log\lceil drLt^2 \rceil - 2\log 2 \qquad (t \in \mathbb{N}). \tag{61}$$

Also, A_7 is bounded above as

$$\begin{aligned}
A_7 =& \mathbb{E}\left[\sum_{t=1}^{T} \eta_t \left|\sigma\left(\boldsymbol{x}_t; \mathcal{D}_{\mathcal{N}_{t-1}}\right) - \sigma\left(\left[\boldsymbol{x}_t\right]_t; \mathcal{D}_{\mathcal{N}_{t-1}}\right)\right|\right] && \text{(62a)}\\
\leq& \sum_{t=1}^{T} \frac{1}{t^2} \mathbb{E}[\eta_t] && \text{(62b)}\\
\leq& \sum_{t=1}^{T} \frac{1}{t^2} \mathbb{E}[\eta_t \mathbb{1}\{\eta_t \geq 0\}] && \text{(62c)}\\
=& \sum_{t=1}^{T} \frac{1}{t^2} \mathbb{E}\left[\sqrt{\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}}\right] && \text{(62d)}\\
\leq& \sum_{t=1}^{T} \frac{1}{t^2} \sqrt{\mathbb{E}[\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}]} && \text{(62e)}\\
\leq& \sum_{t=1}^{T} \frac{1}{t^2} \sqrt{2 + 2\log\left(\left|\mathcal{X}_t\right|/2\right)} && \text{(62f)}\\
\leq& \sum_{t=1}^{T} \frac{1}{t^2} \sqrt{2 + 2\log\left(\left|\mathcal{X}_T\right|/2\right)} && \text{(62g)}\\
=& \frac{\pi^2}{6} \sqrt{s_T}. && \text{(62h)}
\end{aligned}$$

The grounds of the equations are the following: Eq. (62b), Lemma C.2; Eq. (62e), Jensen's inequality; Eq. (62f), Lemma B.3 with $\mathcal{X}' = \mathcal{X}_t$; Eq. (62g), the monotonicity of $|\mathcal{X}_t|$. Finally, the proof is concluded as

$$\begin{aligned}
&\mathrm{BCR}_T = A_1 + A_2 \leq A_1 + A_3 + A_5 + A_6 + A_7, \quad A_2 = A_3 + A_4 + A_5, \quad A_4 \leq A_6 + A_7,\\
&A_1 \leq B_T, \quad A_3 \leq \frac{\pi^2}{6}, \quad A_5 \leq \frac{\pi^2}{6}, \quad A_6 \leq \sqrt{C_1 Q s_T T \bar{\gamma}_{T/Q}}, \quad A_7 \leq \frac{\pi^2}{6}\sqrt{s_T}.
\end{aligned} \tag{63}$$

□

D Proof of Theorem 3

Theorem 4.3 (BSR bound)**.** *Suppose that Assumption 2.1 and Condition 4.1 hold. Additionally, suppose that either $|\mathcal{X}| < \infty$ or Assumption 2.2 holds. Let $\boldsymbol{x}_t = \mathcal{A}\left(\mathcal{D}_{t-1}^{\mathrm{RKB}}\right)$. Then, the following holds:*

$$\mathrm{BSR}_T \leq \frac{B_T}{T} \qquad (T \in \mathbb{N}), \tag{21}$$

where B_T is the same as in Theorem 4.1.

Proof of Theorem 4.3. The regret BSR_t is monotonically decreasing with respect to t as shown in the following:

$$\begin{aligned}
\mathrm{BSR}_t =& \mathbb{E}[f^* - f(\hat{\boldsymbol{x}}_t)]\\
=& \mathbb{E}[f^* - f(\hat{\boldsymbol{x}}_T) + f(\hat{\boldsymbol{x}}_T) - f(\hat{\boldsymbol{x}}_t)]\\
=& \mathbb{E}[f^* - f(\hat{\boldsymbol{x}}_T)] + \mathbb{E}_{\mathcal{D}_T}[\mathbb{E}[f(\hat{\boldsymbol{x}}_T) - f(\hat{\boldsymbol{x}}_t) \mid \mathcal{D}_T]]\\
=& \mathbb{E}[f^* - f(\hat{\boldsymbol{x}}_T)] + \mathbb{E}_{\mathcal{D}_T}[\mu(\hat{\boldsymbol{x}}_T; \mathcal{D}_T) - \mu(\hat{\boldsymbol{x}}_t; \mathcal{D}_T)]\\
\geq& \mathbb{E}[f^* - f(\hat{\boldsymbol{x}}_T)]\\
=& \mathrm{BSR}_T \qquad (T, t \in \mathbb{N}, T > t),
\end{aligned} \tag{64}$$

where the inequality is obtained because $\hat{\boldsymbol{x}}_T$ maximizes $\mu(\,\cdot\,;\mathcal{D}_T)$. Moreover, BSR_t is bounded above as

$$\begin{aligned}
\mathrm{BSR}_t =& \mathbb{E}[f^* - f(\hat{\boldsymbol{x}}_t)] \\
=& \mathbb{E}[f^* - f(\mathcal{A}(\mathcal{D}_{t-1})) + f(\mathcal{A}(\mathcal{D}_{t-1})) - f(\hat{\boldsymbol{x}}_t)] \\
=& \mathbb{E}[f^* - f(\mathcal{A}(\mathcal{D}_{t-1}))] + \mathbb{E}_{\mathcal{D}_t}[\mathbb{E}[f(\mathcal{A}(\mathcal{D}_{t-1})) - f(\hat{\boldsymbol{x}}_t) \mid \mathcal{D}_t]] \\
=& \mathbb{E}[f^* - f(\mathcal{A}(\mathcal{D}_{t-1}))] + \mathbb{E}_{\mathcal{D}_t}[\mu(\mathcal{A}(\mathcal{D}_{t-1}); \mathcal{D}_t) - \mu(\hat{\boldsymbol{x}}_t; \mathcal{D}_t)] \\
\leq& \mathbb{E}[f^* - f(\mathcal{A}(\mathcal{D}_{t-1}))] \\
=& \mathbb{E}_{\mathcal{D}_{\mathcal{N}_{t-1}}, (\boldsymbol{x}_i)_{i=1}^{t-1}}\left[\mathbb{E}\left[f^* - f(\mathcal{A}(\mathcal{D}_{t-1})) \mid \mathcal{D}_{\mathcal{N}_{t-1}}, (\boldsymbol{x}_i)_{i=1}^{t-1}\right]\right] \\
=& \mathbb{E}_{\mathcal{D}_{\mathcal{N}_{t-1}}, (\boldsymbol{x}_i)_{i=1}^{t-1}}\left[\mathbb{E}\left[g_t^* - g_t\big(\mathcal{A}\big(\mathcal{D}_{t-1}^{\mathrm{RKB}}\big)\big) \mid \mathcal{D}_{\mathcal{N}_{t-1}}, (\boldsymbol{x}_i)_{i=1}^{t-1}\right]\right] \\
=& \mathbb{E}\big[g_t^* - g_t\big(\mathcal{A}\big(\mathcal{D}_{t-1}^{\mathrm{RKB}}\big)\big)\big] \\
=& \mathbb{E}[g_t^* - g_t(\boldsymbol{x}_t)] \qquad\qquad (t \in \mathbb{N})\,.
\end{aligned} \tag{65}$$

Line 5 in Eq. (65) is obtained because $\hat{\boldsymbol{x}}_t$ maximizes $\mu(\,\cdot\,;\mathcal{D}_t)$. Line 7 in Eq. (65) is obtained from $g_t \sim p\big(f \mid \mathcal{D}_{\mathcal{N}_{t-1}}\big)$. From Eqs. (64) and (65), we have

$$\mathrm{BSR}_T \leq \frac{1}{T}\sum_{t=1}^{T}\mathrm{BSR}_t \leq \frac{1}{T}\underbrace{\mathbb{E}\left[\sum_{t=1}^{T} g_t^* - g_t(\boldsymbol{x}_t)\right]}_{A_1} \tag{66}$$

From Lemma B.1, we have $A_1 \leq B_T$, which concludes the proof. □

E Auxiliary lemmas

Lemma E.1 (Lemma 5.3 of [38])**.** *Let Assumption 2.1 hold. Then, for any $T \in \mathbb{N}$ and $\mathcal{D}_T = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^T$, the following holds:*

$$I(\boldsymbol{y}_T; \boldsymbol{f}_T) = \frac{1}{2}\sum_{t=1}^{T}\log\big(1 + \sigma_{\mathrm{noise}}^{-2}\sigma^2(\boldsymbol{x}_t; \mathcal{D}_{t-1})\big), \tag{67}$$

where $\boldsymbol{y}_T = [y_1 \ \cdots \ y_T]^\top$ and $\boldsymbol{f}_T = [f(\boldsymbol{x}_1) \ \cdots \ f(\boldsymbol{x}_T)]^\top$.

Lemma E.2. *Let Assumption 2.1 hold. Then, for any $\delta \in (0,1)$, $t \in \mathbb{N}$, $\mathcal{D}_{t-1}$, and finite $\mathcal{X}' \subset \mathcal{X}$, the following holds:*

$$\Pr\Big(\forall \boldsymbol{x} \in \mathcal{X}', f(\boldsymbol{x}) \leq \mu(\boldsymbol{x}; \mathcal{D}_{t-1}) + \beta_\delta^{1/2}\sigma(\boldsymbol{x}; \mathcal{D}_{t-1}) \mid \mathcal{D}_{t-1}\Big) \geq 1 - \delta, \tag{68}$$

where $\beta_\delta = 2\log\left(|\mathcal{X}'| \,/\, (2\delta)\right)$.

Proof of Lemma E.2. We define event $E(\boldsymbol{x})$ as

$$E(\boldsymbol{x}) \Leftrightarrow f(\boldsymbol{x}) \leq \mu(\boldsymbol{x}; \mathcal{D}_{t-1}) + \beta_\delta^{1/2}\sigma(\boldsymbol{x}; \mathcal{D}_{t-1}) \qquad (\boldsymbol{x} \in \mathcal{X}')\,. \tag{69}$$

Then, from Lemma E.4, we have

$$\Pr(\neg E(\boldsymbol{x}) \mid \mathcal{D}_{t-1}) = 1 - \Phi\left(\beta_\delta^{1/2}\right) \leq \frac{\delta}{|\mathcal{X}'|} \qquad (\boldsymbol{x} \in \mathcal{X}')\,. \tag{70}$$

Therefore, Lemma E.2 is proved as

$$\begin{aligned}
&\Pr(\forall \boldsymbol{x} \in \mathcal{X}', E(\boldsymbol{x}) \mid \mathcal{D}_{t-1}) \\
=& 1 - \Pr(\exists \boldsymbol{x} \in \mathcal{X}', \neg E(\boldsymbol{x}) \mid \mathcal{D}_{t-1}) \\
\geq& 1 - \sum_{\boldsymbol{x} \in \mathcal{X}'}\Pr(\neg E(\boldsymbol{x}) \mid \mathcal{D}_{t-1}) \\
\geq& 1 - \delta.
\end{aligned} \tag{71}$$

□

Lemma E.3. *The following inequality holds:*

$$(1+a)^b \leq 1+ab \qquad (a>0, b\in[0,1])\,. \tag{72}$$

Proof of Lemma E.3. Let a be an arbitrary positive number. We define $f:[0,1]\to\mathbb{R}$ as

$$f(x)=(1+a)^x \qquad (x\in[0,1])\,. \tag{73}$$

Then, f is convex because the second derivative f'' of f satisfies

$$f''(x)=(\log(1+a))^2(1+a)^x>0 \qquad (x\in[0,1])\,. \tag{74}$$

Therefore, we have

$$f(b)=(1+a)^b\leq(1-b)f(0)+bf(1)=1+ab \qquad (b\in[0,1])\,, \tag{75}$$

which proves Lemma E.3. □

Lemma E.4. *The survival function of the standard normal distribution can be bounded above as*

$$1-\Phi(c)\leq\frac{1}{2}\exp\left(-\frac{c^2}{2}\right) \qquad (c>0)\,, \tag{76}$$

where Φ is the cumulative distribution function of the standard normal distribution.

Proof of Lemma E.4. The proof is the following:

$$\begin{aligned}
1-\Phi(c)=&\int_c^\infty\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{r^2}{2}\right)dr\\
=&\exp\left(-\frac{c^2}{2}\right)\int_c^\infty\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{r^2-c^2}{2}\right)dr\\
=&\exp\left(-\frac{c^2}{2}\right)\int_c^\infty\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{(r-c)^2}{2}-c(r-c)\right)dr\\
\leq&\exp\left(-\frac{c^2}{2}\right)\int_c^\infty\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{(r-c)^2}{2}\right)dr\\
=&\frac{1}{2}\exp\left(-\frac{c^2}{2}\right).
\end{aligned} \tag{77}$$

□

Lemma E.5 (Theorem E.4 of [24])**.** *Let Assumption 2.1 hold. Suppose that the kernel k is one of the following: linear kernel, Gaussian kernel, or Matérn-ν kernel with $\nu>1$. Define L_σ as*

$$L_\sigma=\begin{cases}1 & \textit{(k is a linear kernel)},\\ \frac{\sqrt{2}}{l} & \textit{(k is a Gaussian kernel)},\\ \frac{\sqrt{2}}{l}\sqrt{\frac{\nu}{\nu-1}} & \textit{(k is a Matérn-ν kernel)}.\end{cases} \tag{78}$$

Then, the following holds:

$$\begin{aligned}
&\forall t\in\mathbb{N}, \forall\mathcal{D}_{t-1}\in 2^{(\mathcal{X}\times\mathbb{R})}, \forall\boldsymbol{x},\boldsymbol{x}'\in\mathcal{X},\\
&|\sigma(\boldsymbol{x};\mathcal{D}_{t-1})-\sigma(\boldsymbol{x}';\mathcal{D}_{t-1})|\leq L_\sigma\left\|\boldsymbol{x}-\boldsymbol{x}'\right\|_1.
\end{aligned} \tag{79}$$

F Examples where condition 4.1 holds

This section provides the details of five BO algorithms, UCB [38], improved randomized UCB (IR-UCB) [43], PIMS [45], EIMS [47], and TS [33] that satisfy Condition 4.1. We also provide u_t, v_t, ζ_t, and ξ_t in those cases. If $\mathcal{X}$ is infinite, we employ Assumption 2.2 and consider a finite subset $\mathcal{X}_t\subset\mathcal{X}$ for some $l_t>0$ as the proof of Theorem 4.2. For $\boldsymbol{x}\in\mathcal{X}$, let $[\boldsymbol{x}]_t$ be one of $\boldsymbol{x}'\in\mathcal{X}_t$ that satisfies $\left\|\boldsymbol{x}-\boldsymbol{x}'\right\|_1\leq l_t$. For convenience, we restate Condition 4.1.

Condition 4.1. Let $t \in \mathbb{N}$ be an arbitrary number. If either (i) Assumption 2.1 holds and $|\mathcal{X}| < \infty$, or (ii) Assumptions 2.1 and 2.2 hold, then the following propositions hold. There exist (possibly random) variables $u_t\left(\mathcal{D}_{t-1}\right)$ and $v_t\left(\mathcal{D}_{t-1}\right) \in \mathbb{R}$ that satisfy the following for all $\mathcal{D}_{t-1} \in 2^{(\mathcal{X}\times\mathbb{R})}$:

$$\mathbb{E}[f^* - f(\boldsymbol{x}_t) \mid \mathcal{D}_{t-1}] \leq \mathbb{E}[u_t(\mathcal{D}_{t-1})\sigma(\boldsymbol{x}_t; \mathcal{D}_{t-1}) + v_t(\mathcal{D}_{t-1}) \mid \mathcal{D}_{t-1}], \tag{15}$$

where $\boldsymbol{x}_t = \mathcal{A}(\mathcal{D}_{t-1})$. Moreover, there exist ζ_t and $\xi_t \in \mathbb{R}$ that satisfy the following for all $\mathcal{D}_{t-1} \in 2^{(\mathcal{X}\times\mathbb{R})}$:

$$\mathbb{E}\left[u_t^2(\mathcal{D}_{t-1})\right] \leq \zeta_t, \quad \mathbb{E}[v_t(\mathcal{D}_{t-1})] \leq \xi_t, \quad \sum_{t=1}^{T} \zeta_t = \tilde{O}(T), \quad \sum_{t=1}^{T} \xi_t = O(1), \tag{16}$$

where $\tilde{O}$ hides polylogarithmic factors.

F.1 UCB

In GP-UCB, the algorithm $\mathcal{A}$ is defined as

$$\mathcal{A}(\mathcal{D}_{t-1}) = \operatorname*{argmax}_{\boldsymbol{x}\in\mathcal{X}} \mu(\boldsymbol{x}; \mathcal{D}_{t-1}) + \beta_t^{1/2}\sigma(\boldsymbol{x}; \mathcal{D}_{t-1}). \tag{80}$$

If $\mathcal{X}$ is finite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \beta_t^{1/2}, \quad v_t\left(\mathcal{D}_{t-1}\right) = f(\boldsymbol{x}^*) - \mu(\boldsymbol{x}^*; \mathcal{D}_{t-1}) - \beta_t^{1/2}\sigma(\boldsymbol{x}^*; \mathcal{D}_{t-1}), \\ &\zeta_t = \beta_t, \quad \xi_t = 1/t^2, \end{aligned} \tag{81}$$

where $\beta_t = 2\log\left(|\mathcal{X}|\, t^2/\sqrt{2\pi}\right)$.

If $\mathcal{X}$ is infinite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \beta_t^{1/2}, \quad v_t\left(\mathcal{D}_{t-1}\right) = f(\boldsymbol{x}^*) - \mu([\boldsymbol{x}^*]_t; \mathcal{D}_{t-1}) - \beta_t^{1/2}\sigma([\boldsymbol{x}^*]_t; \mathcal{D}_{t-1}), \\ &\zeta_t = \beta_t, \quad \xi_t = 2/t^2, \end{aligned} \tag{82}$$

where $\beta_t = 2d\log\lceil dr/l_t\rceil + 2\log\left(t^2/\sqrt{2\pi}\right)$ and $l_t = \left(bt^2\left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)\right)^{-1}$. These results are given in the proof of Theorem B.1 of [43].

F.2 IR-UCB

In IRGP-UCB, the algorithm $\mathcal{A}$ is defined as Eq. (80) with random β_t. If $\mathcal{X}$ is finite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \beta_t^{1/2}, \quad v_t\left(\mathcal{D}_{t-1}\right) = 0, \\ &\zeta_t = 2 + 2\log\left(|\mathcal{X}|/2\right), \quad \xi_t = 0, \end{aligned} \tag{83}$$

where β_t follows the following distribution:

$$p(\beta_t) = \begin{cases} \frac{1}{2}\exp\left(-\frac{\beta_t - s}{2}\right) & (\beta_t \geq s), \\ 0 & (\beta_t < s), \end{cases} \qquad s = 2\log\left(|\mathcal{X}|/2\right). \tag{84}$$

This result is given in the proof of Theorem 4.2 of [43].

If $\mathcal{X}$ is infinite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \beta_t^{1/2}, \quad v_t\left(\mathcal{D}_{t-1}\right) = f(\boldsymbol{x}^*) - f([\boldsymbol{x}^*]_t), \\ &\zeta_t = 2 + 2d\log\lceil dr/l_t\rceil - 2\log 2, \quad \xi_t = 1/t^2, \end{aligned} \tag{85}$$

where β_t follows the following distribution:

$$p(\beta_t) = \begin{cases} \frac{1}{2}\exp\left(-\frac{\beta_t - s_t}{2}\right) & (\beta_t \geq s_t), \\ 0 & (\beta_t < s_t), \end{cases} \qquad s_t = 2d\log\lceil dr/l_t\rceil - 2\log 2, \tag{86}$$

where $l_t = \left(bt^2\left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)\right)^{-1}$. This result is given in the proof of Theorem 4.3 of [43].

F.3 PIMS

In PIMS, the algorithm $\mathcal{A}$ is defined as

$$\mathcal{A}(\mathcal{D}_{t-1}) = \operatorname*{argmax}_{\boldsymbol{x}\in\mathcal{X}} 1 - \Phi\left(\frac{g_t^* - \mu(\boldsymbol{x};\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x};\mathcal{D}_{t-1})}\right), \tag{87}$$

where $g_t \sim p(f \mid \mathcal{D}_{t-1})$ and $g_t^* = \max_{\boldsymbol{x}\in\mathcal{X}} g_t(\boldsymbol{x})$. If $\mathcal{X}$ is finite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \frac{g_t^* - \mu(\boldsymbol{x}_t;\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x}_t;\mathcal{D}_{t-1})}, \quad v_t\left(\mathcal{D}_{t-1}\right) = 0, \\ &\zeta_t = 2 + 2\log\left(|\mathcal{X}|/2\right), \quad \xi_t = 0. \end{aligned} \tag{88}$$

This result is given in the proof of Theorem 4.1 of [45].

If $\mathcal{X}$ is infinite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \frac{\tilde{g}_t^* - \mu(\tilde{\boldsymbol{x}}_t;\mathcal{D}_{t-1})}{\sigma(\tilde{\boldsymbol{x}}_t;\mathcal{D}_{t-1})}, \quad v_t\left(\mathcal{D}_{t-1}\right) = \frac{g_t^* - g_t([\boldsymbol{z}_t^*]_t)}{\sigma(\tilde{\boldsymbol{x}}_t;\mathcal{D}_{t-1})}\sigma(\boldsymbol{x}_t;\mathcal{D}_{t-1}), \\ &\zeta_t = 2 + 2d\log\lceil dr/l_t\rceil - 2\log 2, \quad \xi_t = 1/t^2, \end{aligned} \tag{89}$$

where $\tilde{g}_t^*$, $\boldsymbol{z}_t^*$, $\tilde{\boldsymbol{x}}_t$, and l_t are defined as

$$\begin{aligned} &\tilde{g}_t^* = \max_{\boldsymbol{x}\in\mathcal{X}_t} g_t(\boldsymbol{x}), \quad \boldsymbol{z}_t^* = \operatorname*{argmax}_{\boldsymbol{x}\in\mathcal{X}} g_t(\boldsymbol{x}), \quad \tilde{\boldsymbol{x}}_t = \operatorname*{argmin}_{\boldsymbol{x}\in\mathcal{X}_t} \frac{\tilde{g}_t^* - \mu(\boldsymbol{x};\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x};\mathcal{D}_{t-1})}, \\ &l_t = \left(bt^2\left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)\sqrt{\left(\sigma_{\text{noise}}^2 + t - 1\right)/\sigma_{\text{noise}}^2}\right)^{-1}. \end{aligned} \tag{90}$$

This result is given in the proof of Theorem 4.2 of [45].

F.4 EIMS

In EIMS, the algorithm $\mathcal{A}$ is defined as

$$\mathcal{A}(\mathcal{D}_{t-1}) = \operatorname*{argmax}_{\boldsymbol{x}\in\mathcal{X}} \mathbb{E}\left[\max\left\{f(\boldsymbol{x}) - g_t^*, 0\right\} \mid \mathcal{D}_{t-1}, g_t^*\right], \tag{91}$$

where $g_t \sim p(f \mid \mathcal{D}_{t-1})$ and $g_t^* = \max_{\boldsymbol{x}\in\mathcal{X}} g_t(\boldsymbol{x})$. If $\mathcal{X}$ is finite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \frac{g_t^* - \mu(\boldsymbol{x}_t;\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x}_t;\mathcal{D}_{t-1})}, \quad v_t\left(\mathcal{D}_{t-1}\right) = 0, \\ &\zeta_t = \log\left(\frac{\sigma_{\text{noise}}^2 + t - 1}{\sigma_{\text{noise}}^2}\right) + C_2 + \sqrt{2\pi C_2}, \quad \xi_t = 0. \end{aligned} \tag{92}$$

where $C_2 = 2 + 2\log\left(|\mathcal{X}|/2\right)$. This result is given in the proof of Theorem 4.6 of [47].

If $\mathcal{X}$ is infinite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \frac{g_t^* - \mu(\boldsymbol{x}_t;\mathcal{D}_{t-1})}{\sigma(\boldsymbol{x}_t;\mathcal{D}_{t-1})}, \quad v_t\left(\mathcal{D}_{t-1}\right) = 0, \\ &\zeta_t = \log\left(\frac{\sigma_{\text{noise}}^2 + t - 1}{\sigma_{\text{noise}}^2}\right) + C_t + \sqrt{2\pi C_t}, \quad \xi_t = 0, \end{aligned} \tag{93}$$

where C_t is defined as

$$C_t = 8\left(1 + d\log m_t\right), \quad m_t = \max\left\{2, bdr\sqrt{\left(\sigma_{\text{noise}}^2 + t - 1\right)\log\left(2ad\right)/\sigma_{\text{noise}}^2}\right\}. \tag{94}$$

This result is given in the proof of Theorem 4.8 of [47].

F.5 TS

In TS, the algorithm $\mathcal{A}$ is defined as

$$\mathcal{A}(\mathcal{D}_{t-1}) = \operatorname*{argmax}_{\boldsymbol{x}\in\mathcal{X}} g_t(\boldsymbol{x}), \quad g_t \sim p(f \mid \mathcal{D}_{t-1}). \tag{95}$$

If $\mathcal{X}$ is finite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \frac{g_t(\boldsymbol{x}_t) - \mu(\boldsymbol{x}_t; \mathcal{D}_{t-1})}{\sigma(\boldsymbol{x}_t; \mathcal{D}_{t-1})}, \quad v_t\left(\mathcal{D}_{t-1}\right) = 0, \\ &\zeta_t = 2 + 2\log\left(|\mathcal{X}|/2\right), \quad \xi_t = 0. \end{aligned} \tag{96}$$

This result is given in the proof of Theorem 3.1 of [45].

If $\mathcal{X}$ is infinite, Condition 4.1 holds for the following bounds:

$$\begin{aligned} &u_t\left(\mathcal{D}_{t-1}\right) = \frac{g_t([\boldsymbol{x}_t]_t) - \mu([\boldsymbol{x}_t]_t; \mathcal{D}_{t-1})}{\sigma([\boldsymbol{x}_t]_t; \mathcal{D}_{t-1})}, \\ &v_t\left(\mathcal{D}_{t-1}\right) = f(\boldsymbol{x}^*) - f([\boldsymbol{x}^*]_t) + f([\boldsymbol{x}_t]_t) - f(\boldsymbol{x}_t) + u_t\left(\mathcal{D}_{t-1}\right)\left|\sigma([\boldsymbol{x}_t]_t; \mathcal{D}_{t-1}) - \sigma(\boldsymbol{x}_t; \mathcal{D}_{t-1})\right|, \\ &\zeta_t = 2 + 2d\log\left(\lceil dr/l_t\rceil\right) - 2\log 2, \quad \xi_t = \left(2 + \sqrt{\zeta_t}\right)/t^2, \end{aligned} \tag{97}$$

where l_t is defined using L_σ in Lemma E.5 as

$$l_t = \frac{1}{Lt^2}, \quad L = \max\left\{L_\sigma, b\left(\sqrt{\log(ad)} + \sqrt{\pi}/2\right)\right\}. \tag{98}$$

This result is given in the proof of Theorem 3.2 of [47].

G Result of additional experiments

We present the results of experiments in synchronous and asynchronous settings with various numbers of workers, $Q = \{4, 8, 16, 64\}$.

Setting of synchronous experiments. The setting is the same as that of Section 5, except for Q.

Setting of asynchronous experiments. We assumed that each evaluation takes time that follows an independent standard half-normal distribution. We report the change in the simple regret $f(\boldsymbol{x}^*) - \max_{i\in\mathcal{N}_{t(\tau)}} f(\boldsymbol{x}_i)$ or the best objective value $\max_{i\in\mathcal{N}_{t(\tau)}} f(\boldsymbol{x}_i)$ over time $\tau \geq 0$ instead of iterations, where $t(\tau)$ is the last iteration completed by time τ. Each experiment was conducted 64 times under different conditions and terminated after a certain number of evaluations. The number of evaluations is 256 for synthetic objectives, benchmark objectives, and emulators "fullerenes" and "hplc". Other emulators are evaluated 128 times in each experiment. Because the end time varies for each experiment, we present the results up to the earliest end time across the 64 experimental trials. In the benchmark and emulator experiments, we selected the hyperparameters of the predictive model by marginal likelihood maximization every 4 iterations. Other settings are the same as those of the experiments in the synchronous setting.

Result. The experimental results are presented in the following figures: for synthetic objectives, Figs. 4 and 5; for benchmark objectives, Figs. 6 and 7; and for emulators, Figs. 8, 9, 10, and 11. Moreover, Fig. 12 summarizes the results across all conditions, focusing on the improvement in the value of the objective function throughout the experiments. As shown in Fig. 12, RKB-PIMS and RKB-UCB perform at least as well as other theoretically guaranteed methods, such as PTS and BUCB. In particular, RKB-PIMS consistently demonstrates strong performance, even compared to qEI, a state-of-the-art method. These results indicate that RKB has not only theoretical support but also strong practical performance. It can also be observed that, when the degree of parallelism is high, the performance of KB often degrades, whereas RKB is much less affected. This trend can be observed, for example, in the experiments on the Ackley function and the "snar" and "suzuki" emulators. We conjecture that this preferable property of RKB comes from its inherent randomness. That is, since KB must be overly confident that the current predictive mean is correct, KB gradually approaches an overly exploitative strategy if the degree of parallelism becomes high. Our RKB can avoid this tendency by using the random posterior samples, which reflect the remaining uncertainty of the posterior.

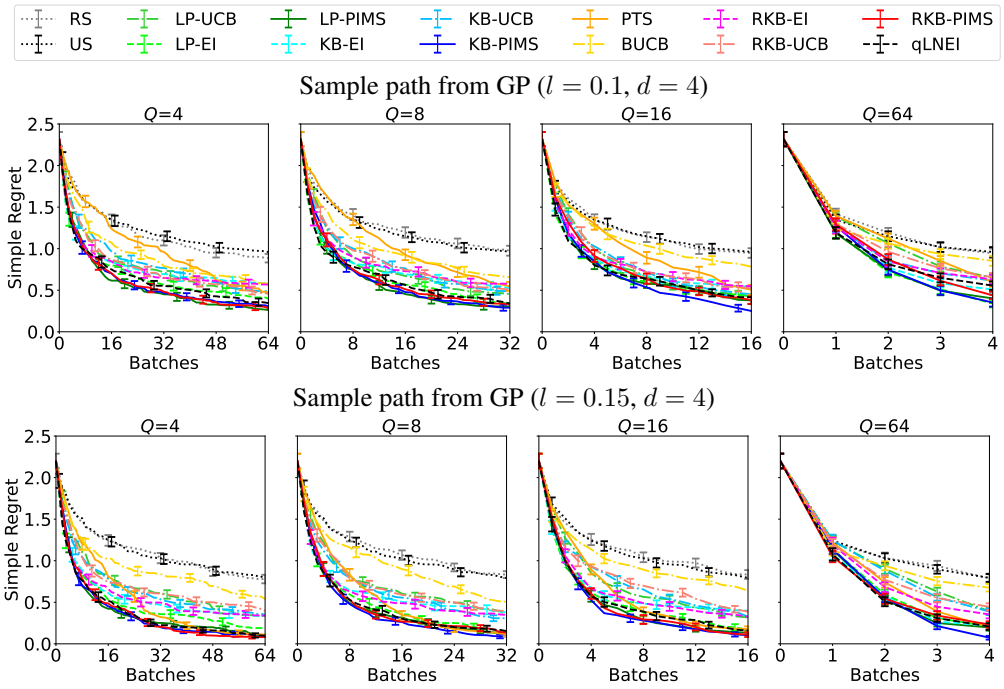


Figure 4: Result of experiments on synthetic objectives with synchronous setting. The lines and error bars mean average and standard error of the simple regret $f^* - \max_{i \in [t]} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition. One batch corresponds to Q iterations.

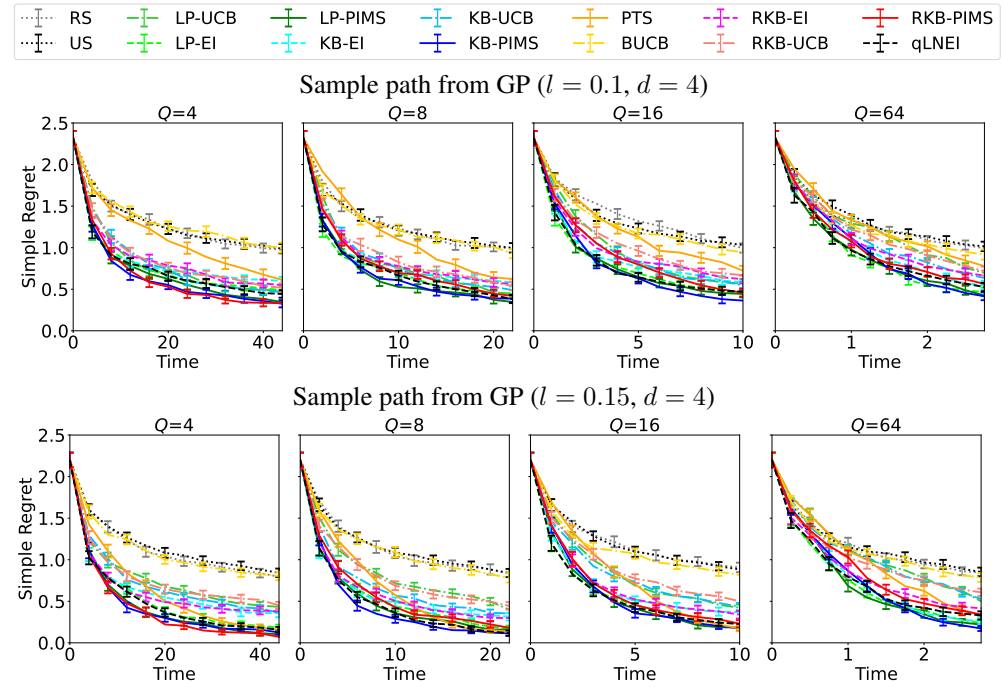


Figure 5: Result of experiments on synthetic objectives with asynchronous setting. The lines and error bars mean average and standard error of the simple regret $f^* - \max_{i \in \mathcal{N}_{t(\tau)}} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition, where $t(\tau)$ is the last iteration completed by time τ.

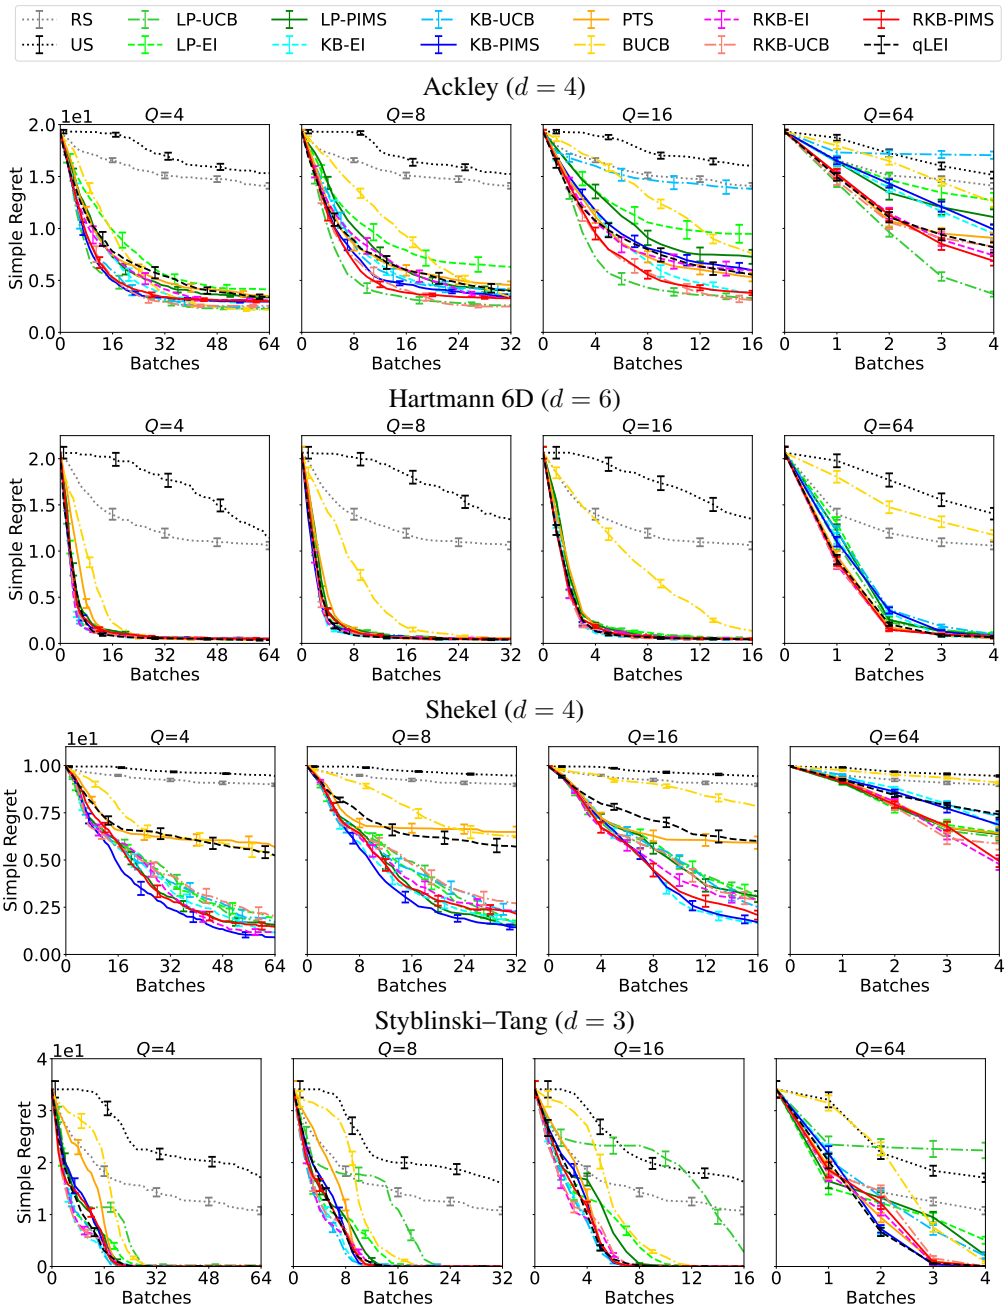


Figure 6: Result of experiments on benchmark functions with synchronous setting. The lines and error bars mean average and standard error of the simple regret $f^* - \max_{i \in [t]} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition. One batch corresponds to Q iterations.

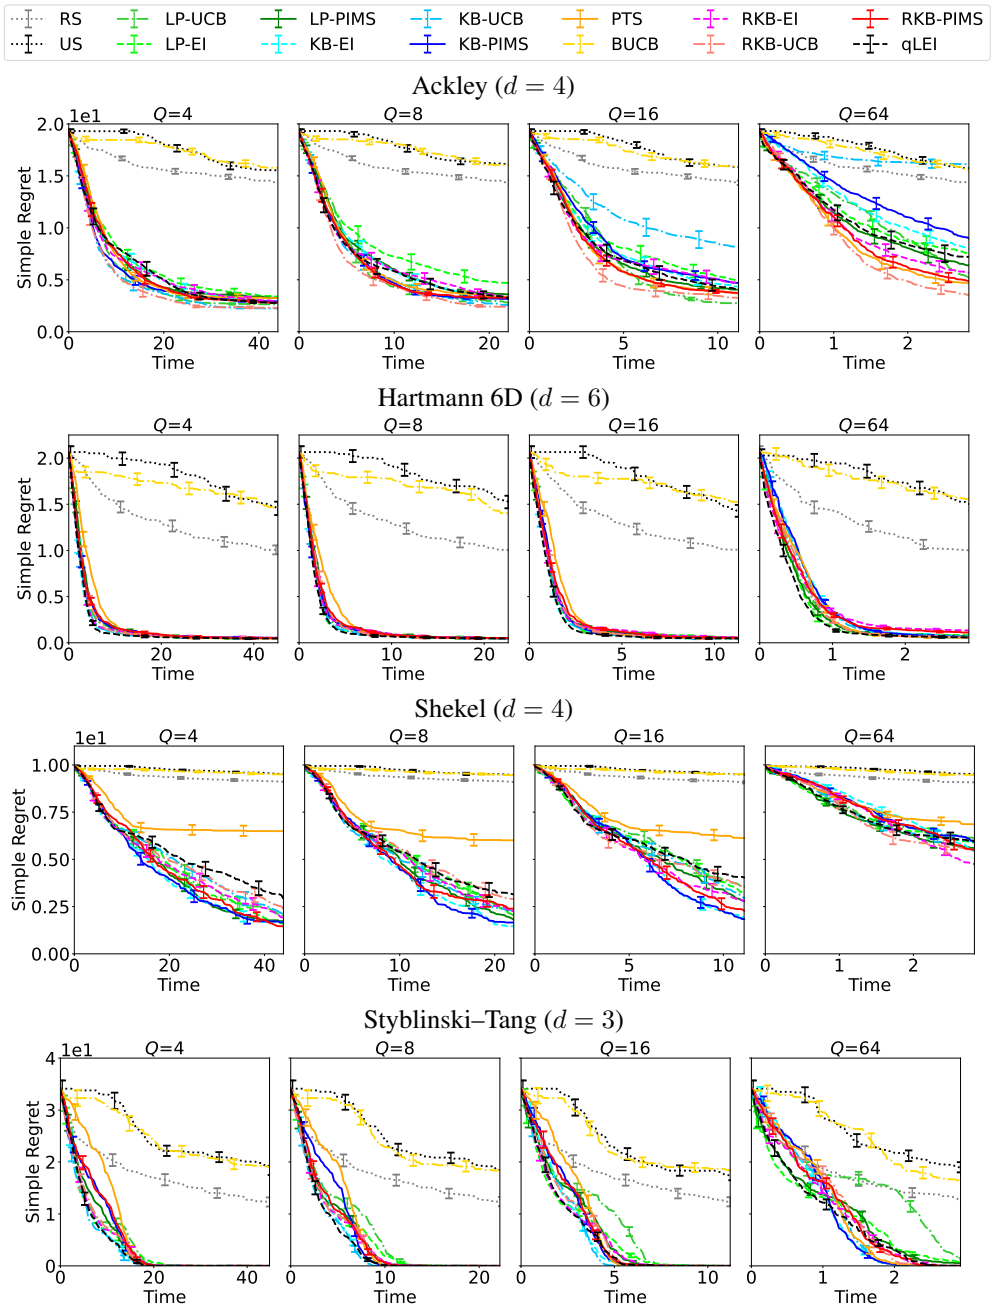


Figure 7: Result of experiments on benchmark functions with asynchronous setting. The lines and error bars mean average and standard error of the simple regret $f^* - \max_{i \in \mathcal{N}_{t(\tau)}} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition, where $t(\tau)$ is the last iteration completed by time τ.

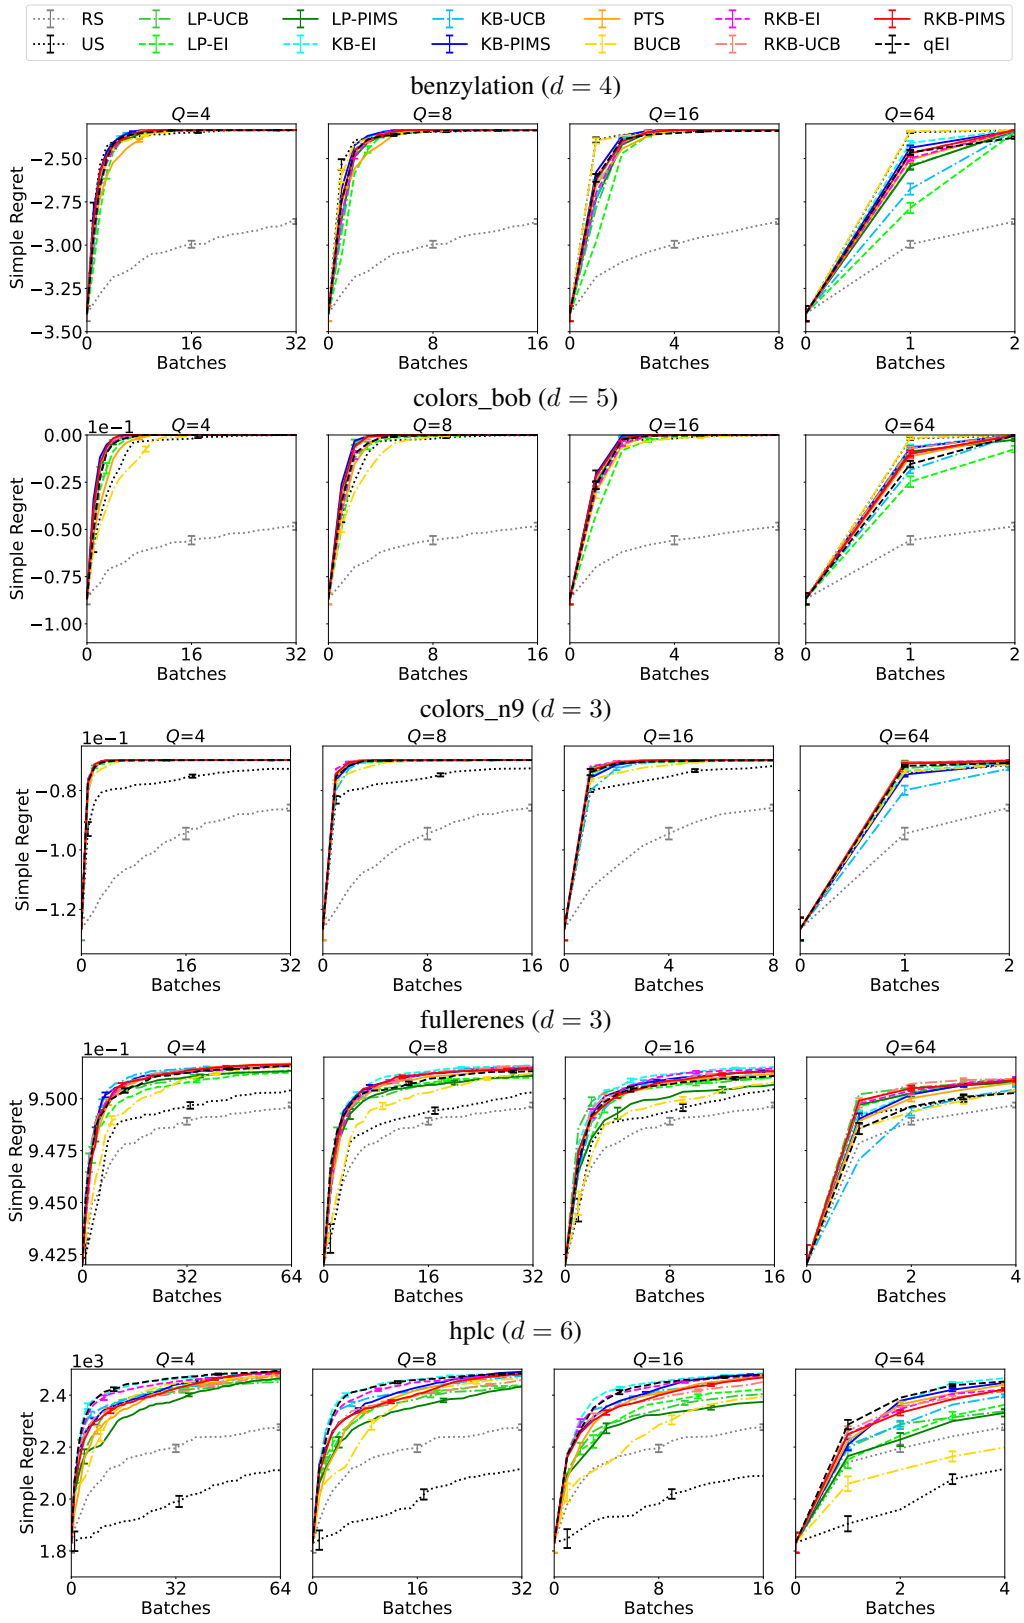


Figure 8: Result of experiments on emulators with synchronous setting. The lines and error bars mean average and standard error of the best objective value $\max_{i \in [t]} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition. One batch corresponds to Q iterations.

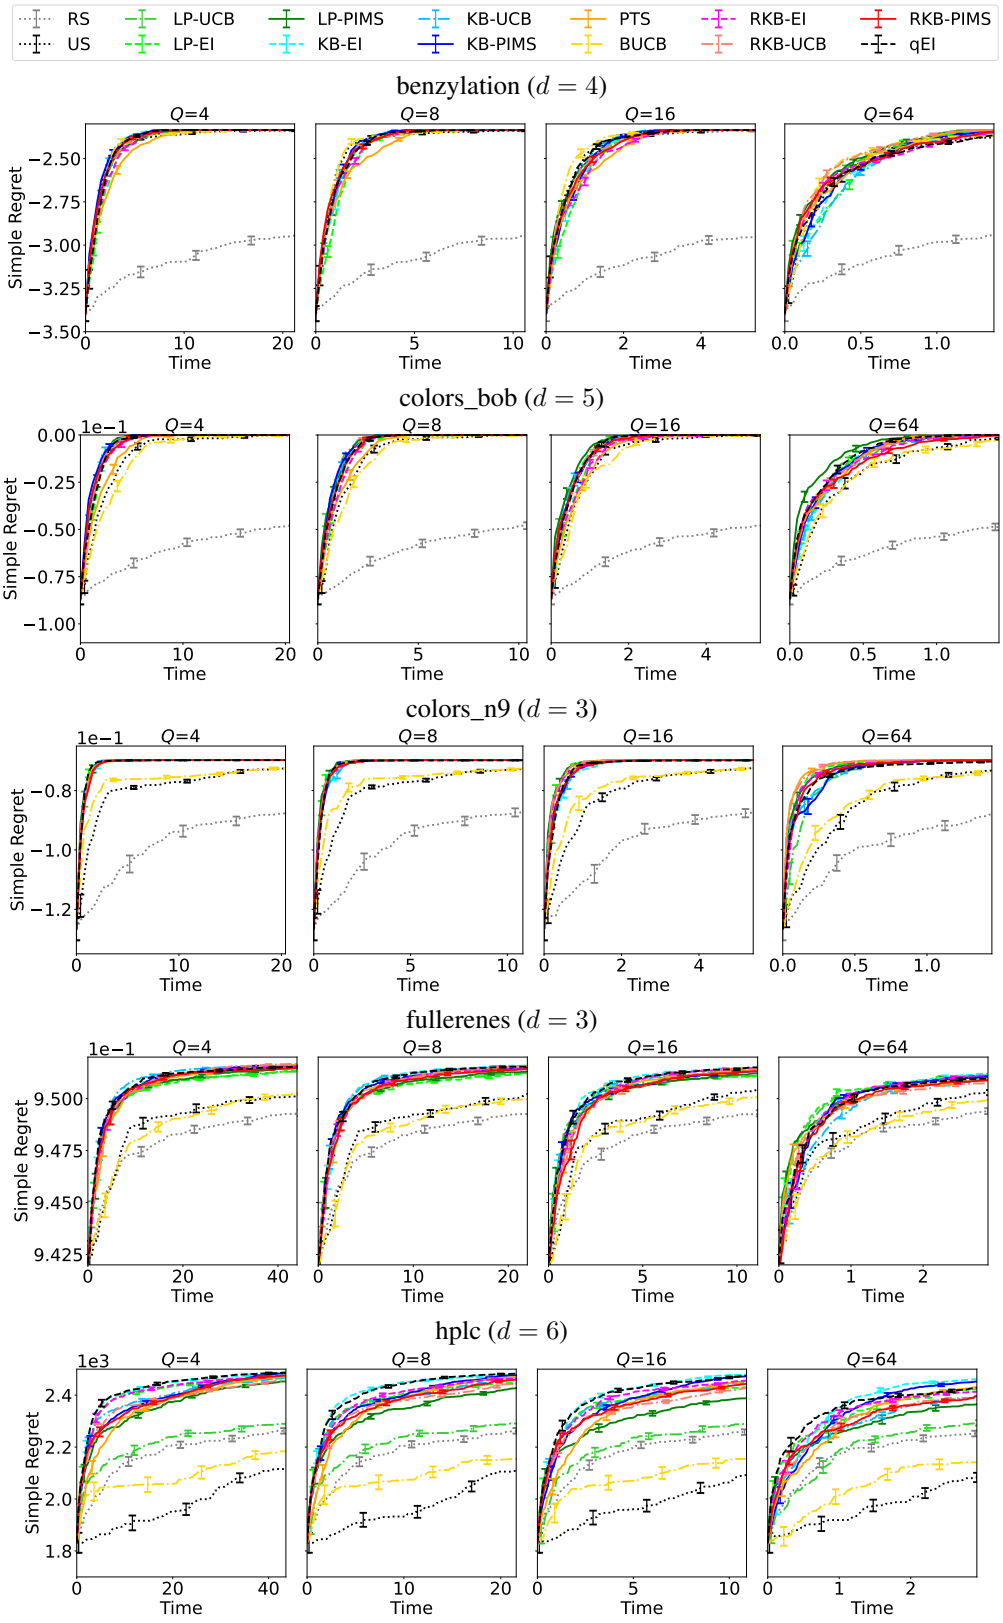


Figure 9: Result of experiments on emulators with asynchronous setting. The lines and error bars mean average and standard error of the best objective value $\max_{i \in \mathcal{N}_{t(\tau)}} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition, where $t(\tau)$ is the last iteration completed by time τ.

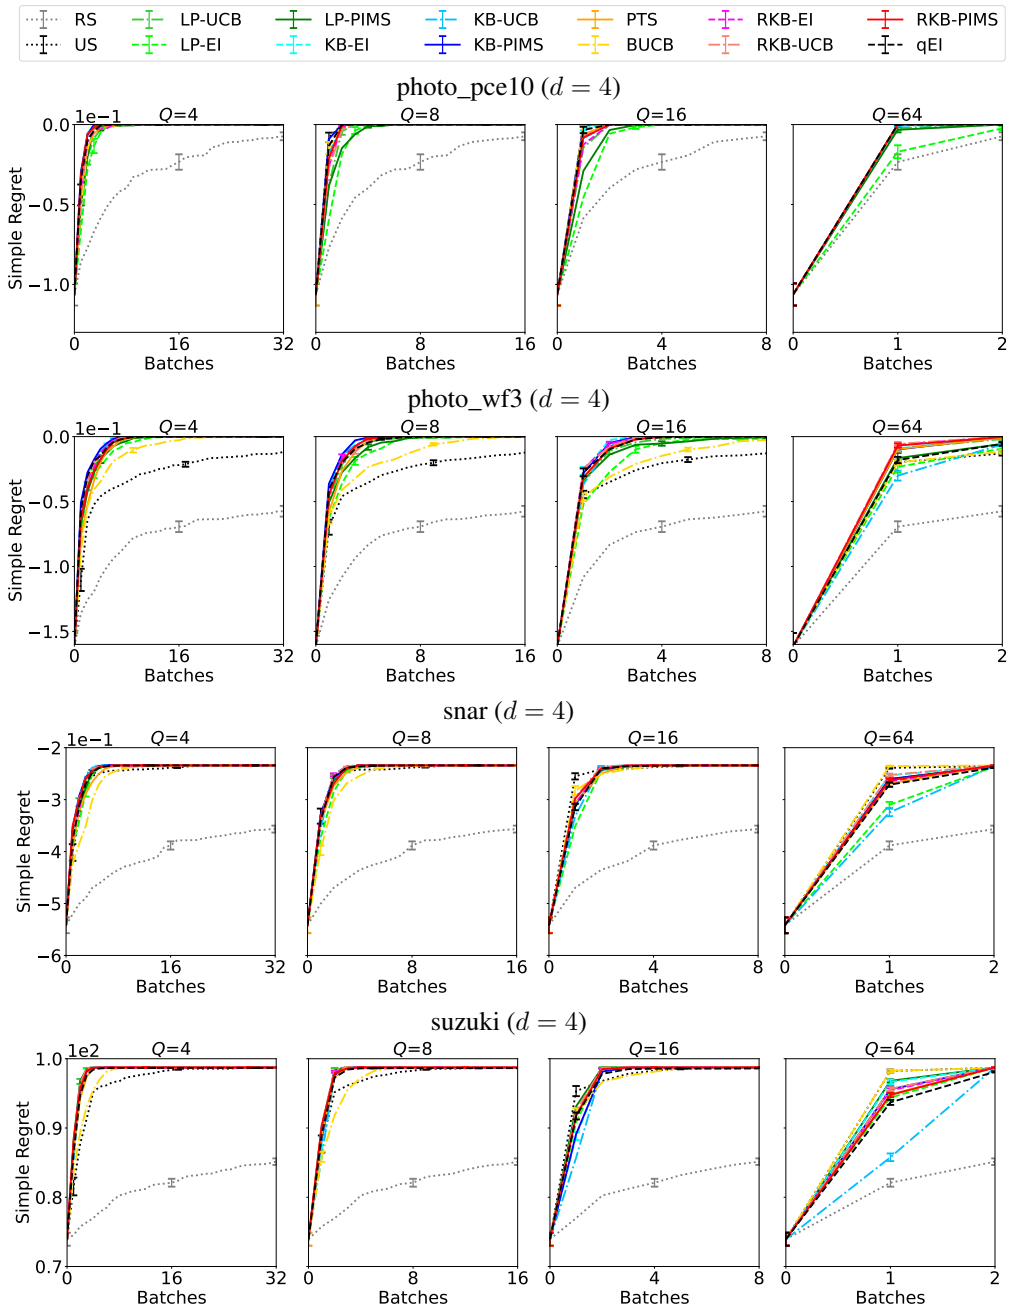


Figure 10: Result of experiments on emulators with synchronous setting. The lines and error bars mean average and standard error of the best objective value $\max_{i \in [t]} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition. One batch corresponds to Q iterations.

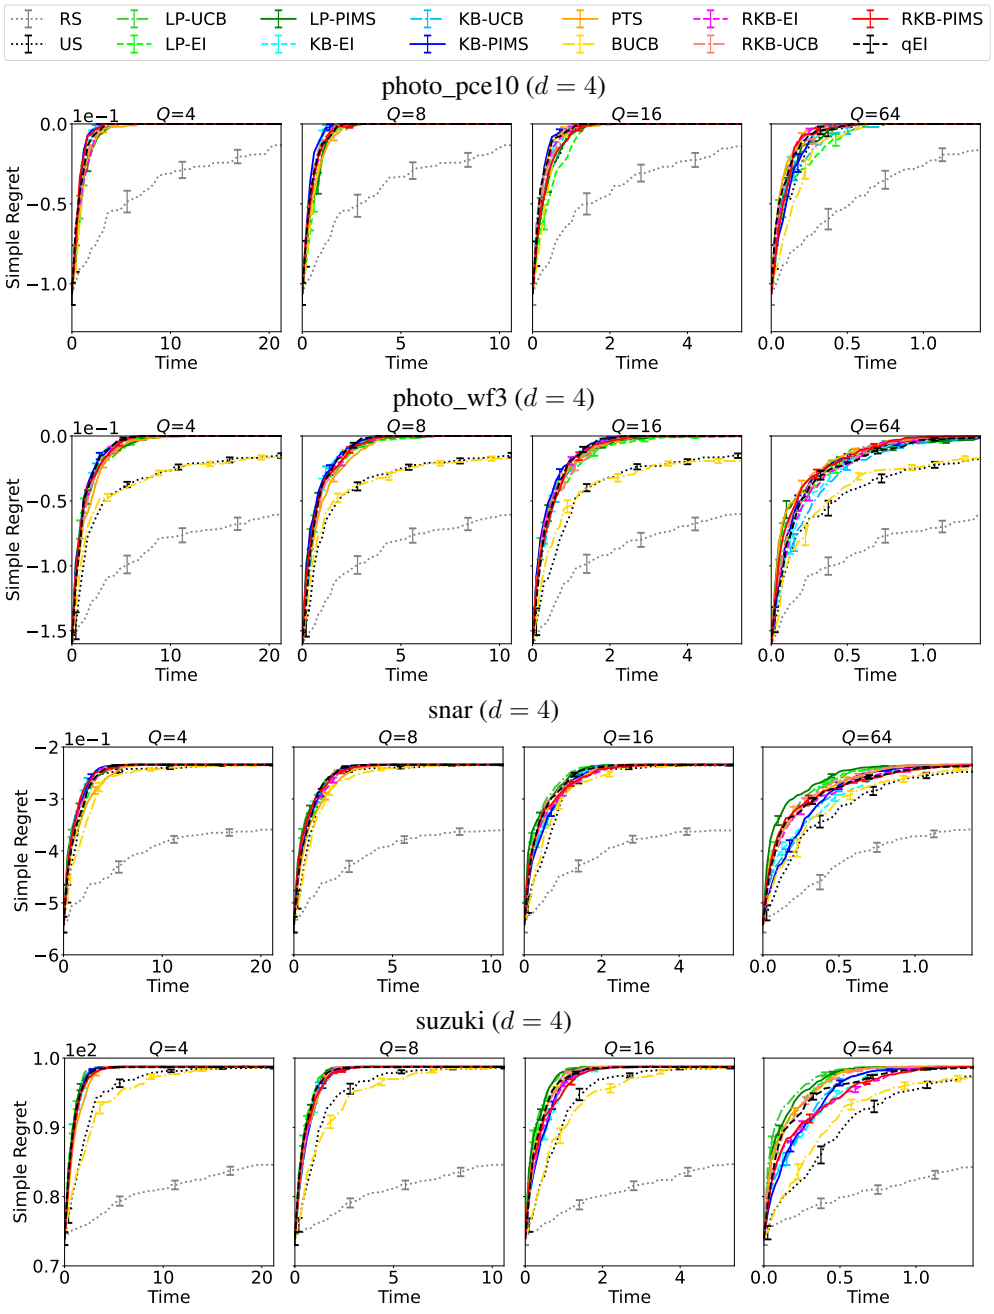


Figure 11: Result of experiments on emulators with asynchronous setting. The lines and error bars mean average and standard error of the best objective value $\max_{i \in \mathcal{N}_{t(\tau)}} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition, where $t(\tau)$ is the last iteration completed by time τ.

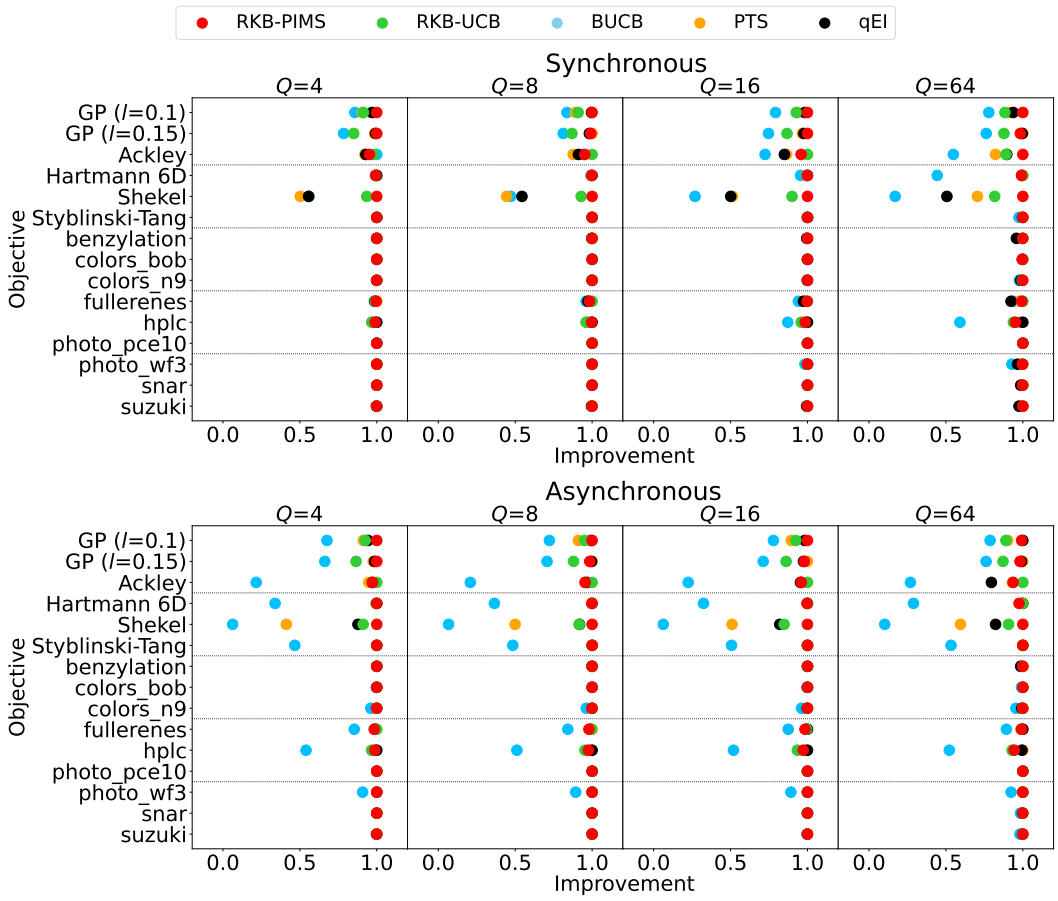


Figure 12: Mean improvement of the best objective value $\max_{i \in [t]} f(\boldsymbol{x}_i)$ across the 64 experiments on each condition. The improvement amount has been normalized so that its maximum value among the five compared methods is one. In asynchronous settings, the improvement is measured up to the end time of each experiment. Methods qLEI and qLNEI are denoted as qEI in the legend.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The paper includes theorems (Section 4) and proofs (Appendix) , as well as experimental results (Section 5 and Appendix), that support its main claims.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The limitation is stated in Section 6 as future work.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: The assumptions of the theorems are in Sections 2 and 4, and the proof is in the appendix.

Guidelines:

- The answer [N/A] means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Details of the experimental setup are described in Section 5.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code used in the experiment, along with instructions on how to use it, will be released as supplemental material.

Guidelines:

- The answer [N/A] means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Details of the experimental setup are described in Section 5.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: The experimental results are analyzed in Section 5 using the mean and standard error.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [No]

   Justification: The experiments do not require any computational resources that warrant special mention.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The paper uses only publicly available code and data.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [No]

    Justification: This research concerns fundamental optimization and will not directly have any negative impact on society.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The research does not deal with data or models that could potentially be misused.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: The paper references the sources of the code and data used in the research.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not introduce any new assets.

    Guidelines:

    - The answer [N/A] means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: The paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.