

Regret Analysis for Randomized Gaussian Process Upper Confidence Bound

SHION TAKENO*, Nagoya University, Japan and RIKEN AIP, Japan

YU INATSU, Nagoya Institute of Technology, Japan

MASAYUKI KARASUYAMA, Nagoya Institute of Technology, Japan

Gaussian process upper confidence bound (GP-UCB) is a theoretically established algorithm for Bayesian optimization (BO), where we assume the objective function f follows a GP. One notable drawback of GP-UCB is that the theoretical confidence parameter β increases along with the iterations and is too large. To alleviate this drawback, this paper analyzes the randomized variant of GP-UCB called improved randomized GP-UCB (IRGP-UCB), which uses the confidence parameter generated from the shifted exponential distribution. We analyze the expected regret and conditional expected regret, where the expectation and the probability are taken respectively with f and noise and with the randomness of the BO algorithm. In both regret analyses, IRGP-UCB achieves a sub-linear regret upper bound without increasing the confidence parameter if the input domain is finite. Furthermore, we show that randomization plays a key role in avoiding an increase in confidence parameter by showing that GP-UCB using a constant confidence parameter can incur linearly growing expected cumulative regret. Finally, we show numerical experiments using synthetic and benchmark functions and real-world emulators.

JAIR Associate Editor: Mikko Koivisto

JAIR Reference Format:

Shion Takeno, Yu Inatsu, and Masayuki Karasuyama. 2025. Regret Analysis for Randomized Gaussian Process Upper Confidence Bound. *Journal of Artificial Intelligence Research* 84, Article 18 (November 2025), 33 pages. DOI: [10.1613/jair.1.19393](https://doi.org/10.1613/jair.1.19393)

1 Introduction

Bayesian optimization (BO) has become a popular black-box optimization method. BO aims to optimize with fewer function evaluations for an objective function, which is expensive to evaluate. For this purpose, BO sequentially evaluates the function value at the recommended input, which is chosen by maximizing an acquisition function based on a surrogate Bayesian model. This paper focuses on BO using a Gaussian process (GP) surrogate model. GP-based BO has been widely used in many applications, such as automatic machine learning (Snoek et al. 2012), materials informatics (Ueno et al. 2016), and drug design (Griffiths and Hernández-Lobato 2020; Korovina et al. 2020).

GP upper confidence bound (GP-UCB) (Kushner 1964; Srinivas et al. 2010) is a seminal work for regret-guaranteed BO. Srinivas et al. show the sub-linear regret bounds of GP-UCB, which has become a fundamental technique for regret analysis in this field. On the other hand, GP-UCB suffers from the tuning of the confidence parameter, which is often written as β_t . The confidence parameter for the theoretical analysis is too large and

*Corresponding Author.

Authors' Contact Information: Shion Takeno, ORCID: [0009-0000-3638-8658](https://orcid.org/0009-0000-3638-8658), takeno.shion.m6@f.mail.nagoya-u.ac.jp, Nagoya University, Nagoya, Aichi, Japan and RIKEN AIP, Nihonbashi, Tokyo, Japan; Yu Inatsu, ORCID: [0000-0001-5655-2558](https://orcid.org/0000-0001-5655-2558), inatsu.yu@nitech.ac.jp, Nagoya Institute of Technology, Nagoya, Aichi, Japan; Masayuki Karasuyama, ORCID: [0000-0002-6177-3686](https://orcid.org/0000-0002-6177-3686), karasuyama@nitech.ac.jp, Nagoya Institute of Technology, Nagoya, Aichi, Japan.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2025 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.19393](https://doi.org/10.1613/jair.1.19393)

needs to increase over time. Hence, although the analysis of GP-UCB is solid and general, there is still an essential gap between theory and practice.

To resolve this problem, Berk et al. (2020) propose the randomized variants of GP-UCB called RGP-UCB, where the confidence parameter follows the Gamma distribution. Although Berk et al. show the superior performance of randomized GP-UCB, their regret analysis appears to contain several technical issues as discussed in Appendix C of the previous conference version of this paper (Takeno et al. 2023). We rectified the analysis and showed that GP-UCB with randomized confidence parameters achieves sub-linear expected regret upper bounds under several conditions, discussed in Section 4.1 and Appendix D of Takeno et al. However, even in this analysis, the confidence parameter is too large and needs to increase over time, as with GP-UCB.

In this paper, we analyze randomized GP-UCB using confidence parameters generated from the shifted exponential distribution, referred to as improved randomized GP-UCB (IRGP-UCB). First, we show sub-linear expected regret upper bounds of IRGP-UCB, where the regret is averaged over all the randomness of the GP, the noise, and the BO algorithm itself. In particular, our analysis reveals that IRGP-UCB can achieve a sub-linear expected regret without increasing the confidence parameter when the input domain is finite. On the other hand, since the randomness of the BO algorithm is averaged, the expected regret analysis is not fair in comparing the expected regret upper bounds of deterministic BO algorithms. That is, depending on the realization of the BO algorithm, the regret averaged by the randomness of GP and the noise may get worse. Therefore, we further analyze the conditional expected regret averaged over the randomness of the GP and the noise and conditioned on the randomness of the BO algorithm. Even in this analysis, IRGP-UCB does not require increasing the confidence parameter and maintains the same order of upper bounds. In addition, although the expectation of the randomized confidence parameters needs to increase, we show that randomized GP-UCB can achieve high-probability regret upper bounds, which implies that randomizing confidence parameters does not diminish the theoretical performance compared with the existing result (Srinivas et al. 2010).

The above expected and conditional expected regret analyses raise the question of whether randomization is necessary for achieving a sub-linear expected regret bound. We show an affirmative answer for this question by showing a problem instance where GP-UCB using a constant confidence parameter incurs linear expected cumulative regret even in finite input domains. Hence, at least for the expected regret analysis, we can see that randomizing the confidence parameters plays a key role in avoiding increasing confidence parameters.

This paper is organized as follows. Section 1.1 reviews the existing regret-guaranteed BO methods in the Bayesian setting, in which the objective function f follows the GP. Section 2 provides the technical background. Section 3 shows the algorithm of the proposed method. Section 4 shows the analyses for the expected regret upper bound of IRGP-UCB (Section 4.1), the conditional expected regret upper bound of IRGP-UCB (Section 4.2), the high-probability regret upper bound of randomized GP-UCB (Section 4.3), and the expected regret lower bound of GP-UCB with a constant confidence parameter on a specific problem instance (Section 4.4). Section 5 shows the wide range of numerical and real-world dataset experiments. Finally, Section 6 concludes.

Compared with the previous conference version of this paper (Takeno et al. 2023), we add the following analyses in addition to reorganizing the entire paper. First, we add the conditional expected regret analysis of IRGP-UCB in Section 4.2. Second, we show that randomized GP-UCB with increasing confidence parameters can achieve high-probability sub-linear cumulative regret bounds in Section 4.3. Third, in Section 4.4, we show that GP-UCB using a constant confidence parameter can incur linear expected cumulative regret. In addition, we have removed an unnecessary assumption, the kernel's stationarity, from theorems and reorganized Lemma 4.2 from Lemma 4.2 of Takeno et al. to improve the lemma's usefulness for the proof. Finally, we have added one baseline method called the probability of improvement from the maximum of sample path (Takeno et al. 2024) in experiments.

1.1 Related Work

The theoretical assumptions for BO are mainly twofold: the Bayesian setting (Kandasamy et al. 2018; Russo and Van Roy 2014; Srinivas et al. 2010), where the objective function f is assumed to be a sample path from a GP, and the frequentist setting (Chowdhury and Gopalan 2017; Janz et al. 2020; Srinivas et al. 2010), where the objective function f is an element of a reproducing kernel Hilbert space (RKHS). Hence, we can construct credible intervals directly in the Bayesian setting since f follows the posterior distribution. On the other hand, although f does not follow the posterior distribution in the frequentist setting, confidence intervals for f based on GP models have been developed (Chowdhury and Gopalan 2017; Srinivas et al. 2010). Based on these tools, we can design BO algorithms that have regret guarantees using GP models. Note that since a sample path from a GP does not belong to the RKHS with bounded norm, neither analysis subsumes the other, as discussed in the last paragraph of Section 4 of Srinivas et al. This study considers BO in the Bayesian setting.

Various BO methods have been developed in the literature, for example, expected improvement (EI) (Mockus et al. 1978), entropy search (ES) (Hennig and Schuler 2012), and predictive entropy search (PES) (Hernández-Lobato et al. 2014). The theoretical analysis of EI for the noiseless case, where function values are not contaminated by noise, has been performed in the existing works (Bull 2011; Grünewälder et al. 2010; Vazquez and Bect 2010) for both Bayesian and frequentist settings. However, the regret analysis of EI for the noisy case is still limited to the frequentist setting (Tran-The et al. 2022; Wang and Freitas 2014). Further, although the practical performance of ES and PES has been shown repeatedly, their regret analysis is an open problem.

GP-UCB is one of the prominent studies for theoretically guaranteed BO methods. Srinivas et al. (2010) show the high-probability bound for cumulative regret. Furthermore, GP-UCB achieves the sub-linear expected regret upper bounds as shown in Theorem 2.1 and B.1 of Takeno et al. (2023). However, the confidence parameter must be scheduled as $\beta_t \propto \log t$ in both regret analyses. Although many studies (Chowdhury and Gopalan 2017; Janz et al. 2020; Srinivas et al. 2010) considered the frequentist setting, this study concentrates on the Bayesian setting. On the other hand, truncated variance reduction (TRUVAR) (Bogunovic et al. 2016) achieves high-probability regret bounds. TRUVAR is based on the variance reduction in potential maximizers determined by UCB and lower confidence bounds. Hence, TRUVAR may be improved using randomization similar to GP-UCB.

Another well-known regret-guaranteed BO method is Thompson sampling (TS) (Kandasamy et al. 2018; Russo and Van Roy 2014; Takeno et al. 2024). TS achieves sub-linear BCR bounds by sequentially observing the optimal point of the GP's posterior sample path. Although TS does not require any hyperparameter, TS often deteriorates by over-exploration, as discussed in the existing study (Shahriari et al. 2016). Recently, another posterior sampling-based BO called probability improvement from the maximum of sample path (PIMS) has been proposed. Known-best expected regret analyses of TS and PIMS are based on the proof technique of our previous conference version of this paper (Takeno et al. 2023).

Wang et al. (2016) show regret analysis of the GP estimation (GP-EST) algorithm, which can be interpreted as GP-UCB with the confidence parameter defined using \hat{m} , an estimator of $\mathbb{E}[\max f(\mathbf{x})]$. Their analysis requires an assumption $\hat{m} \geq \mathbb{E}[\max f(\mathbf{x})]$, whose sufficient condition is provided in Corollary 3.5 of Wang et al. However, this sufficient condition does not typically hold, as discussed immediately after the corollary in Wang et al. Furthermore, the final convergence rate contains \hat{m} itself, whose convergence rate is not clarified. In contrast, our Lemma 4.2 shows the bound for $\mathbb{E}[\max f(\mathbf{x})]$ under common regularity conditions. Wang and Jegelka (2017) show regret analysis of max-value entropy search (MES). However, it is pointed out that their proof contains several technical issues (Takeno, Tamura, et al. 2022).

Scarlett (2018) shows that, under several conditions, there exists an algorithm that achieves a tighter high-probability regret upper bound compared with other well-known algorithms, such as GP-UCB and TS. However, the algorithm shown by Scarlett depends on the increasing confidence parameters for the construction of

confidence bounds. Therefore, their algorithm also has the same gap between theoretical analysis and practical effectiveness as GP-UCB.

Another relevant line of work is level set estimation (LSE) (Gotovos et al. 2013). LSE considers the classification problem of the expensive-to-evaluate black-box functions. After our preceding conference version of this paper (Takeno et al. 2023), a randomized version of LSE has been proposed by Inatsu, Takeno, Kutsukake, et al. (2024). Even for LSE, the randomization avoids increasing the confidence parameter in UCB-based LSE methods.

2 Background

This section presents an overview of BO, our theoretical assumptions, evaluation measures, and the GP-UCB algorithm.

2.1 Bayesian Optimization

We consider an optimization problem

$$\mathbf{x}^* = \arg \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}),$$

where f is a black-box function, $\mathcal{X} \subset \mathbb{R}^d$ is an input domain, and d is an input dimension. Here, black-box means that we can access only function values and cannot obtain other information, such as gradients. Furthermore, we assume that observations are contaminated by noise. In addition, we consider the case that f is expensive to evaluate. Therefore, BO sequentially observes the noisy function value with some policy so that the number of function evaluations is minimized as far as possible. This policy is determined by an acquisition function (AF) $a : \mathcal{X} \rightarrow \mathbb{R}$, which quantifies the utility of function evaluation. That is, at each iteration t , BO queries $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} a(\mathbf{x})$ and obtains $y_t = f(\mathbf{x}_t) + \epsilon_t$, where ϵ_t is additive noise. Regularity assumptions of f and ϵ_t are described in Section 2.2.

BO generally employs the GP surrogate model (Rasmussen and Williams 2005) due to its flexible prediction capability and analytical computations. The GP surrogate model can be derived under assumptions that f is a sample path from a GP with a zero mean and a kernel function $k : \mathcal{X} \times \mathcal{X} \mapsto \mathbb{R}$, and $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ is i.i.d. Gaussian noise with a positive variance $\sigma^2 > 0$. At each iteration t , a dataset $\mathcal{D}_{t-1} := \{(\mathbf{x}_i, y_i)\}_{i=1}^{t-1}$ is already obtained from the nature of BO. Then, the posterior distribution $p(f | \mathcal{D}_{t-1})$ is a GP again. The posterior mean and variance of $f(\mathbf{x})$ are derived as follows:

$$\begin{aligned} \mu_{t-1}(\mathbf{x}) &= \mathbf{k}_{t-1}(\mathbf{x})^\top (\mathbf{K} + \sigma^2 \mathbf{I}_{t-1})^{-1} \mathbf{y}_{t-1}, \\ \sigma_{t-1}^2(\mathbf{x}) &= k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_{t-1}(\mathbf{x})^\top (\mathbf{K} + \sigma^2 \mathbf{I}_{t-1})^{-1} \mathbf{k}_{t-1}(\mathbf{x}), \end{aligned}$$

where $\mathbf{k}_{t-1}(\mathbf{x}) := (k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_{t-1}))^\top \in \mathbb{R}^{t-1}$, $\mathbf{K} \in \mathbb{R}^{(t-1) \times (t-1)}$ is the kernel matrix whose (i, j) -element is $k(\mathbf{x}_i, \mathbf{x}_j)$, $\mathbf{I}_{t-1} \in \mathbb{R}^{(t-1) \times (t-1)}$ is the identity matrix, and $\mathbf{y}_{t-1} := (y_1, \dots, y_{t-1})^\top \in \mathbb{R}^{t-1}$. Hereafter, we denote that the probability density function (PDF) $p(\cdot | \mathcal{D}_{t-1}) = p_t(\cdot)$, the probability $\Pr(\cdot | \mathcal{D}_{t-1}) = \Pr_t(\cdot)$, and the expectation $\mathbb{E}[\cdot | \mathcal{D}_{t-1}] = \mathbb{E}_t[\cdot]$ for brevity.

2.2 Regularity Assumptions

We assume that the GP surrogate model is correct. That is, $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ is i.i.d. Gaussian noise with a positive variance $\sigma^2 > 0$, and the underlying true function f is a sample path from a known GP with a zero mean and a kernel function k . Here, the zero mean is without loss of generality (Rasmussen and Williams 2005). In addition, we assume that the kernel function k is normalized. Therefore, the kernel $k(\mathbf{x}, \mathbf{x}')$ satisfies the condition $k(\mathbf{x}, \mathbf{x}') \leq 1$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$.

Furthermore, when \mathcal{X} is continuous, we consider the following smoothness assumption:

Assumption 2.1. Let $\mathcal{X} \subset [0, r]^d$ be a compact and convex set, where $r > 0$. Assume that the kernel k satisfies the following condition on the derivatives of a sample path f . There exists the constants $a, b > 0$ such that,

$$\Pr \left(\sup_{\mathbf{x} \in \mathcal{X}} \left| \frac{\partial f}{\partial x_j} \right| > L \right) \leq a \exp \left(- \left(\frac{L}{b} \right)^2 \right), \text{ for } j \in [d],$$

where $[d] = \{1, \dots, d\}$.

This assumption holds for stationary (Genton 2001) and four times differentiable kernels, such as RBF kernel and Matérn kernels with a parameter ν , from Theorem 5 of Ghosal and Roy (2006) as discussed in Section 4 of Srinivas et al. (2010). Therefore, this assumption is commonly used, for example, in Srinivas et al., Kandasamy et al. (2018), Paria et al. (2020), and Takeno et al. (2024).

2.3 Regret

In this paper, we evaluate the performance of BO methods by *cumulative regret* (Kandasamy et al. 2018; Paria et al. 2020; Russo and Van Roy 2014; Srinivas et al. 2010; Takeno et al. 2024). We incur the instantaneous regret $r_t = f(\mathbf{x}^*) - f(\mathbf{x}_t)$ for each function evaluation. The cumulative regret for T iterations is defined as the sum of instantaneous regrets over T function evaluations, defined as follows:

$$R_T = \sum_{t=1}^T r_t.$$

In particular, BO methods that achieve sub-linear cumulative regret (i.e., such that $\lim_{T \rightarrow \infty} R_T/T = 0$) are called *no-regret*. No-regret property also indicates that the best regret $\min_{t \in [T]} r_t$ converges to 0 since $\min_{t \in [T]} r_t < R_T/T$. Therefore, the main goal of theoretical analysis is to show sub-linearity of cumulative regret. Note that regret is a random quantity due to f , $\{\epsilon_t\}_{t \geq 1}$, and the randomness of the BO algorithm.

Maximum information gain. : The convergence rates of regret bounds are characterized by *maximum information gain* (MIG) (Srinivas et al. 2010; Vakili et al. 2021). MIG γ_T is defined as follows:

Definition 2.1 (Maximum information gain). Let $f \sim \mathcal{GP}(0, k)$. Let $A = \{\mathbf{a}_i\}_{i=1}^T \subset \mathcal{X}$. Let $\mathbf{f}_A = (f(\mathbf{a}_i))_{i=1}^T$, $\epsilon_A = (\epsilon_i)_{i=1}^T$, where $\forall i, \epsilon_i \sim \mathcal{N}(0, \sigma^2)$, and $\mathbf{y}_A = \mathbf{f}_A + \epsilon_A \in \mathbb{R}^T$. Then, MIG γ_T is defined as follows:

$$\gamma_T := \max_{A \subset \mathcal{X}; |A|=T} I(\mathbf{y}_A; \mathbf{f}_A),$$

where I is the Shannon mutual information.

MIG is known to be sub-linear for commonly used kernel functions, for example, $\gamma_T = O((\log T)^{d+1})$ for RBF kernels and $\gamma_T = O(T^{\frac{d}{2\nu+d}} (\log T)^{\frac{2\nu}{2\nu+d}})$ for Matérn- ν kernels (Srinivas et al. 2010; Vakili et al. 2021).

2.4 GP-UCB

GP-UCB selects the next input by maximizing UCB as follows:

$$\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}).$$

In the existing regret analyses (Srinivas et al. 2010; Takeno et al. 2023), to obtain the sub-linear cumulative regret bounds, the confidence parameter β_t must be scheduled as $\beta_t = \Theta(\log(|\mathcal{X}|t))$ for finite input domains.

Algorithm 1 IRGP-UCB

Require: Input space \mathcal{X} , Parameters $\{s_t\}_{t \geq 1}$ and λ for ζ_t , GP prior $\mu = 0$ and k

- 1: $\mathcal{D}_0 \leftarrow \emptyset$
 - 2: **for** $t = 1, \dots$ **do**
 - 3: Fit GP to \mathcal{D}_{t-1}
 - 4: Generate ζ_t by Eq. (1)
 - 5: $\mathbf{x}_t \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x})$
 - 6: Observe $y_t = f(\mathbf{x}_t) + \epsilon_t$ and $\mathcal{D}_t \leftarrow \mathcal{D}_{t-1} \cup (\mathbf{x}_t, y_t)$
 - 7: **end for**
-

3 Algorithm

Algorithm 1 shows the pseudo-code of IRGP-UCB. The difference from the original GP-UCB is that $\{\zeta_t\}$ is a sequence of random variables, not a monotonically increasing (non-random) sequence. Using the generated ζ_t , the next input is chosen as follows:

$$\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}).$$

Our choice of the distribution of ζ_t is a *shifted exponential distribution* (for example, Beg 1980; Lam et al. 1994), whose PDF is written as follows:

$$p(\zeta; s_t, \lambda) = \begin{cases} \lambda \exp(-\lambda(\zeta - s_t)) & \text{if } \zeta \geq s_t, \\ 0 & \text{otherwise,} \end{cases}$$

where λ is a rate parameter. This can be seen as the distribution of the sum of s_t and $Z \sim \text{Exp}(\lambda)$. Thus, sampling from the shifted exponential distribution can be easily performed using an exponential distribution as follows:

$$\zeta_t \leftarrow s_t + Z_t, \text{ where } Z_t \sim \text{Exp}(\lambda). \quad (1)$$

The theoretical choice of $\{s_t\}_{t \geq 1}$ and λ will be shown in Section 4.

4 Regret Analysis

In this section, we analyze the performance of IRGP-UCB and the necessity of randomizing confidence parameters. First, we analyze the average-case performance through expected regret averaged over all the randomness of f , $\{\epsilon_t\}_{t \geq 1}$, and $\{\zeta_t\}_{t \geq 1}$. Second, we analyze conditional expected regret averaged over f and $\{\epsilon_t\}_{t \geq 1}$ and conditioned on $\{\zeta_t\}_{t \geq 1}$. Third, we derive high-probability regret bounds of randomized GP-UCB with increasing confidence parameters. Lastly, we show that GP-UCB using a constant confidence parameter can incur linearly growing expected cumulative regret.

4.1 Expected Regret Bounds

Here, we derive the expected cumulative regret bound, which is also called Bayesian cumulative regret (BCR) (Russo and Van Roy 2014). The definition of BCR is

$$\text{BCR}_T := \mathbb{E} \left[\sum_{t=1}^T r_t \right], \quad (2)$$

where the expectation is taken with all the randomness, that is, f , $\{\epsilon_t\}_{t \geq 1}$, and $\{\zeta_t\}_{t \geq 1}$.

First, for completeness, we provide a modified version of the basic lemma derived in Srinivas et al. (2010):

Lemma 4.1. *Suppose that f is a sample path from a GP with zero mean and a positive semidefinite kernel k , and \mathcal{X} is finite. Pick $\delta \in (0, 1)$ and $t \geq 1$. Then, for any given \mathcal{D}_{t-1} ,*

$$\Pr_t \left(f(\mathbf{x}) \leq \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X} \right) \geq 1 - \delta,$$

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

See Appendix A.1 for the proof. This lemma differs slightly from Lemma 5.1 of Srinivas et al. (2010), since, in Lemma 4.1, the iteration t is fixed, and β_δ does not depend on t .

Based on Lemma 4.1, we show the following lemma, which will play a key role in the subsequent analyses:

Lemma 4.2. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel, and \mathcal{X} be finite. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|\mathcal{X}|/2)$ and $\lambda = 1/2$. Then, for any given \mathcal{D}_{t-1} , the following inequality holds:*

$$\mathbb{E}_t [f(\mathbf{x}^*)] \leq \mathbb{E}_t \left[\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}) \right],$$

for all $t \geq 1$.

PROOF. We here show a short proof of Lemma 4.2, although a detailed proof is shown in Appendix A.2. From Lemma 4.1, we can obtain the following inequality:

$$F_t^{-1}(1 - \delta) \leq \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}),$$

where $F_t(\cdot) := \Pr_t(f(\mathbf{x}_*) < \cdot)$ is a cumulative distribution function of $f(\mathbf{x}_*)$, and F_t^{-1} is its inverse function. Then, substituting $U \sim \text{Uni}(0, 1)$ into δ and taking the expectation, we obtain the following inequality:

$$\mathbb{E}_t [f(\mathbf{x}_*)] \leq \mathbb{E}_U \left[\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_U^{1/2} \sigma_{t-1}(\mathbf{x}) \right],$$

which can be derived in a similar way to inverse transform sampling. Hence, $\beta_U^{1/2} = 2 \log(|\mathcal{X}|/2) - 2 \log(U)$ results in a random variable, which follows the shifted exponential distribution. Consequently, we can obtain the desired bound. \square

Lemma 4.2 shows a direct upper bound of the expectation of $f(\mathbf{x}^*)$. A notable consequence of Lemma 4.2 is an upper bound of BCR as follows:

$$\begin{aligned} \text{BCR}_T &= \sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - v_t(\mathbf{x}_t)] + \mathbb{E} [v_t(\mathbf{x}_t) - f(\mathbf{x}_t)] \\ &= \sum_{t=1}^T \mathbb{E} [\mathbb{E}_t [f(\mathbf{x}^*) - v_t(\mathbf{x}_t)]] + \mathbb{E} [v_t(\mathbf{x}_t) - f(\mathbf{x}_t)] \\ &\leq \sum_{t=1}^T \mathbb{E} [v_t(\mathbf{x}_t) - f(\mathbf{x}_t)], \end{aligned}$$

where $v_t(\mathbf{x}) := \mu_{t-1}(\mathbf{x}) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x})$. The above inequality is obtained by the selection rule of IRGP-UCB. Importantly, $\mathbb{E}[\zeta_t] = 2 + s_t$ is a constant since $s_t = 2 \log(|\mathcal{X}|/2)$ does not depend on t in Lemma 4.2. Thus, using Lemma 4.2, we can obtain the BCR bound for finite input domains:

Theorem 4.1 (BCR bound for finite domain). *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|\mathcal{X}|/2)$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , BCR can be bounded as follows:*

$$\text{BCR}_T \leq \sqrt{C_1 C_2 T \gamma_T},$$

where $C_1 := 2/\log(1 + \sigma^{-2})$ and $C_2 := 2 + 2 \log(|\mathcal{X}|/2)$.

See Appendix A.3 for detailed proof.

Theorem 4.1 has an important implication that IRGP-UCB does not need to schedule the parameters s_t and λ of ζ_t , in contrast to GP-UCB and RGP-UCB. Therefore, through randomization, IRGP-UCB essentially alleviates the problem that the well-known $\beta_t \propto \log(t)$ strategy results in a practically too large confidence parameter. Further, over-exploration in the later iterations can be avoided. IRGP-UCB is the first GP-UCB-based method that enjoys the above practical benefits without sacrificing regret guarantee.

On the other hand, note that the $\sqrt{\log |\mathcal{X}|}$ dependence remains as with the prior works (Kandasamy et al. 2018; Paria et al. 2020; Russo and Van Roy 2014; Srinivas et al. 2010). In BO, we are usually interested in the case that the total number of iterations $T < |\mathcal{X}|$. Therefore, in such case, although our regret upper bound $O(\sqrt{\gamma_T T \log |\mathcal{X}|})$ is slightly better than the existing one $O(\sqrt{\gamma_T T \log(|\mathcal{X}|T)})$, the difference between those regret upper bounds is only the constant term.

Next, we show the BCR bound for the infinite (continuous) domain:

Theorem 4.2 (BCR bound for infinite domain). *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel, $k(\mathbf{x}, \mathbf{x}) \leq 1$, and Assumption 2.1 holds. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2d \log(bdrt^2(\sqrt{\log(ad)} + \sqrt{\pi}/2)) - 2 \log 2$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB, BCR can be bounded as follows:*

$$\text{BCR}_T \leq \frac{\pi^2}{6} + \sqrt{C_1 T \gamma_T (2 + s_T)},$$

where $C_1 := 2/\log(1 + \sigma^{-2})$.

See Appendix A.4 for the proof. Unfortunately, in this case, $\mathbb{E}[\zeta_t] = O(\log t)$ and the resulting BCR bound is $O(\sqrt{T \gamma_T \log T})$. This is because the discretization of input domain \mathcal{X}_t , updated as $|\mathcal{X}_t| = O(t^{2d})$, is required to obtain the BCR bound for the infinite domain. Since we cannot avoid the dependence on $|\mathcal{X}|$ also in Theorem 4.1, the resulting BCR bound in Theorem 4.2 requires the term $\log(|\mathcal{X}_t|) = O(2d \log(t))$.

4.2 Conditional Expected Regret Bounds

In this section, we analyze the conditional expected regret of IRGP-UCB. This analysis is newly provided compared with the previous conference version of our paper (Takeno et al. 2023). The regret depends on the randomness of f , $\{\epsilon_t\}_{t \geq 1}$, and $\{\zeta_t\}_{t \geq 1}$. Although the randomness of f and $\{\epsilon_t\}_{t \geq 1}$ arise from the problem setup, $\{\zeta_t\}_{t \geq 1}$ is produced by the user. Hence, distinguishing the randomness, we analyze the following conditional expected regret:

$$\mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} [f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1}],$$

where the expectation is taken with f and $\{\epsilon_t\}_{t \geq 1}$. Note that $\mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} [f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1}]$ is a random variable depending on $\{\zeta_t\}_{t \geq 1}$. Hence, showing the sub-linear high-probability upper bound of the conditional expected regret implies that the BO algorithm achieves a good performance, averaged over a random environment, with a high-probability concerning the randomness of the algorithm.

Let \mathcal{P} be a set of random variables included in the problem setup ($\mathcal{P} = \{f, \{\epsilon_t\}_{t \geq 1}\}$) and \mathcal{A} be a set of random variables included in the algorithm ($\mathcal{A} = \{\zeta_t\}_{t \geq 1}$ in the case of IRGP-UCB). In the usual expected regret

analysis of the deterministic BO algorithm, such as GP-UCB (for example, Paria et al. 2020; Takeno et al. 2023), the expectation over the problem setup $\mathbb{E}_{\mathcal{P}}[R_T]$ is considered. When we interpret the confidence parameter in GP-UCB as a random variable equal to β_t with probability 1 ($\mathcal{A} = \{\beta_t\}_{t \geq 1}$), $\mathbb{E}_{\mathcal{P}}[R_T]$ is equivalent to the joint expectation $\mathbb{E}_{\mathcal{P}, \mathcal{A}}[R_T]$. Further, we obviously see that the conditional expected regret is also equivalent: $\mathbb{E}_{\mathcal{P}, \mathcal{A}}[R_T] = \mathbb{E}_{\mathcal{P}}[R_T | \mathcal{A}]$. On the other hand, since IRGP-UCB is stochastic, possible definitions of the expected regret are $\mathbb{E}_{\mathcal{P}, \mathcal{A}}[R_T]$ and $\mathbb{E}_{\mathcal{P}}[R_T | \mathcal{A}]$, but these two regrets can be different. Theorems 4.1 and 4.2 show the theoretical rate of $\mathbb{E}_{\mathcal{P}, \mathcal{A}}[R_T]$. However, these theorems do not guarantee anything about the conditional expected regret $\mathbb{E}_{\mathcal{P}}[R_T | \mathcal{A}]$. The following remark provides a motivation for deriving a conditional expected regret bound:

Remark 4.1. Although $\mathbb{E}_{\mathcal{P}, \mathcal{A}}[R_T]$ derived by Theorems 4.1 and 4.2 reveal average behavior of the regret of IRGP-UCB, it obscures variability caused by $\{\zeta_t\}_{t \geq 1}$. Hence, comparing $\mathbb{E}_{\mathcal{P}}[R_T]$ of GP-UCB and $\mathbb{E}_{\mathcal{P}, \mathcal{A}}[R_T]$ of IRGP-UCB can be unfair. On the other hand, $\mathbb{E}_{\mathcal{P}}[R_T | \mathcal{A}]$ is a random variable, whose randomness stems from $\{\zeta_t\}_{t \geq 1}$. Therefore, by showing a high-probability bound of $\mathbb{E}_{\mathcal{P}}[R_T | \mathcal{A}]$, we can 1) fairly compare the regret of IRGP-UCB with $\mathbb{E}_{\mathcal{P}}[R_T]$ of GP-UCB in the sense that the expectation is taken over the same random variables \mathcal{P} , and 2) demonstrate that the expected regret of IRGP-UCB can be stably bounded with respect to the random behavior of the algorithm.

First, we analyze the case of finite \mathcal{X} . We show the following theorem:

Theorem 4.3. Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|\mathcal{X}|/2)$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , the conditional expected regret can be bounded with probability at least $1 - \delta$ as follows:

$$\Pr \left(\forall T \geq 1, \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right] \leq U(T, \delta) \right) \geq 1 - \delta,$$

where

$$U(T, \delta) = 6\sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right)} + \sqrt{C_1 \gamma_T \left(Ts_T + T + 2\sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right)} + 2 \log \left(\frac{\pi^2 T^2}{3\delta} \right) \right)}$$

and $C_1 := 2/\log(1 + \sigma^{-2})$.

PROOF. We show the proof sketch here only. See Appendix B.1 for the details.

The first key point of the proof is transforming the conditional expected regret using Lemma 4.2 and the properties of the expectation as follows:

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} [f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1}] \\ &= \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_{f | \mathcal{D}_{t-1}} [f(\mathbf{x}^*)] - \mu_{t-1}(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}] \\ &\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_{\zeta_t | \mathcal{D}_{t-1}} [v_t(\mathbf{x}_t)] - v_t(\mathbf{x}_t) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}] \quad (\because \text{Lemma 4.2}) \\ &= \underbrace{\sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \{\zeta_i\}_{i < t}] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}]}_{A_1} + \underbrace{\mathbb{E} \left[\sum_{t=1}^T \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \mid \{\zeta_t\}_{t > 1} \right]}_{A_2}, \end{aligned}$$

where $v_t(\mathbf{x}_t) = \mu_{t-1}(\mathbf{x}_t) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t)$. Then, A_2 can be bounded from above using MIG and the high-probability upper bound of chi-square random variables (Laurent and Massart 2000) shown in Lemma E.4.

The second technical key point is bounding A_1 by Azuma's inequality (Azuma 1967), which requires the conditions that A_1 is a sum of the martingale difference sequence with conditional subgaussian property (See Lemma B.1 for details). We show that A_1 satisfies this condition using several well-known lemmas shown in Appendix B.1. Finally, we can obtain the desired result by applying the union bound for the upper bounds of A_1 and A_2 for all $T \geq 1$. \square

Theorem 4.3 shows the same benefits as Theorem 4.1. First, under the above mild condition $\min\{s_T^2/2, s_T\}T > \log(\pi^2 T^2 / (3\delta))$, the dominant term in $U(T, \delta)$ is $\sqrt{2C_1 T \gamma_T s_T}$, which is the same as Theorem 4.1. Therefore, Theorem 4.3 maintains the rate $O(\sqrt{T \gamma_T \log(|\mathcal{X}|)})$ under mild condition. Second, the distribution of ζ_t is the same as in Theorem 4.1. Hence, increasing the confidence parameter is also avoided even if we consider the high-probability bound regarding the randomness of the algorithm.

Next, as with the expected regret bounds, we extend Theorem 4.3 to the continuous case:

Theorem 4.4. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and Assumption 2.1 holds. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2d \log(bdrt^2(\sqrt{\log(ad)} + \sqrt{\pi}/2)) - 2 \log 2$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , the conditional expected regret can be bounded with probability at least $1 - \delta$ as follows:*

$$\Pr \left(\forall T \geq 1, \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right] \leq U(T, \delta) \right) \geq 1 - \delta,$$

where

$$U(T, \delta) = \frac{\pi^2}{6} + 6\sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right)} + \sqrt{C_1 \gamma_T \left(T s_T + T + 2\sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right)} + 2 \log \left(\frac{\pi^2 T^2}{3\delta} \right) \right)},$$

and $C_1 := 2/\log(1 + \sigma^{-2})$.

See Appendix B.2 for the proof. The resulting upper bound maintains the same rate as the expected regret bound shown in Theorem 4.2 if $\min\{s_T^2/2, s_T\}T > 2 \log(\pi^2 T^2 / (3\delta))$.

4.3 High-Probability Regret Bounds

In this section, we consider the high-probability regret bounds of the randomized GP-UCB, which is newly added compared with the previous conference version of this paper. For simplicity, we focus on the case of the finite \mathcal{X} . As discussed in Russo and Van Roy (2014), a direct consequence of Markov's inequality is

$$\Pr(R_T \leq \text{BCR}_T / \delta) \geq 1 - \delta,$$

for any $\delta \in (0, 1)$. Therefore, the following proposition can be easily obtained:

Proposition 4.1. *Assume the same condition as in Theorem 4.1. Then, by running IRGP-UCB, for some fixed $T \geq 1$, the following holds:*

$$\Pr \left(R_T \leq \sqrt{C_1 C_2 T \gamma_T} / \delta \right) \geq 1 - \delta,$$

where $C_1 := 2/\log(1 + \sigma^{-2})$ and $C_2 := 2 + 2 \log(|\mathcal{X}|/2)$.

However, this proposition shows the regret bound for some fixed $T \geq 1$. Ideally, an anytime algorithm that guarantees that the sub-linear regret bound holds for all $T \geq 1$ with high-probability is desired. For example, the analysis of GP-UCB (Srinivas et al. 2010) is anytime and shows regret upper bound $O(\sqrt{T\gamma_T \log(T|\mathcal{X}|/\delta)})$. In addition, the resulting rate $O(1/\delta)$ regarding δ is worse than that of GP-UCB $O(\log(1/\delta))$. Therefore, we focus on obtaining an anytime regret upper bound with the $\log(1/\delta)$ dependence regarding δ .

First, we consider high-probability bounds using the randomized confidence parameter. Since the randomized confidence parameter ζ_t is scheduled as $\mathbb{E}[\zeta_t] = \Theta(\log t)$, we can derive the following lemma as with Lemma 5.1 in Srinivas et al. (2010):

Lemma 4.3. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel, and \mathcal{X} be finite. Pick $\delta \in (0, 1)$. Assume that $\zeta_t = 2 \log(1/U_t)$, where $\{U_t\}_{t \geq 1}$ is a sequence of mutually independent random variables that satisfy $\mathbb{E}[U_t] = \frac{6\delta}{|\mathcal{X}|t^2\pi^2}$ and $\Pr(U_t \in (0, 1)) = 1$ for all $t \geq 1$. Then, the following inequality holds:*

$$\Pr\left(\forall t \geq 1, \forall \mathbf{x} \in \mathcal{X}, |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x})\right) \geq 1 - \delta.$$

PROOF. As with Lemma 5.1 in Srinivas et al. (2010), for any $t \geq 1$, $\mathbf{x} \in \mathcal{X}$, \mathcal{D}_{t-1} , and U_t , we obtain

$$\Pr\left(|f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| > \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}) \mid U_t, \mathcal{D}_{t-1}\right) \leq U_t.$$

Taking the expectation with respect to $\{U_t\}_{t \geq 1}$ and \mathcal{D}_{t-1} , we obtain

$$\Pr\left(|f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| > \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x})\right) \leq \mathbb{E}[U_t].$$

Then, by applying the union bound for all $t \geq 1$ and $\mathbf{x} \in \mathcal{X}$, we obtain

$$\Pr\left(\forall t \geq 1, \forall \mathbf{x} \in \mathcal{X}, |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x})\right) \geq 1 - |\mathcal{X}| \sum_{t=1}^T \mathbb{E}[U_t] \geq 1 - \delta. \quad \square$$

This analysis can be seen as a generalized version of the analysis of GP-UCB since the randomized confidence parameter is equivalent to GP-UCB if we set $U_t = \frac{6\delta}{|\mathcal{X}|t^2\pi^2}$ almost surely.

Then, by combining Lemma 4.3 and existing analysis techniques, we can obtain the following theorems:

Theorem 4.5. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Pick $\delta \in (0, 1)$. Assume that $\zeta_t = 2 \log(1/U_t)$, where $\{U_t\}_{t \geq 1}$ is a sequence of mutually independent random variables that satisfy $\mathbb{E}[U_t] = \frac{6\delta}{|\mathcal{X}|t^2\pi^2}$ and $\Pr(U_t \in (0, 1)) = 1$ for all $t \geq 1$. Then, by running the randomized GP-UCB with ζ_t , the cumulative regret can be bounded with probability at least $1 - \delta$ as follows:*

$$\Pr\left(\forall T \geq 1, R_T \leq 2\sqrt{C_1\gamma_T \sum_{t=1}^T \zeta_t}\right) \geq 1 - \delta,$$

where $C_1 := 2/\log(1 + \sigma^{-2})$.

See Appendix C for the proof. Note that the resulting upper bound contains the random quantity ζ_t . Thus, this theorem still does not rigorously imply a sub-linear regret upper bound. On the other hand, we can see that Theorem 4.5 is a generalized version of the analysis of GP-UCB (Srinivas et al. 2010) since the randomized GP-UCB is equivalent to GP-UCB if $U_t = \frac{6\delta}{|\mathcal{X}|t^2\pi^2}$ almost surely. If we adopt Theorem 4.5 to IRGP-UCB and additionally take the upper bound regarding $\sum_{t=1}^T \zeta_t$, we can obtain the following theorem:

Corollary 4.1. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Pick $\delta \in (0, 1)$. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|\mathcal{X}|t^2\pi^2/(6\delta))$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , the cumulative regret can be bounded with probability at least $1 - \delta$ as follows:*

$$\Pr \left(\forall T \geq 1, R_T \leq 2 \sqrt{C_1 \gamma_T \left(T s_T + T + 2 \sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right) + 2 \log \left(\frac{\pi^2 T^2}{3\delta} \right)} \right)} \right) \geq 1 - \delta,$$

where $C_1 = 2/\log(1 + \sigma^{-2})$.

See Appendix C for the proof. We can see that even if confidence parameters are randomized, the same order regret upper bound as that of GP-UCB can be obtained. Thus, this result implies that randomization does not deteriorate the high-probability regret upper bounds compared with the existing result shown by Srinivas et al. (2010).

However, in Theorem 4.5, we can see that

$$\mathbb{E}[\zeta_t] = \mathbb{E}[-2 \log(U_t)] \geq -2 \log(\mathbb{E}[U_t]) = 2 \log \left(\frac{|\mathcal{X}|t^2\pi^2}{3\delta} \right).$$

Therefore, although ζ_t may be smaller than the usual β_t in GP-UCB, its expectation $\mathbb{E}[\zeta_t]$ is larger than the usual β_t . We leave the high-probability regret analysis without increasing the confidence parameter as future work.

4.4 Expected Regret of GP-UCB with Constant Confidence Parameter

In this section, we show an expected regret lower bound of GP-UCB with any constant confidence parameter. This analysis is added from the previous conference version of this paper. We can show that GP-UCB, using any constant confidence parameter, incurs linear expected cumulative regret in the following simple problem instance:

Theorem 4.6. *Assume that $\mathcal{X} = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}\}$, $\epsilon_t \sim \mathcal{N}(0, 1)$ for all $t \geq 1$, and*

$$\begin{pmatrix} f(\mathbf{x}^{(1)}) \\ f(\mathbf{x}^{(2)}) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 0.99 \end{pmatrix} \right),$$

where $\rho < 1$ is a covariance parameter. Then, if GP-UCB with any constant confidence parameter $\beta_t^{(1/2)} = c > 0$ runs, then BCR grows linearly, that is,

$$\text{BCR}_T = \mathbb{E} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \right] = \Omega(T).$$

See Appendix D for the proof. In the proof, we show that the probability of the extreme case, where the GP-UCB algorithm continue to evaluate $\mathbf{x}^{(1)}$ even if $f(\mathbf{x}^{(1)}) < f(\mathbf{x}^{(2)})$, is rigorously positive.

Theorem 4.6 suggests that, under assumptions of Theorems 4.1 and 4.3, the confidence parameter must be increased in the GP-UCB algorithm to obtain sub-linear expected regret upper bounds. Therefore, randomizing confidence parameters as in IRGP-UCB is necessary to avoid an increase in confidence parameters.




Fig. 1. This figure shows the confidence parameters of GP-UCB-based methods. For GP-UCB, the solid line represents β_t . For RGP-UCB and IRGP-UCB, the solid lines and shaded areas represent the expectations $\mathbb{E}[\zeta_t]$ and 95% credible intervals, respectively.




Fig. 2. This figure shows the average and standard errors of simple regrets.

5 Experiments

We demonstrate the experimental results on synthetic and benchmark functions and the materials dataset provided in Liang et al. (2021). As a baseline, we performed EI (Mockus et al. 1978), TS (Russo and Van Roy 2014), MES (Wang and Jegelka 2017), joint entropy search (JES) (Hvarfner et al. 2022), GP-UCB (Srinivas et al. 2010), and PIMS (Takeno et al. 2024). For the posterior sampling in TS, PIMS, MES, and JES, we used random Fourier features (Rahimi and Recht 2008). For Monte Carlo estimation in MES and JES, we used ten samples. We evaluate the practical performance of BO by simple regret $f(\mathbf{x}^*) - \max_{t \leq T} f(\mathbf{x}_t)$.




Fig. 3. Average and standard errors of simple regret on benchmark functions.

5.1 Synthetic Function Experiment

We use synthetic functions generated from $\mathcal{GP}(0, k)$, where k is a Gaussian kernel with a length scale parameter $\ell = 0.1$ and the input dimension $d = 3$. We set the noise variance $\sigma^2 = 10^{-4}$. The input domain consists of equally divided points in $[0, 0.9]$, that is, $\mathcal{X} = \{0, 0.1, \dots, 0.9\}^d$ and $|\mathcal{X}| = 1000$. Figure 1 shows the theoretical confidence parameters of GP-UCB, RGP-UCB with Gamma($\kappa_t = \log(|\mathcal{X}|t^2)/\log(1.5)$, $\theta = 1$), and IRGP-UCB with a two-parameter exponential distribution. We can see that the confidence parameters for GP-UCB and RGP-UCB are considerably large due to a logarithmic increase, particularly in the later iterations. In contrast, the confidence parameter of IRGP-UCB is drastically smaller than that of GP-UCB and does not change with respect to the iteration. Therefore, we can observe that IRGP-UCB alleviates the over-exploration.

Next, we report the simple regret in Figure 2. Ten synthetic functions and ten initial training datasets are randomly generated. Thus, the average and standard error for 10×10 trials are reported. The hyperparameters for GP are fixed to those used to generate the synthetic functions. The confidence parameters for GP-UCB, RGP-UCB, and IRGP-UCB are set as in Figure 1. We can see that the regrets of GP-UCB, RGP-UCB, TS, and JES converge slowly. We conjecture that this slow convergence came from over-exploration. On the other hand, EI, PIMS, MES, and IRGP-UCB show fast convergence. In particular, IRGP-UCB achieves the best average in most iterations. These results suggest that IRGP-UCB bridges the gap between theory and practice in contrast to GP-UCB and RGP-UCB.

5.2 Benchmark Function Experiments

We employ three benchmark functions called Holder table ($d = 2$), Cross in tray ($d = 2$), and Ackley ($d = 4$) functions, whose analytical forms are shown at <https://www.sfu.ca/~ssurjano/optimization.html>. For each function, we report the average and standard error for 10 trials using ten random initial datasets \mathcal{D}_0 , where $|\mathcal{D}_0| = 2^d$. We set the noise variance $\sigma^2 = 10^{-4}$. We used the Gaussian kernel with automatic relevance determination, whose hyperparameter was selected by marginal likelihood maximization (Rasmussen and Williams 2005) per 5 iterations.

In this experiment, since the input domain is continuous, the theoretical choice of the confidence parameter contains an unknown variable. Thus, we use the heuristic choice for confidence parameters. For GP-UCB, we set the confidence parameter as $\beta_t = 0.2d \log(2t)$, which is the heuristic used in Kandasamy et al. (2017, 2015). For RGP-UCB, we set $\zeta_t \sim \text{Gamma}(\kappa_t, \theta = 1)$ with $\kappa_t = 0.2d \log(2t)$ since $\mathbb{E}[\zeta_t]$ must have the same order as β_t (note that $\mathbb{E}[\zeta_t] = \theta\kappa_t$). For IRGP-UCB, we set $s = d/2$ and $\lambda = 1/2$. Note that λ equals the theoretical setting, and s captures the dependence on d , as shown in Corollary 4.2.




Fig. 4. Average and standard errors of simple regret on real-world datasets.

Figure 3 shows the results. In the Holder table function, JES shows faster convergence until 40 iterations. However, the regrets of all baseline methods stagnate in $[10^{-1}, 10^0]$. In contrast, the regret of IRGP-UCB converges to 10^{-3} until 60 iterations. In the Cross in tray function, IRGP-UCB showed rapid convergence at approximately 50 iterations. In the Ackley function, IRGP-UCB constantly belonged to the top group and showed minimum regret after 125 iterations. Since the input dimension is relatively large, TS and JES, which depend on the sampled maxima \mathbf{x}_* , deteriorate by over-exploration. Throughout the experiments, IRGP-UCB outperforms TS, PIMS, GP-UCB, and RGP-UCB, which achieve sub-linear BCR bounds, and EI and MES, which are practically well-used. This result supports the effectiveness of randomization using the two-parameter exponential distribution.

5.3 Real-World Dataset Experiments

This section provides the experimental results on the materials datasets provided in Liang et al. (2021). In the perovskite dataset (Sun et al. 2021), we optimize environmental stability with respect to composition parameters for halide perovskite ($d = 3$ and $|\mathcal{X}| = 94$). In the P3HT/CNT dataset (Bash et al. 2021), we optimize electrical conductivity with respect to composition parameters for the carbon nanotube polymer blend ($d = 5$ and $|\mathcal{X}| = 178$). In the AgNP dataset (Mekki-Berrada et al. 2021), we optimize the absorbance spectrum of synthesized silver nanoparticles with respect to processing parameters for synthesizing triangular nanoprisms ($d = 5$ and $|\mathcal{X}| = 164$). See Liang et al. (2021) for more details about each dataset.

We set the initial dataset size $|\mathcal{D}_0| = 2$ as with Liang et al. (2021). Since the dataset size is small at earlier iterations and the dataset contains fluctuations from real-world experiments, we observed that hyperparameter tuning could be unstable. Thus, we optimized the hyperparameters of the RBF kernel in each iteration to avoid repeatedly obtaining samples using an inappropriate hyperparameter. The other settings matched those used in the benchmark function experiments.

Figure 4 shows the results. In the perovskite dataset, IRGP-UCB constantly belonged to the top group and showed the best performance after 20 iterations. In the P3HT/CNT dataset, EI converged to the smallest value after 60 iterations. On the other hand, IRGP-UCB shows faster convergence during the first 20 iterations. In the AgNP dataset, IRGP-UCB found the optimal point after 42 iterations in all the trials. In this experiment, heuristic methods, EI, MES, and JES, showed worse performance and required at least 60 function evaluations to find the optimal point. Consequently, we can observe the effectiveness of IRGP-UCB against real-world datasets.

6 Conclusion

This paper analyzes the expected regret of the randomized variant of GP-UCB called IRGP-UCB in the Bayesian setting. Moreover, to fairly compare with the deterministic algorithm, we obtain the high-probability upper bound

of the conditional expected regret averaged over the randomness of the GP and the noise, and conditioned on the algorithm's randomness. In both analyses, if the input domain is finite, increasing the confidence parameter, which can harm the practical efficiency, is avoided. In addition, by showing the specific instance where GP-UCB using a constant confidence parameter incurs linear expected cumulative regret, we confirmed that randomization plays a key role in avoiding increasing confidence parameters. Finally, we show the practical effectiveness of the IRGP-UCB through the numerical experiments, including the application to the materials dataset.

Several directions for future work can be considered. First, whether we can show high-probability regret bounds without increasing the confidence parameter is of interest. Second, avoiding an increase in the confidence parameter for the continuous domain is also an important direction. Last, as with the randomized LSE proposed recently (Inatsu, Takeno, Kutsukake, et al. 2024), we may be able to extend IRGP-UCB to various other practical settings, where the usual GP-UCB-based methods have been extended, for example, multi-objective BO (Inatsu, Takeno, Hanada, et al. 2024; Paria et al. 2020; Suzuki et al. 2020; Zuluaga et al. 2016), multi-fidelity BO (Kandasamy et al. 2016; Kandasamy et al. 2017; Takeno, Fukuoka, et al. 2022; Takeno et al. 2020), parallel BO (Contal et al. 2013; Desautels et al. 2014), high-dimensional BO (Kandasamy et al. 2015), cascade BO (Kusakawa et al. 2022), and robust BO (Bogunovic et al. 2018; Inatsu et al. 2022; Iwazaki et al. 2021).

Acknowledgments

This work was supported by MEXT KAKENHI 21H03498, 22H00300, MEXT Program: Data Creation and Utilization-Type Material Research and Development Project Grant Number JPMXP1122712807, JSPS KAKENHI Grant Number JP20H00601, JP21J14673, JP23K16943, JP23K19967 and JP24K20847, and JST ACT-X Grant Number JPMJAX23CD, RIKEN Center for Advanced Intelligence Project.

References

- K. Azuma. 1967. "Weighted sums of certain dependent random variables." *Tohoku Mathematical Journal, Second Series*, 19, 3, 357–367.
- D. Bash et al. 2021. "Multi-Fidelity High-Throughput Optimization of Electrical Conductivity in P3HT-CNT Composites." *Advanced Functional Materials*, 31, 2102606.
- M. A. Beg. 1980. "On the estimation of $\Pr \{Y < X\}$ for the two-parameter exponential distribution." *Metrika*, 27, 1, 29–34.
- J. Berk, S. Gupta, S. Rana, and S. Venkatesh. 2020. "Randomised Gaussian process upper confidence bound for Bayesian optimisation." In: *Proceedings of the 29th International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization, 2284–2290.
- M. Biskup. 2020. "Extrema of the two-dimensional discrete Gaussian free field." In: *Random Graphs, Phase Transitions, and the Gaussian Free Field: PIMS-CRM Summer School in Probability, Vancouver, Canada, June 5–30, 2017*. Springer, 163–407.
- I. Bogunovic, J. Scarlett, S. Jegelka, and V. Cevher. 2018. "Adversarially robust optimization with Gaussian processes." In: *Advances in Neural Information Processing Systems 31*. Curran Associates, Inc., 5760–5770.
- I. Bogunovic, J. Scarlett, A. Krause, and V. Cevher. 2016. "Truncated variance reduction: A unified approach to Bayesian optimization and level-set estimation." In: *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc., 1507–1515.
- A. D. Bull. 2011. "Convergence rates of efficient global optimization algorithms." *Journal of Machine Learning Research*, 12, 10, 2879–2904.
- S. R. Chowdhury and A. Gopalan. 2017. "On kernelized multi-armed bandits." In: *Proceedings of the 34th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 70. PMLR, 844–853.
- E. Contal, D. Buffoni, A. Robicquet, and N. Vayatis. 2013. "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*. Springer-Verlag, 225–240.
- T. Desautels, A. Krause, and J. W. Burdick. 2014. "Parallelizing exploration-exploitation tradeoffs in Gaussian process bandit optimization." *Journal of Machine Learning Research*, 15, 4053–4103.
- M. G. Genton. 2001. "Classes of kernels for machine learning: a statistics perspective." *Journal of Machine Learning Research*, 2, 299–312.
- S. Ghosal and A. Roy. 2006. "Posterior consistency of Gaussian process prior for nonparametric binary regression." *The Annals of Statistics*, 34, 5, 2413–2429.
- A. Gotovos, N. Casati, G. Hitz, and A. Krause. 2013. "Active learning for level set estimation." In: *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*. AAAI Press, 1344–1350.

- R.-R. Griffiths and J. M. Hernández-Lobato. 2020. “Constrained Bayesian optimization for automatic chemical design using variational autoencoders.” *Chemical science*, 11, 2, 577–586.
- S. Grünewälder, J.-Y. Audibert, M. Opper, and J. Shawe–Taylor. 2010. “Regret bounds for Gaussian process bandit problems.” In: *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 9. PMLR, 273–280.
- R. V. Handel. 2016. *Probability in High Dimension*. Lecture notes. Available in <https://web.math.princeton.edu/~rvan/APC550.pdf>. (2016).
- P. Hennig and C. J. Schuler. 2012. “Entropy search for information-efficient global optimization.” *Journal of Machine Learning Research*, 13, 57, 1809–1837.
- J. M. Hernández-Lobato, M. W. Hoffman, and Z. Ghahramani. 2014. “Predictive entropy search for efficient global optimization of black-box functions.” In: *Advances in Neural Information Processing Systems 27*. Curran Associates, Inc., 918–926.
- C. Hvarfner, F. Hutter, and L. Nardi. 2022. “Joint Entropy Search For Maximally-Informed Bayesian Optimization.” In: *Advances in Neural Information Processing Systems 35*. Curran Associates, Inc., 11494–11506.
- Y. Inatsu, S. Takeno, H. Hanada, K. Iwata, and I. Takeuchi. 2024. “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* (Proceedings of Machine Learning Research). Vol. 238. PMLR, 4564–4572.
- Y. Inatsu, S. Takeno, M. Karasuyama, and I. Takeuchi. 2022. “Bayesian optimization for distributionally robust chance-constrained problem.” In: *Proceedings of the 39th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 162. PMLR, 9602–9621.
- Y. Inatsu, S. Takeno, K. Kutsukake, and I. Takeuchi. 2024. “active learning for level set estimation using randomized straddle algorithms.” *Transactions on Machine Learning Research*. <https://openreview.net/forum?id=N8M2yqRicS>.
- S. Iwazaki, Y. Inatsu, and I. Takeuchi. 2021. “mean-variance analysis in Bayesian optimization under uncertainty.” In: *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 130. PMLR, 973–981.
- D. Janz, D. Burt, and J. Gonzalez. 2020. “Bandit optimisation of functions in the Matérn kernel RKHS.” In: *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 108. PMLR, 2486–2495.
- K. Kandasamy, G. Dasarathy, J. Oliva, J. Schneider, and B. Póczos. 2016. “Gaussian process bandit optimisation with multi-fidelity evaluations.” In: *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc., 1000–1008.
- K. Kandasamy, G. Dasarathy, J. Schneider, and B. Póczos. 2017. “Multi-fidelity Bayesian optimisation with continuous approximations.” In: *Proceedings of the 34th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 70. PMLR, 1799–1808.
- K. Kandasamy, A. Krishnamurthy, J. Schneider, and B. Póczos. 2018. “Parallelised Bayesian optimisation via Thompson sampling.” In: *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 84. PMLR, 133–142.
- K. Kandasamy, J. Schneider, and B. Póczos. 2015. “High dimensional Bayesian optimisation and bandits via additive models.” In: *Proceedings of the 32nd International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 37. PMLR, 295–304.
- K. Korovina, S. Xu, K. Kandasamy, W. Neiswanger, B. Póczos, J. Schneider, and E. Xing. 2020. “ChemBO: Bayesian optimization of small organic molecules with synthesizable recommendations.” In: *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 108. PMLR, 3393–3403.
- S. Kusakawa, S. Takeno, Y. Inatsu, K. Kutsukake, S. Iwazaki, T. Nakano, T. Ujihara, M. Karasuyama, and I. Takeuchi. 2022. “Bayesian optimization for cascade-type multistage processes.” *Neural Computation*, 34, 12, 2408–2431.
- H. J. Kushner. 1964. “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.
- K. Lam, B. K. Sinha, and Z. Wu. 1994. “Estimation of parameters in a two-parameter exponential distribution using ranked set sample.” *Annals of the Institute of Statistical Mathematics*, 46, 4, 723–736.
- B. Laurent and P. Massart. 2000. “Adaptive estimation of a quadratic functional by model selection.” *The Annals of Statistics*, 28, 5, 1302–1338.
- Q. Liang et al. 2021. “Benchmarking the performance of Bayesian optimization across multiple experimental materials science domains.” *npj Computational Material*, 7, 188, 10 pages.
- F. Mekki-Berrada et al.. 2021. “Two-step machine learning enables optimized nanoparticle synthesis.” *npj Computational Materials*, 7, 1, 1–10.
- J. Mockus, V. Tiesis, and A. Zilinskas. 1978. “The application of Bayesian methods for seeking the extremum.” *Towards Global Optimization*, 2, 117–129, 2.
- B. Paria, K. Kandasamy, and B. Póczos. 2020. “A flexible framework for multi-objective Bayesian optimization using random scalarizations.” In: *Proceedings of the 35th Uncertainty in Artificial Intelligence Conference* (Proceedings of Machine Learning Research). Vol. 115. PMLR, 766–776.
- A. Rahimi and B. Recht. 2008. “Random features for large-scale kernel machines.” In: *Advances in Neural Information Processing Systems 20*. Curran Associates, Inc., 1177–1184.
- C. E. Rasmussen and C. K. I. Williams. 2005. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

- D. Russo and B. Van Roy. 2014. “Learning to optimize via posterior sampling.” *Mathematics of Operations Research*, 39, 4, 1221–1243.
- J. Scarlett. 2018. “Tight regret bounds for Bayesian optimization in one dimension.” In: *Proceedings of the 35th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 80. PMLR, 4500–4508.
- B. Shahriari, K. Swersky, Z. Wang, R. Adams, and N. De Freitas. 2016. “Taking the human out of the loop: A review of Bayesian optimization.” *Proceedings of the IEEE*, 104, 1, 148–175.
- J. Snoek, H. Larochelle, and R. P. Adams. 2012. “Practical Bayesian optimization of machine learning algorithms.” In: *Advances in Neural Information Processing Systems 25*. Curran Associates, Inc., 2951–2959.
- N. Srinivas, A. Krause, S. M. Kakade, and M. W. Seeger. 2010. “Gaussian process optimization in the bandit setting: no regret and experimental design.” In: *Proceedings of the 27th International Conference on Machine Learning*. Omnipress, 1015–1022.
- S. Sun et al. 2021. “A data fusion approach to optimize compositional stability of halide perovskites.” *Matter*, 4, 4, 1305–1322.
- S. Suzuki, S. Takeno, T. Tamura, K. Shitara, and M. Karasuyama. 2020. “Multi-objective Bayesian optimization using Pareto-frontier entropy.” In: *Proceedings of the 37th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 119. PMLR, 9279–9288.
- S. Takeno, H. Fukuoka, Y. Tsukada, T. Koyama, M. Shiga, I. Takeuchi, and M. Karasuyama. 2022. “A generalized framework of multifidelity max-value entropy search through joint entropy.” *Neural Computation*, 34, 10, 2145–2203.
- S. Takeno, H. Fukuoka, Y. Tsukada, T. Koyama, M. Shiga, I. Takeuchi, and M. Karasuyama. 2020. “Multi-fidelity Bayesian optimization with max-value entropy search and its parallelization.” In: *Proceedings of the 37th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 119. PMLR, 9334–9345.
- S. Takeno, Y. Inatsu, and M. Karasuyama. 2023. “Randomized Gaussian process upper confidence bound with tighter Bayesian regret bounds.” In: *Proceedings of the 40th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 202. PMLR, 33490–33515.
- S. Takeno, Y. Inatsu, M. Karasuyama, and I. Takeuchi. 2024. “Posterior sampling-based Bayesian optimization with tighter Bayesian regret bounds.” In: *Proceedings of the 41st International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 235. PMLR, 47510–47534.
- S. Takeno, T. Tamura, K. Shitara, and M. Karasuyama. 2022. “Sequential and parallel constrained max-value entropy search via information lower bound.” In: *Proceedings of the 39th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 162. PMLR, 20960–20986.
- H. Tran-The, S. Gupta, S. Rana, and S. Venkatesh. 2022. “regret bounds for expected improvement algorithms in Gaussian process bandit optimization.” In: *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 151. PMLR, 8715–8737.
- T. Ueno, T. D. Rhone, Z. Hou, T. Mizoguchi, and K. Tsuda. 2016. “COMBO: An efficient Bayesian optimization library for materials science.” *Materials discovery*, 4, 18–21.
- S. Vakili, K. Khezeli, and V. Picheny. 2021. “on information gain and regret bounds in Gaussian process bandits.” In: *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 130. PMLR, 82–90.
- E. Vazquez and J. Bect. 2010. “Convergence properties of the expected improvement algorithm with fixed mean and covariance functions.” *Journal of Statistical Planning and Inference*, 140, 11, 3088–3095.
- R. Vershynin. 2018. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Z. Wang and S. Jegelka. 2017. “Max-value entropy search for efficient Bayesian optimization.” In: *Proceedings of the 34th International Conference on Machine Learning* (Proceedings of Machine Learning Research). Vol. 70. PMLR, 3627–3635.
- Z. Wang, B. Zhou, and S. Jegelka. 2016. “Optimization as estimation with Gaussian processes in bandit settings.” In: *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics* (Proceedings of Machine Learning Research). Vol. 51. PMLR, 1022–1031.
- Z. Wang and N. de Freitas. 2014. “Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters.” *arXiv:1406.7758*.
- O. Zeitouni. 2014. *Gaussian fields*. Notes for Lectures. Available in <https://www.wisdom.weizmann.ac.il/~zeitouni/notesGauss.pdf>. (2014).
- M. Zuluaga, A. Krause, and M. Püschel. 2016. “e-PAL: An active learning approach to the multi-objective optimization problem.” *Journal of Machine Learning Research*, 17, 104, 1–32.

A Proofs for Expected Regret Bounds of IRGP-UCB

This section provides proofs for the expected regret bounds of IRGP-UCB.

A.1 Proof of Lemma 4.1

The main difference between Lemma 4.1 and Lemma 5.1 of Srinivas et al. (2010) is that Lemma 4.1 does not consider the union bound for all $t \geq 1$. That is, Lemma 5.1 of Srinivas et al. (2010) considers the following probability:

$$\Pr \left(f(\mathbf{x}) \leq \mu_{t-1}(\mathbf{x}) + \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X}, \forall t \geq 1 \right).$$

On the other hand, Lemma 4.1 bounds the following probability with one fixed t :

$$\Pr_t \left(f(\mathbf{x}) \leq \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X} \right),$$

where $\beta_\delta^{1/2}$ does not depend on t as described below.

Assume the same condition as in Lemma 4.1, and thus, $\beta_\delta = 2 \log(|\mathcal{X}|/(2\delta))$. Then, for all $\mathbf{x} \in \mathcal{X}$ and \mathcal{D}_{t-1} , we see that

$$\Pr_t \left(f(\mathbf{x}) > \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}) \right) \leq \frac{\delta}{|\mathcal{X}|},$$

where we use an elementary result of the Gaussian distribution shown in Lemma E.3 in Appendix E. Then, we can obtain the following bound:

$$\begin{aligned} \Pr_t \left(f(\mathbf{x}) > \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}), \exists \mathbf{x} \in \mathcal{X} \right) &\leq \sum_{\mathbf{x} \in \mathcal{X}} \Pr_t \left(f(\mathbf{x}) > \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}) \right) \\ &\leq \delta. \end{aligned}$$

Therefore,

$$\begin{aligned} &\Pr_t \left(f(\mathbf{x}) \leq \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X} \right) \\ &= 1 - \Pr_t \left(f(\mathbf{x}) > \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}), \exists \mathbf{x} \in \mathcal{X} \right) \geq 1 - \delta, \end{aligned}$$

which concludes the proof.

A.2 Proof of Lemma 4.2

This section provides detailed proof of the following Lemma 4.2:

Lemma 4.2. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel, and \mathcal{X} be finite. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|\mathcal{X}|/2)$ and $\lambda = 1/2$. Then, for any given \mathcal{D}_{t-1} , the following inequality holds:*

$$\mathbb{E}_t[f(\mathbf{x}^*)] \leq \mathbb{E}_t \left[\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}) \right],$$

for all $t \geq 1$.

PROOF. Fix the dataset \mathcal{D}_{t-1} and $\delta \in (0, 1)$. From Lemma 4.1, we can see that

$$\Pr_t \left(f(\mathbf{x}^*) \leq \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}) \right) \geq 1 - \delta,$$

where $\beta_\delta = 2 \log(|\mathcal{X}|/(2\delta))$. Note that the probability is taken only with $f(\mathbf{x}^*)$ since \mathcal{D}_{t-1} is fixed. Using the cumulative distribution function $F_t(\cdot) := \Pr_t(f(\mathbf{x}^*) \leq \cdot)$ and its inverse function F_t^{-1} , we can rewrite

$$\begin{aligned} F_t \left(\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}) \right) &\geq 1 - \delta \\ \iff \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_\delta^{1/2} \sigma_{t-1}(\mathbf{x}) &\geq F_t^{-1}(1 - \delta), \end{aligned}$$

since F_t^{-1} is a monotone increasing function. Since $\delta \in (0, 1)$ is arbitrary, we can see that

$$\Pr \left(\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_U^{1/2} \sigma_{t-1}(\mathbf{x}) \geq F_t^{-1}(1 - U) \right) = 1,$$

where $U \sim \text{Uni}(0, 1)$ and probability is taken by the randomness of U . Then, because of the basic property of the expectation $\mathbb{E}[X] \leq \mathbb{E}[Y]$ if $\Pr(X \leq Y) = 1$, we can obtain,

$$\begin{aligned} \mathbb{E}_U \left[\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_U^{1/2} \sigma_{t-1}(\mathbf{x}) \right] &\geq \mathbb{E}_U [F_t^{-1}(1 - U)] \\ &= \mathbb{E}_U [F_t^{-1}(U)], \end{aligned}$$

where we use the fact that $1 - U$ also follows $\text{Uni}(0, 1)$. Since $F_t^{-1}(U)$ and $f(\mathbf{x}^*)$ are identically distributed as with the inverse transform sampling, we obtain

$$\mathbb{E}_U \left[\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \beta_U^{1/2} \sigma_{t-1}(\mathbf{x}) \right] \geq \mathbb{E}_t [f(\mathbf{x}^*)].$$

Then, β_U can be decomposed as follows:

$$\beta_U = 2 \log(|\mathcal{X}|/2) - 2 \log(U).$$

From the inverse transform sampling, $-2 \log(U) \sim \text{Exp}(\lambda = 1/2)$. Hence, β_U follows a shifted exponential distribution with $s = 2 \log(|\mathcal{X}|/2)$ and $\lambda = 1/2$. Therefore, we obtain the following:

$$\mathbb{E}_t \left[\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) + \zeta^{1/2} \sigma_{t-1}(\mathbf{x}) \right] \geq \mathbb{E}_t [f(\mathbf{x}^*)].$$

□

A.3 Proof of Theorem 4.1

Using Lemma 4.2, we can obtain the proof of the following theorem for finite input domains:

Theorem 4.1. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|\mathcal{X}|/2)$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , BCR can be bounded as follows:*

$$\text{BCR}_T \leq \sqrt{C_1 C_2 T \gamma_T},$$

where $C_1 := 2/\log(1 + \sigma^{-2})$ and $C_2 := 2 + 2 \log(|\mathcal{X}|/2)$.

PROOF. From Lemma 4.2, we obtain

$$\text{BCR}_T = \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - f(\mathbf{x}_t)] \quad (3)$$

$$= \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - v_t(\mathbf{x}_t) + v_t(\mathbf{x}_t) - f(\mathbf{x}_t)] \quad (4)$$

$$\leq \sum_{t=1}^T \mathbb{E}[v_t(\mathbf{x}_t) - f(\mathbf{x}_t)], \quad (5)$$

where $v_t(\mathbf{x}) = \mu_{t-1}(\mathbf{x}) + \zeta_t^{1/2}(\mathbf{x})$. Then, we can easily derive the bound as follows:

$$\text{BCR}_T \leq \sum_{t=1}^T \mathbb{E}[v_t(\mathbf{x}_t) - f(\mathbf{x}_t)] \quad (6)$$

$$= \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [\mathbb{E}[v_t(\mathbf{x}_t) - f(\mathbf{x}_t) | \mathcal{D}_{t-1}, \zeta_t]] \quad (7)$$

$$= \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_{\zeta_t} [v_t(\mathbf{x}_t) - \mu_{t-1}(\mathbf{x}_t)]] \quad (8)$$

$$= \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [\zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t)] \quad (9)$$

$$= \mathbb{E} \left[\sum_{t=1}^T \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \right] \quad (10)$$

$$\leq \mathbb{E} \left[\sqrt{\sum_{t=1}^T \zeta_t \sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}_t)} \right] \quad (\because \text{Cauchy-Schwarz inequality}) \quad (11)$$

$$\leq \mathbb{E} \left[\sqrt{\sum_{t=1}^T \zeta_t} \right] \sqrt{C_1 \gamma_T} \quad (\because \text{Lemma 5.4 in Srinivas et al. (2010)}) \quad (12)$$

$$\leq \sqrt{\sum_{t=1}^T \mathbb{E}[\zeta_t]} \sqrt{C_1 \gamma_T} \quad (\because \text{Jensen's inequality}) \quad (13)$$

$$= \sqrt{C_1 C_2 T \gamma_T}, \quad (\because \text{Definition of } \zeta_t) \quad (14)$$

where $C_2 = \mathbb{E}[\zeta_t] = 2 + 2 \log(|X|/2)$. This concludes the proof. \square

A.4 Proof of Theorem 4.2

Theorem 4.2. Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel, $k(\mathbf{x}, \mathbf{x}) \leq 1$, and Assumption 2.1 holds. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2d \log(bdr t^2 (\sqrt{\log(ad)} + \sqrt{\pi}/2)) - 2 \log 2$ and

$\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB, BCR can be bounded as follows:

$$\text{BCR}_T \leq \frac{\pi^2}{6} + \sqrt{C_1 T \gamma_T (2 + s_T)},$$

where $C_1 := 2/\log(1 + \sigma^{-2})$.

PROOF. We consider a finite discretized input set $\mathcal{X}_t \subset \mathcal{X}$ for $t \geq 1$ for theoretical analysis, though \mathcal{X}_t is not related to the actual algorithm. Let $\mathcal{X}_t = \{\frac{r}{\tau_{t+1}}, \dots, \frac{\tau_t r}{\tau_{t+1}}\}^d \subset \mathcal{X}$, where each coordinate takes values from a uniform grid of $\tau_t = bdr u_t (\sqrt{\log(ad)} + \sqrt{\pi}/2)$ points. Thus, $|\mathcal{X}_t| = \tau_t^d$. In addition, we define $[\mathbf{x}]_t$ as the nearest point in \mathcal{X}_t of $\mathbf{x} \in \mathcal{X}$.

Then, we decompose BCR as follows:

$$\text{BCR}_T = \sum_{t=1}^T \mathbb{E} \left[f(\mathbf{x}^*) - f([\mathbf{x}^*]_t) + f([\mathbf{x}^*]_t) - \max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) + \max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) - v_t(\mathbf{x}_t) + v_t(\mathbf{x}_t) - f(\mathbf{x}_t) \right]. \quad (15)$$

Obviously, $\mathbb{E}[\max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) - v_t(\mathbf{x}_t)] \leq 0$ from the definition of \mathbf{x}_t . Furthermore, we observe the following:

$$\mathbb{E} \left[f([\mathbf{x}^*]_t) - \max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) \right] \leq \mathbb{E} \left[\max_{\mathbf{x} \in \mathcal{X}_t} f(\mathbf{x}) - \max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) \right],$$

which can be bounded above by 0 by setting $s_t = 2 \log(|\mathcal{X}_t|/2)$ and $\lambda = 1/2$ from Lemma 4.2. Therefore, we obtain the following:

$$\text{BCR}_T \leq \sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - f([\mathbf{x}^*]_t) + v_t(\mathbf{x}_t) - f(\mathbf{x}_t)] \quad (16)$$

$$= \sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - f([\mathbf{x}^*]_t)] + \sum_{t=1}^T \mathbb{E} [v_t(\mathbf{x}_t) - f(\mathbf{x}_t)]. \quad (17)$$

First, we consider the first term $\sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - f([\mathbf{x}^*]_t)]$. From Lemma E.2, we can obtain the following:

$$\sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - f([\mathbf{x}^*]_t)] \leq \sum_{t=1}^T \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - f([\mathbf{x}]_t) \right] \quad (18)$$

$$\leq \sum_{t=1}^T \frac{1}{t^2} \quad (\because \text{Lemma E.2}) \quad (19)$$

$$\leq \frac{\pi^2}{6}, \quad \left(\because \sum_{t=1}^{\infty} \frac{1}{t^2} = \frac{\pi^2}{6} \right) \quad (20)$$

where u_t in Lemma E.2 corresponds to t^2 .

The second term is bounded as follows:

$$\sum_{t=1}^T \mathbb{E} [v_t(\mathbf{x}_t) - f(\mathbf{x}_t)] \leq \sqrt{\sum_{t=1}^T \mathbb{E}[\zeta_t] C_1 \gamma_T}, \quad (\because \text{See the proof of Theorem 4.1}) \quad (21)$$

where

$$\mathbb{E}[\zeta_t] = 2 + 2 \log(|\mathcal{X}_t|/2) \quad (22)$$

$$= 2 + 2d \log(bdr t^2 (\sqrt{\log(ad)} + \sqrt{\pi}/2)) - 2 \log 2, \quad (23)$$

which is monotone increasing. Therefore, we obtain the following:

$$\sum_{t=1}^T \mathbb{E} [v_t(\mathbf{x}_t) - f(\mathbf{x}_t)] \leq \sqrt{C_1 T \gamma_T \mathbb{E}[\zeta_T]}, \quad (24)$$

where $\mathbb{E}[\zeta_T] = 2 + s_T = 2 + 2d \log(bdrT^2(\sqrt{\log(ad)} + \sqrt{\pi}/2)) - 2 \log 2$. Consequently, combining Eqs. (17), (20), and (21) concludes the proof. \square

B Proofs for Conditional Expected Regret Bounds of IRGP-UCB

This section provides proofs for the conditional expected regret bounds of IRGP-UCB.

B.1 Proof of Theorem 4.3

For the proof of Theorem 4.3, we use the properties of subgaussian random variables. Therefore, before the proof of Theorem 4.3, we provide several propositions.

The definition of the subgaussian random variable is shown below, following Definition 3.5 of Handel (2016):

Definition B.1. A random variable is called σ^2 -subgaussian if the following is satisfied:

$$\mathbb{E}[\exp(\lambda X)] \leq \exp(\sigma^2 \lambda^2 / 2),$$

for all $\lambda \in \mathbb{R}$. Then, σ^2 is called the variance proxy.

Several equivalent definitions are known as in Proposition 2.5.2 of Vershynin (2018). In particular, we use the following relations:

Proposition B.1. Assume that the random variable X satisfies the following property:

$$\Pr(|X| > c) \leq 2 \exp(-c^2 / A^2),$$

for all $c \geq 0$ and $\mathbb{E}[X] = 0$. Then, X is $9A^2$ -subgaussian.

PROOF. For example, by combining Problems 3.1 (d) and (e) of Handel (2016), we can obtain this result. \square

Proposition B.2. Let X be σ^2 -subgaussian. Then,

$$\Pr(X > c) \leq \exp(-c^2 / 2\sigma^2),$$

and

$$\Pr(|X| > c) \leq 2 \exp(-c^2 / 2\sigma^2).$$

PROOF. See, for example, Problem 3.1 (c) of Handel (2016) or the proof for (v) \Rightarrow (i) of Proposition 2.5.2 in Vershynin (2018). \square

Furthermore, we use the following form of Azuma's inequality (Azuma 1967) shown in Lemma 3.7 of Handel (2016):

Lemma B.1 (Azuma's inequality). Let $\{\mathcal{F}_t\}_{t \leq T}$ be any filtration, and let X_1, \dots, X_T be random variables that satisfy the following properties for $t = 1, \dots, T$:

- (1) Martingale difference property: X_t is \mathcal{F}_t -measurable and $\mathbb{E}[X_t | \mathcal{F}_{t-1}] = 0$.
- (2) Conditional subgaussian property: $\mathbb{E}[\exp(\lambda X_t) | \mathcal{F}_{t-1}] \leq \exp(\sigma_t^2 \lambda^2 / 2)$ for all $\lambda \in \mathbb{R}$ almost surely.

Then, the sum $\sum_{t=1}^T X_t$ is subgaussian with variance proxy $\sum_{t=1}^T \sigma_t^2$.

To show the subgaussian property the variable A_1 in the following proof of Theorem 4.3, we use the following proposition shown in, for example, Theorem 5.2.2 of Vershynin (2018), Proposition 4 of Zeitouni (2014), and Lemma 6.4 of Biskup (2020):

Proposition B.3 (Concentration for Lipschitz functions). *Let $\mathbf{X} = (X_1, \dots, X_k)$ be a random vector follows $\mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, where \mathbf{I}_k is k -dimensional identity matrix. Suppose that $h : \mathbb{R}^k \rightarrow \mathbb{R}$ is Lipschitz; that is, there exists a constant $L_h \in \mathbb{R}$ such that*

$$\|h(\mathbf{a}) - h(\mathbf{b})\|_2 \leq L_h \|\mathbf{a} - \mathbf{b}\|_2,$$

for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}^k$, where $\|\cdot\|_2$ is L2 norm. Then,

$$\Pr(|h(\mathbf{X}) - \mathbb{E}[h(\mathbf{X})]| > c) \leq 2 \exp\left\{-\frac{c^2}{2L_h^2}\right\}.$$

Then, we show the following theorem:

Theorem 4.3. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|\mathcal{X}|/2)$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , the conditional expected regret can be bounded with probability at least $1 - \delta$ as follows:*

$$\Pr\left(\forall T \geq 1, \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right] \leq U(T, \delta)\right) \geq 1 - \delta,$$

where

$$U(T, \delta) = 6\sqrt{T \log\left(\frac{\pi^2 T^2}{3\delta}\right)} + \sqrt{C_1 \gamma_T \left(T s_T + T + 2\sqrt{T \log\left(\frac{\pi^2 T^2}{3\delta}\right)} + 2 \log\left(\frac{\pi^2 T^2}{3\delta}\right) \right)}$$

and $C_1 := 2/\log(1 + \sigma^{-2})$.

PROOF. First, we show an upper bound for some fixed $T \geq 1$. For any $\{\zeta_t\}_{t \geq 1}$, we obtain

$$\begin{aligned} & \sum_{t=1}^T \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} [f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1}] \\ &= \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_{f \mid \mathcal{D}_{t-1}} [f(\mathbf{x}^*)] - \mu_{t-1}(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}] \\ &= \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_{f \mid \mathcal{D}_{t-1}} [f(\mathbf{x}^*)] - \mathbb{E}_{\zeta_t \mid \mathcal{D}_{t-1}} [v_t(\mathbf{x}_t)] + \mathbb{E}_{\zeta_t \mid \mathcal{D}_{t-1}} [v_t(\mathbf{x}_t)] - v_t(\mathbf{x}_t) + v_t(\mathbf{x}_t) - \mu_{t-1}(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}] \\ &\leq \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_{\zeta_t \mid \mathcal{D}_{t-1}} [v_t(\mathbf{x}_t)] - v_t(\mathbf{x}_t) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}] \quad (\because \text{Lemma 4.2}) \\ &= \underbrace{\sum_{t=1}^T \left\{ \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \{\zeta_i\}_{i < t}] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}] \right\}}_{A_1} + \underbrace{\mathbb{E} \left[\sum_{t=1}^T \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \mid \{\zeta_t\}_{t > 1} \right]}_{A_2}, \end{aligned}$$

where $v_t(\mathbf{x}_t) = \mu_{t-1}(\mathbf{x}_t) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t)$.

First, regarding A_2 , the following inequality holds with probability at least $1 - \delta_2$,

$$\begin{aligned}
 & \mathbb{E} \left[\sum_{t=1}^T \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \mid \{\zeta_t\}_{t>1} \right] \\
 & \leq \sqrt{\sum_{t=1}^T \zeta_t} \mathbb{E} \left[\sqrt{\sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}_t)} \mid \{\zeta_t\}_{t>1} \right] && (\because \text{Cauchy-Schwartz inequality}) \\
 & \leq \sqrt{C_1 \gamma_T \sum_{t=1}^T \zeta_t} && (\because \text{Lemma 5.4 of Srinivas et al. (2010)}) \\
 & \leq \sqrt{C_1 \gamma_T \left(T s_T + T + 2\sqrt{T} \log(1/\delta_2) + 2 \log(1/\delta_2) \right)}, && (\because \text{Lemma E.4})
 \end{aligned}$$

where the last inequality holds with high-probability since $\sum_{t=1}^T \{\zeta_t - s_t\}$ follows a chi-square distribution with T degrees of freedom. Note that, regarding the first inequality, $\sqrt{\sum_{t=1}^T \zeta_t}$ can be put outside of the conditional expectation since the expectation is taken with f and $\{\epsilon_t\}_{t \geq 1}$.

Second, we consider A_1 . We will show that A_1 is an $18T$ -subgaussian random variable. Let history $\mathcal{H}_t = \{\zeta_i\}_{i \leq t}$. Then, we can see that

$$\begin{aligned}
 & \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \{\zeta_i\}_{i < t}] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \{\zeta_i\}_{i \leq t}] \\
 & = \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \mathcal{H}_{t-1}] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \mathcal{H}_t],
 \end{aligned}$$

and, for any \mathcal{H}_{t-1} , the conditional expectation is zero as follows:

$$\begin{aligned}
 & \mathbb{E}_{\zeta_t} \left[\mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \mathcal{H}_{t-1}] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \mathcal{H}_t] \mid \mathcal{H}_{t-1} \right] \\
 & = \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \mathcal{H}_{t-1}] - \mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \mathcal{H}_{t-1}] = 0.
 \end{aligned}$$

In addition, from the definition of measurability, $\mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \mathcal{H}_t]$ is \mathcal{H}_t -mesurable.

Then, we derive the conditional subgaussian property of $\mathbb{E}_{\mathcal{D}_{t-1}, \zeta_t} [v_t(\mathbf{x}_t) \mid \mathcal{H}_{t-1}] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \mathcal{H}_t]$. Let a function $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ be

$$h(\mathbf{a}) = \mathbb{E}_{\mathcal{D}_{t-1}} \left[\max_{\mathbf{x} \in \mathcal{X}} \left\{ \mu_{t-1}(\mathbf{x}) + \sqrt{2 \log(|\mathcal{X}|/2) + \|\mathbf{a}\|_2^2} \sigma_{t-1}(\mathbf{x}) \right\} \mid \mathcal{H}_{t-1} \right].$$

Note that $(h(X) \mid \mathcal{H}_{t-1}) \stackrel{d}{=} (\mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \zeta_t] \mid \mathcal{H}_{t-1})$ for any fixed \mathcal{H}_{t-1} , where $X \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$ and $\stackrel{d}{=}$ means the equivalence of the distribution, since $\|X\|_2^2$ follows the exponential distribution with $\lambda = 1/2$. Then, for any fixed

\mathcal{H}_{t-1} , h is 1-Lipschitz function as shown below:

$$\begin{aligned}
& |h(\mathbf{a}) - h(\mathbf{b})| \\
&= \left| \mathbb{E}_{\mathcal{D}_{t-1}} \left[\max_{\mathbf{x} \in \mathcal{X}} \left\{ \mu_{t-1}(\mathbf{x}) + \sqrt{2 \log(|\mathcal{X}|/2) + \|\mathbf{a}\|_2^2 \sigma_{t-1}(\mathbf{x})} \right\} \right. \right. \\
&\quad \left. \left. - \max_{\mathbf{x} \in \mathcal{X}} \left\{ \mu_{t-1}(\mathbf{x}) + \sqrt{2 \log(|\mathcal{X}|/2) + \|\mathbf{b}\|_2^2 \sigma_{t-1}(\mathbf{x})} \right\} \mid \mathcal{H}_{t-1} \right] \right| \\
&\leq \left| \mathbb{E}_{\mathcal{D}_{t-1}} \left[\max_{\mathbf{x} \in \mathcal{X}} \left\{ \sqrt{2 \log(|\mathcal{X}|/2) + \|\mathbf{a}\|_2^2 \sigma_{t-1}(\mathbf{x})} - \sqrt{2 \log(|\mathcal{X}|/2) + \|\mathbf{b}\|_2^2 \sigma_{t-1}(\mathbf{x})} \right\} \mid \mathcal{H}_{t-1} \right] \right| \\
&\stackrel{(a)}{\leq} \left| \sqrt{2 \log(|\mathcal{X}|/2) + \|\mathbf{a}\|_2^2} - \sqrt{2 \log(|\mathcal{X}|/2) + \|\mathbf{b}\|_2^2} \right| \\
&\stackrel{(b)}{\leq} \|\mathbf{a}\|_2 - \|\mathbf{b}\|_2 \\
&\stackrel{(c)}{\leq} \|\mathbf{a} - \mathbf{b}\|_2.
\end{aligned}$$

The inequalities (a), (b), and (c) are derived by (a) $\sigma_{t-1}(\mathbf{x}) \leq 1$, (b) the monotone decreasing property of the gradient of $\sqrt{\cdot}$, and (c) the triangle inequality. Therefore, from Proposition B.3 and $(h(\mathbf{X}) \mid \mathcal{H}_{t-1}) \stackrel{d}{=} (\mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \zeta_t] \mid \mathcal{H}_{t-1})$,

$$\Pr_{\zeta_t} \left(\left| \mathbb{E}_{\mathcal{D}_{t-1, \zeta_t}} [v_t(\mathbf{x}_t)] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t)] \right| > c \mid \mathcal{H}_{t-1} \right) \leq 2 \exp \{-c^2/2\},$$

for any \mathcal{H}_{t-1} . Therefore, using Proposition B.1, $\mathbb{E}_{\mathcal{D}_{t-1, \zeta_t}} [v_t(\mathbf{x}_t) \mid \mathcal{H}_{t-1}] - \mathbb{E}_{\mathcal{D}_{t-1}} [v_t(\mathbf{x}_t) \mid \mathcal{H}_{t-1}] \mid \mathcal{H}_{t-1}$ is 18-subgaussian.

Finally, since we see that A_1 is the sum of random variables that satisfy the condition of Azuma's inequality, A_1 is $18T$ -subgaussian. Therefore, from Proposition B.2, the following inequality holds with probability at least $1 - \delta_1$,

$$A_1 \leq 6\sqrt{T \log(1/\delta_1)}.$$

Hence, for some fixed $T \geq 1$, we obtain the high-probability bounds of A_1 and A_2 . Then, substituting $6\delta/(\pi^2 T^2)$ into δ_1 and δ_2 and taking the union bounds for all $T \geq 1$, we conclude the proof. \square

B.2 Proof of Theorem 4.4

We show the following theorem:

Theorem 4.4. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and Assumption 2.1 holds. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2d \log(bdr t^2(\sqrt{\log(ad)} + \sqrt{\pi}/2)) - 2 \log 2$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , the conditional expected regret can be bounded with probability at least $1 - \delta$ as follows:*

$$\Pr \left(\forall T \geq 1, \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right] \leq U(T, \delta) \right) \geq 1 - \delta,$$

where

$$U(T, \delta) = \frac{\pi^2}{6} + 6\sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right)} + \sqrt{C_1 \gamma_T \left(T s_T + T + 2\sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right)} + 2 \log \left(\frac{\pi^2 T^2}{3\delta} \right) \right)},$$

and $C_1 := 2/\log(1 + \sigma^{-2})$.

PROOF. We consider a finite discretized input set $\mathcal{X}_t \subset \mathcal{X}$ for $t \geq 1$ for theoretical analysis, though \mathcal{X}_t is not related to the actual algorithm. Let $\mathcal{X}_t = \{\frac{r}{\tau_{t+1}}, \dots, \frac{\tau_t r}{\tau_{t+1}}\}^d \subset \mathcal{X}$, where each coordinate takes values from a uniform grid of $\tau_t = bdr u_t (\sqrt{\log(ad)} + \sqrt{\pi}/2)$ points. Thus, $|\mathcal{X}_t| = \tau_t^d$. In addition, we define $[\mathbf{x}]_t$ as the nearest point in \mathcal{X}_t of $\mathbf{x} \in \mathcal{X}$.

As with the proof of Theorem 4.2, we obtain

$$\begin{aligned} & \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right] \\ & \leq \mathbb{E}_f \left[\sum_{t=1}^T f(\mathbf{x}^*) - f([\mathbf{x}^*]_t) \right] + \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} \left[\sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{X}_t} f(\mathbf{x}) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right]. \end{aligned}$$

Then, the first term can be bounded from above by $\pi^2/6$ as in the proof of Theorem 4.2. For the second term, we see that

$$\begin{aligned} & \mathbb{E}_{f, \{\epsilon_t\}_{t \geq 1}} \left[\sum_{t=1}^T \max_{\mathbf{x} \in \mathcal{X}_t} f(\mathbf{x}) - f(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right] \\ & = \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} \left[\mathbb{E}_{f \mid \mathcal{D}_{t-1}} \left[\max_{\mathbf{x} \in \mathcal{X}_t} f(\mathbf{x}) \right] - \mathbb{E}_{\zeta_t} \left[\max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) \right] + \mathbb{E}_{\zeta_t} \left[\max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) \right] - \mathbb{E}_{\zeta_t} [v_t(\mathbf{x}_t)] \right. \\ & \quad \left. + \mathbb{E}_{\zeta_t} [v_t(\mathbf{x}_t)] - v_t(\mathbf{x}_t) + v_t(\mathbf{x}_t) - \mu_{t-1}(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right] \\ & \leq \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} \left[\mathbb{E}_{\zeta_t} [v_t(\mathbf{x}_t)] - v_t(\mathbf{x}_t) + \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \mid \{\zeta_t\}_{t \geq 1} \right], \end{aligned}$$

where we use

$$\begin{aligned} \mathbb{E}_{f \mid \mathcal{D}_{t-1}} \left[\max_{\mathbf{x} \in \mathcal{X}_t} f(\mathbf{x}) \right] & \leq \mathbb{E}_{\zeta_t} \left[\max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) \right], & (\because \text{Lemma 4.2}) \\ \mathbb{E}_{\zeta_t} \left[\max_{\mathbf{x} \in \mathcal{X}_t} v_t(\mathbf{x}) \right] & \leq \mathbb{E}_{\zeta_t} [v_t(\mathbf{x}_t)]. & (\because \mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} v_t(\mathbf{x})) \end{aligned}$$

Then, the rest of the proof is almost the same as Theorem 4.3. \square

C Proofs for High-Probability Regret Bounds of IRGP-UCB

Here, we show the proof of the following Theorem 4.5:

Theorem 4.5. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Pick $\delta \in (0, 1)$. Assume that $\zeta_t = 2 \log(1/U_t)$, where $\{U_t\}_{t \geq 1}$ is a sequence of mutually independent random variables that satisfy $\mathbb{E}[U_t] = \frac{6\delta}{|\mathcal{X}|t^2\pi^2}$ and $\Pr(U_t \in (0, 1)) = 1$ for all $t \geq 1$. Then, by running the randomized GP-UCB with ζ_t , the cumulative regret can be bounded with probability at least $1 - \delta$ as follows:*

$$\Pr \left(\forall T \geq 1, R_T \leq 2\sqrt{C_1 \gamma_T \sum_{t=1}^T \zeta_t} \right) \geq 1 - \delta,$$

where $C_1 := 2/\log(1 + \sigma^{-2})$.

PROOF. We can transform the regret as follows:

$$\begin{aligned}
R_T &= \sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \\
&= \sum_{t=1}^T f(\mathbf{x}^*) - \mu_{t-1}(\mathbf{x}_t) + \mu_{t-1}(\mathbf{x}_t) - f(\mathbf{x}_t) \\
&\leq 2 \sum_{t=1}^T \zeta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) && (\because \text{Lemma 4.3}) \\
&\leq 2 \sqrt{\sum_{t=1}^T \zeta_t \sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}_t)} && (\because \text{Cauchy-Schwartz inequality}) \\
&\leq 2 \sqrt{C_1 \gamma_T \sum_{t=1}^T \zeta_t}. && (\because \text{Lemma 5.4 in Srinivas et al. (2010)})
\end{aligned}$$

□

Next, we show the following corollary:

Corollary 4.1. *Let $f \sim \mathcal{GP}(0, k)$, where k is a positive semidefinite kernel and $k(\mathbf{x}, \mathbf{x}) \leq 1$, and X be finite. Pick $\delta \in (0, 1)$. Assume that ζ_t follows a shifted exponential distribution with $s_t = 2 \log(|X|t^2\pi^2/(6\delta))$ and $\lambda = 1/2$ for any $t \geq 1$. Then, by running IRGP-UCB with ζ_t , the cumulative regret can be bounded with probability at least $1 - \delta$ as follows:*

$$\Pr \left(\forall T \geq 1, R_T \leq 2 \sqrt{C_1 \gamma_T \left(T s_T + T + 2 \sqrt{T \log \left(\frac{\pi^2 T^2}{3\delta} \right)} + 2 \log \left(\frac{\pi^2 T^2}{3\delta} \right) \right)} \right) \geq 1 - \delta,$$

where $C_1 := 2/\log(1 + \sigma^{-2})$.

PROOF. From the assumption that $s_t = 2 \log(|X|t^2\pi^2/(6\delta))$ and $\lambda = 1/2$, we obtain that $U_t \sim \text{Uni}(0, \frac{6\delta}{|X|t^2\pi^2})$ and $\mathbb{E}[U_t] = \frac{3\delta}{|X|t^2\pi^2}$, where $\zeta_t = 2 \log(1/U_t)$ as in Lemma 4.3. Moreover, from Theorem 4.5, we obtain

$$\Pr \left(\forall T \geq 1, R_T \leq 2 \sqrt{C_1 \gamma_T \sum_{t=1}^T \zeta_t} \right) \geq 1 - \frac{\delta}{2}.$$

Furthermore, since $\sum_{t=1}^T \{\zeta_t - s_t\}$ follows a chi-square distribution with T degrees of freedom, using the union bound and Lemma E.4, we can obtain

$$\Pr \left(\forall T \geq 1, \sum_{t=1}^T \zeta_t \leq T s_T + T + 2 \sqrt{T \log \frac{\pi^2 T^2}{3\delta}} + 2 \log \frac{\pi^2 T^2}{3\delta} \right) \geq 1 - \frac{\delta}{2}.$$

Therefore, taking the union bound concludes the proof. □

D Proof for Expected Regret Bound of GP-UCB with Constant Confidence Parameter

First, we derive the following lemma to show Theorem 4.6:

Lemma D.1. *Let $\{\epsilon_i\}_{i \geq 1}$ be a mutually independent sequence of standard normal random variables, that is, $\epsilon_i \sim \mathcal{N}(0, 1)$ for all $i \geq 1$. Then, there exists a positive constant C such that*

$$\Pr\left(\forall t \leq T, \frac{\sum_{i=1}^t \epsilon_i}{t} \geq -1\right) \geq C.$$

PROOF. From the properties of the probability and the union bound, we can see that

$$\begin{aligned} \Pr\left(\forall t \leq T, \frac{\sum_{i=1}^t \epsilon_i}{t} \geq -1\right) &= 1 - \Pr\left(\exists t \leq T, \frac{\sum_{i=1}^t \epsilon_i}{t} < -1\right) \\ &\geq 1 - \sum_{t=1}^T \Pr\left(\frac{\sum_{i=1}^t \epsilon_i}{t} < -1\right) \\ &= 1 - \sum_{t=1}^T \Pr\left(\frac{\sum_{i=1}^t \epsilon_i}{t} > 1\right), \end{aligned}$$

where the last equality is obtained from the symmetry of the Gaussian distribution. Since $\epsilon_i \sim \mathcal{N}(0, 1)$ for all $i \geq 1$, the average $\sum_{i=1}^t \epsilon_i/t \sim \mathcal{N}(0, 1/t)$. Therefore, by using Lemma E.3, we obtain that

$$\Pr\left(\frac{\sum_{i=1}^t \epsilon_i}{t} > 1\right) \leq \frac{1}{2} \exp\left(-\frac{t}{2}\right).$$

Since $\left(\frac{1}{2} \exp\left(-\frac{t}{2}\right)\right)_{t \geq 1}$ is a geometric series with common ratio $\exp(-\frac{1}{2})$ and initial term $\frac{1}{2} \exp(-\frac{1}{2})$, its sum of infinite series converges as follows:

$$\begin{aligned} \sum_{t=1}^T \Pr\left(\frac{\sum_{i=1}^t \epsilon_i}{t} > 1\right) &\leq \sum_{t=1}^T \frac{1}{2} \exp\left(-\frac{t}{2}\right) \\ &\leq \sum_{t=1}^{\infty} \frac{1}{2} \exp\left(-\frac{t}{2}\right) \\ &= \frac{1}{2 \exp(1/2)(1 - \exp(-1/2))} \approx 0.771. \end{aligned}$$

Consequently, we can obtain the following lower bound of the probability:

$$\begin{aligned} \Pr\left(\forall t \leq T, \frac{\sum_{i=1}^t \epsilon_i}{t} \geq -1\right) &\geq 1 - \sum_{t=1}^T \Pr\left(\frac{\sum_{i=1}^t \epsilon_i}{t} > 1\right) \\ &\geq 1 - \frac{1}{2 \exp(1/2)(1 - \exp(-1/2))} \approx 0.229. \end{aligned}$$

□

Then, we show the following theorem:

Theorem 4.6. *Assume that $\mathcal{X} = \{\mathbf{x}^{(1)}, \mathbf{x}^{(2)}\}$, $\epsilon_t \sim \mathcal{N}(0, 1)$ for all $t \geq 1$, and*

$$\begin{pmatrix} f(\mathbf{x}^{(1)}) \\ f(\mathbf{x}^{(2)}) \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 0.99 \end{pmatrix}\right),$$

where $\rho < 1$ is a covariance parameter. Then, if GP-UCB with any constant confidence parameter $\beta_t^{(1/2)} = c > 0$ runs, then BCR grows linearly, that is,

$$\text{BCR}_T = \mathbb{E} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \right] = \Omega(T).$$

PROOF. We consider the following event:

$$E_T = \left\{ f(\mathbf{x}^{(1)}) \geq \frac{2 \max\{1, c\}}{1 - \rho} + 1, f(\mathbf{x}^{(2)}) > f(\mathbf{x}^{(1)}) + 1, \text{ and } \sum_{i=1}^t \epsilon_i/t \geq -1 \text{ for all } t \geq 1 \right\}.$$

From Lemma D.1 and the independentness between f and $(\epsilon_i)_{i \geq 1}$, the probability of this event $\Pr(E_T) > 0$. Then, the BCR can be bounded from below as follows:

$$\text{BCR}_T = \mathbb{E} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \right] \quad (25)$$

$$= \mathbb{E} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid E_T \right] \Pr(E_T) + \mathbb{E} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid E_T^c \right] \Pr(E_T^c) \quad (26)$$

$$\geq \mathbb{E} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid E_T \right] \Pr(E_T), \quad (27)$$

where E_T^c is a complementary event of E_T .

We show that, when E_T holds, $\mathbf{x}_t = \mathbf{x}^{(1)}$ for all $t \leq T$. First, since the variance of $\mathbf{x}^{(1)}$ is larger than that of $\mathbf{x}^{(2)}$ in the prior distribution, the first chosen input $\mathbf{x}_1 = \mathbf{x}^{(1)}$. Then, we show $\mathbf{x}_{t+1} = \mathbf{x}^{(1)}$ for all $t \geq 1$. Fix $t \geq 1$ and assume that $\mathbf{x}_i = \mathbf{x}^{(1)}$ for all $i \leq t$. The posterior mean can be derived as follows:

$$\begin{aligned} \mu_t(\mathbf{x}^{(1)}) &= \mathbf{1}_t^\top (\mathbf{1}_t \mathbf{1}_t^\top + \mathbf{I}_t)^{-1} \mathbf{y}_t, \\ \mu_t(\mathbf{x}^{(2)}) &= \rho \mathbf{1}_t^\top (\mathbf{1}_t \mathbf{1}_t^\top + \mathbf{I}_t)^{-1} \mathbf{y}_t, \end{aligned}$$

where $\mathbf{1}_t = (1, \dots, 1)^\top \in \mathbb{R}^t$ and $\mathbf{I}_t \in \mathbb{R}^{t \times t}$ is the identity matrix. From the Woodbury formula, we obtain

$$(\mathbf{1}_t \mathbf{1}_t^\top + \mathbf{I}_t)^{-1} = \mathbf{I}_t - \frac{\mathbf{1}_t \mathbf{1}_t^\top}{t+1}.$$

Therefore, the posterior means can be transformed as

$$\begin{aligned} \mu_t(\mathbf{x}^{(1)}) &= \sum_{i=1}^t y_i - \frac{t}{t+1} \sum_{i=1}^t y_i = \frac{t}{t+1} \frac{\sum_{i=1}^t y_i}{t}, \\ \mu_t(\mathbf{x}^{(2)}) &= \frac{\rho t}{t+1} \frac{\sum_{i=1}^t y_i}{t}. \end{aligned}$$

Thus, the difference between the posterior means is bounded from below as

$$\begin{aligned}
 \mu_t(\mathbf{x}^{(1)}) - \mu_t(\mathbf{x}^{(2)}) &= (1 - \rho) \frac{t}{t+1} \frac{\sum_{i=1}^t y_i}{t} \\
 &\geq \frac{1 - \rho}{2} \frac{tf(\mathbf{x}^{(1)}) + \sum_{i=1}^t \epsilon_i}{t} && (\because t/(t+1) \geq 1/2), \\
 &\geq \frac{1 - \rho}{2} \frac{\frac{2t \max\{1, c\}}{1 - \rho} + 1 - 1}{t} && (\because E_T \text{ holds}) \\
 &= \max\{1, c\} \\
 &\geq c \\
 &= \beta_t^{(1/2)} \\
 &\geq \beta_t^{(1/2)} \sigma_t(\mathbf{x}^{(2)}) && (\because \sigma_t(\mathbf{x}^{(2)}) \leq 1) \\
 &\geq \beta_t^{(1/2)} \sigma_t(\mathbf{x}^{(2)}) - \beta_t^{(1/2)} \sigma_t(\mathbf{x}^{(1)}).
 \end{aligned}$$

Hence, we see that $\mathbf{x}_{t+1} = \mathbf{x}^{(1)}$ since $\mu_t(\mathbf{x}^{(1)}) + \beta_t^{(1/2)} \sigma_t(\mathbf{x}^{(1)}) \geq \mu_t(\mathbf{x}^{(2)}) + \beta_t^{(1/2)} \sigma_t(\mathbf{x}^{(2)})$.

Since E_T holds, the optimal point $\mathbf{x}^* = \mathbf{x}^{(2)}$. However, as shown above, $\mathbf{x}_t = \mathbf{x}^{(1)}$ for all $t \geq 1$ if E_T holds. Therefore, $\mathbb{E} \left[\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \mid E_T \right] \geq T$. Consequently, since $\Pr(E_T) > 0$, we can conclude that

$$\text{BCR}_T = \Omega(T).$$

□

E Auxiliary Lemmas

For convenience, we here show the assumption again:

Assumption 2.1. Let $\mathcal{X} \subset [0, r]^d$ be a compact and convex set, where $r > 0$. Assume that the kernel k satisfies the following condition on the derivatives of sample path f . There exist the constants $a, b > 0$ such that,

$$\Pr \left(\sup_{\mathbf{x} \in \mathcal{X}} \left| \frac{\partial f}{\partial x_j} \right| > L \right) \leq a \exp \left(- \left(\frac{L}{b} \right)^2 \right), \text{ for } j \in [d],$$

where $[d] = \{1, \dots, d\}$.

Then, from Assumption 2.1, we obtain several lemmas. First, we show an upper bound of the supremum of partial derivatives. This result is tighter than Lemma 12 in Kandasamy et al. (2018). Note that, in the proof of Lemma 12 in Kandasamy et al. (2018), there is a typo that $\Pr(L \geq t)$ is bounded above by $a \exp\{t^2/b^2\}$, which should be $ad \exp\{-t^2/b^2\}$.

Lemma E.1. Let $f \sim \mathcal{GP}(0, k)$ and Assumption 2.1 holds. Let the supremum of the partial derivatives $L_{\max} := \sup_{j \in [d]} \sup_{\mathbf{x} \in \mathcal{X}} \left| \frac{\partial f}{\partial x_j} \right|$. Then, $\mathbb{E}[L_{\max}]$ can be bounded above as follows:

$$\mathbb{E}[L_{\max}] \leq b(\sqrt{\log(ad)} + \sqrt{\pi}/2). \quad (28)$$

PROOF. From the assumption, using the union bound, we obtain a bound on the probability,

$$\Pr(L_{\max} > L) \leq \sum_{j=1}^d a \exp \left(- \left(\frac{L}{b} \right)^2 \right) = ad \exp \left(- \left(\frac{L}{b} \right)^2 \right). \quad (29)$$

Then, using expectation integral identity, we can bound $\mathbb{E}[L_{\max}]$ as follows:

$$\mathbb{E}[L_{\max}] = \int_0^{\infty} \Pr(L_{\max} > L) dL \quad (30)$$

$$\leq \int_0^{\infty} \min\{1, ade^{-(L/b)^2}\} dL \quad (\because \text{Eq. (29)}) \quad (31)$$

$$= b\sqrt{\log(ad)} + \int_{b\sqrt{\log(ad)}}^{\infty} ade^{-(L/b)^2} dL \quad (32)$$

$$= b\sqrt{\log(ad)} + abd\sqrt{\pi} \int_{b\sqrt{\log(ad)}}^{\infty} \frac{1}{\sqrt{2\pi}(b^2/2)} e^{-(L/b)^2} dL \quad (33)$$

$$= b\sqrt{\log(ad)} + abd\sqrt{\pi} \left(1 - \Phi \left(\frac{b\sqrt{\log(ad)}}{b/\sqrt{2}} \right) \right) \quad (34)$$

$$\leq b\sqrt{\log(ad)} + \frac{b\sqrt{\pi}}{2}, \quad (\because \text{Lemma E.3}) \quad (35)$$

where Φ is a cumulative distribution function of the standard normal distribution. \square

Next, we show the bound for the difference of discretized outputs:

Lemma E.2. *Let $f \sim \mathcal{GP}(0, k)$ and Assumption 2.1 holds. Let $\mathcal{X}_t = \{\frac{r}{\tau_{t+1}}, \dots, \frac{\tau_t r}{\tau_{t+1}}\}^d \subset \mathcal{X}$, where each coordinate takes values from a uniform grid of $\tau_t = bdru_t(\sqrt{\log(ad)} + \sqrt{\pi}/2)$ points. Then, we can bound the expectation of differences,*

$$\sum_{t=1}^T \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - f([\mathbf{x}]_t) \right] \leq \sum_{t=1}^T \frac{1}{u_t}, \quad (36)$$

where $[\mathbf{x}]_t$ is the nearest point in \mathcal{X}_t of $\mathbf{x} \in \mathcal{X}$.

PROOF. From the construction of \mathcal{X}_t , we can obtain the upper bound of L1 distance between \mathbf{x} and $[\mathbf{x}]_t$ as follows:

$$\sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x} - [\mathbf{x}]_t\|_1 \leq \frac{dr}{bdru_t(\sqrt{\log(ad)} + \sqrt{\pi}/2)} \quad (37)$$

$$= \frac{1}{bu_t(\sqrt{\log(ad)} + \sqrt{\pi}/2)}. \quad (38)$$

Note that this discretization does not depend on any randomness and is fixed beforehand.

Then, we obtain the following:

$$\sum_{t=1}^T \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - f([\mathbf{x}]_t) \right] \leq \sum_{t=1}^T \mathbb{E} \left[L_{\max} \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x} - [\mathbf{x}]_t\|_1 \right] \quad (39)$$

$$\leq \sum_{t=1}^T \mathbb{E} [L_{\max}] \frac{1}{bu_t(\sqrt{\log(ad)} + \sqrt{\pi}/2)} \quad (\because \text{Eq. (38)}) \quad (40)$$

$$\leq \sum_{t=1}^T \frac{1}{u_t}. \quad (\because \text{Lemma E.1}) \quad (41)$$

\square

We used the following useful lemma:

Lemma E.3 (in Lemma 5.2 of Srinivas et al. (2010)). *For $c > 0$, the survival function of the standard normal distribution can be bounded above as follows:*

$$1 - \Phi(c) \leq \frac{1}{2} \exp(-c^2/2). \quad (42)$$

Lemma E.4 (Eq. (4.3) of Laurent and Massart (2000)). *Pick $\delta \in (0, 1)$. Suppose that X follows a chi-square distribution with D degrees of freedom. Then, the following upper bound for the tail probability holds:*

$$\Pr\left(X \geq D + 2\sqrt{Dc} + 2c\right) \leq \exp(-c),$$

for any $c > 0$. Therefore, the following inequality holds with probability at least $1 - \delta$:

$$X \leq D + 2\sqrt{D \log(1/\delta)} + 2 \log(1/\delta).$$

Received 03 June 2025; accepted 13 October 2025