

Regret Analysis of Posterior Sampling-Based Expected Improvement for Bayesian Optimization

Anonymous authors
Paper under double-blind review

Abstract

Bayesian optimization is a powerful tool for optimizing an expensive-to-evaluate black-box function. In particular, the effectiveness of expected improvement (EI) has been demonstrated in a wide range of applications. However, theoretical analyses of EI are limited compared with other theoretically established algorithms. This paper analyzes a randomized variant of EI, which evaluates the EI from the maximum of the posterior sample path. We show that this posterior sampling-based random EI achieves the sublinear Bayesian cumulative regret bounds under the assumption that the black-box function follows a Gaussian process. Finally, we demonstrate the effectiveness of the proposed method through numerical experiments.

1 Introduction

Bayesian optimization (BO) (Shahriari et al., 2016) is a powerful tool for optimizing expensive-to-evaluate black-box functions. BO aims for optimization with fewer observations due to the expensive evaluation cost of the objective function. For this purpose, BO sequentially queries a candidate determined by maximizing an acquisition function (AF) based on a Bayesian model. A Gaussian process (GP) model (Rasmussen & Williams, 2005) is typically employed in BO. BO has been applied to various fields, such as materials informatics (Ueno et al., 2016), robotics (Berkenkamp et al., 2023), and hyperparameter tuning (Snoek et al., 2012).

Several widely used BO algorithms are based on an improvement from a reference value, which is often set as the current best observation. The GP-based probability of improvement (GP-PI) algorithm (Kushner, 1964) evaluates the AF defined as the probability that a new observation exceeds the current best observation. However, it is known that GP-PI is too exploitative depending on the reference value (Jones, 2001; Shahriari et al., 2016). Therefore, GP-based expected improvement (GP-EI) (Mockus et al., 1978; Jones et al., 1998), which evaluates the expectation of improvement from the current best observation, has been frequently used. The effectiveness of GP-EI has been repeatedly demonstrated not only in the usual BO (Snoek et al., 2012; Ament et al., 2023) but also in several extended BO settings, such as multi-fidelity BO (Huang et al., 2006), constrained BO (Gardner et al., 2014; Gelbart et al., 2014), multi-objective BO (Emmerich et al., 2011), and BO with large noises (Letham et al., 2019). However, since the usual GP-EI still depends on the current best observation, which can include noise, its heuristic hyperparameter tuning approach has been discussed; for example, see Section 2.3.2 in (Brochu et al., 2010) and (Lizotte, 2008). Hence, GP-EI is often regarded as a heuristic method with strong empirical performance.

Other lines of study have tackled theoretical guarantees on regret (for example, Srinivas et al., 2010; Russo & Van Roy, 2014). Seminal work (Srinivas et al., 2010) showed the cumulative regret upper bound for the GP upper confidence bound (GP-UCB) algorithm. Russo & Van Roy (2014); Chowdhury & Gopalan (2017) derived the cumulative regret upper bounds for the GP-based Thompson sampling algorithm (GP-TS). Furthermore, Scarlett (2018) showed the algorithm-independent regret lower bound and a tighter regret upper bound under several conditions. Recently, Takeno et al. (2024) showed that the variant of GP-PI called GP-PI from the maximum of sample path (GP-PIMS) achieves a sublinear Bayesian cumulative regret (BCR) upper bound. GP-PIMS avoids the problem of selecting the reference value by replacing it

with the maximum of the posterior sample path, inspired by GP-TS. Furthermore, a recent study (Iwazaki, 2025) showed a tighter high-probability regret upper bound for the usual GP-UCB.

Several studies (Bull, 2011; Wang & de Freitas, 2014; Nguyen et al., 2017; Berk et al., 2019) have attempted to obtain regret guarantees of GP-EI-based algorithms, in which the reference value is often changed from the current best observation to, for example, the maximum of the posterior mean (Wang & de Freitas, 2014), and the maximum of the posterior sample path (Berk et al., 2019). However, as reported in existing studies (Tran-The et al., 2022; Bedi et al., 2022; Hu et al., 2025), several analyses appear incorrect (See Section 3 and Appendix A for more details). Other theoretical analyses in (Wang & de Freitas, 2014; Tran-The et al., 2022) require rescaling the posterior variance, which can deteriorate the practical optimization performance. Hence, a theoretical analysis for GP-EI-based algorithms without posterior variance rescaling has not been established except for the noiseless case (Bull, 2011).

This paper analyzes the posterior sampling-based GP-EI algorithm inspired by GP-PIMS (Takeno et al., 2024), which is GP-EI with the reference value set by the maximum of the posterior sample path. We refer to this algorithm as GP-EI from the maximum of the sample path (GP-EIMS). We show that GP-EIMS achieves the sublinear BCR under the Bayesian setting, where the objective function follows the GP. Moreover, GP-EIMS does not require rescaling of the posterior variance. Finally, we demonstrate the effectiveness of GP-EIMS, which stems from avoiding posterior variance rescaling, through numerical experiments.

2 Preliminary

Bayesian Optimization: We consider the black-box optimization for $f : \mathcal{X} \rightarrow \mathbb{R}$ formulated as

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

where $\mathcal{X} \subset \mathbb{R}^d$ is an input domain and d is a dimension. We assume that only observations $y = f(\mathbf{x}) + \epsilon$, where ϵ is a noise, can be obtained, but are costly. Thus, our goal is to optimize f with fewer observations. For this goal, BO sequentially performs querying an input \mathbf{x}_t and observing $y_t = f(\mathbf{x}_t) + \epsilon_t$, where t is an iteration. The AF chooses queried input based on the Bayesian model, which is updated by the collected dataset.

Gaussian Process Model: In this paper, we assume that f follows a GP with a predefined kernel $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, denoted as $f \sim \mathcal{GP}(0, k)$. We further assume that the observation is contaminated by Gaussian noise; that is, $y_t = f(\mathbf{x}_t) + \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(0, \sigma^2)$ with a positive variance $\sigma^2 > 0$. Let $\mathcal{D}_{t-1} = \{(\mathbf{x}_i, y_i)\}_{i=1}^{t-1}$ be the training data obtained until the beginning of t -th iteration. Then, the posterior distribution $p(f \mid \mathcal{D}_{t-1})$ is a GP again (Rasmussen & Williams, 2005), whose mean and variance can be obtained as follows:

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

$$\sigma_{t-1}^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_{t-1}(\mathbf{x})^\top (\mathbf{K} + \sigma^2 \mathbf{I}_{t-1})^{-1} \mathbf{k}_{t-1}(\mathbf{x}), \quad (2)$$

where $\mathbf{k}_{t-1}(\mathbf{x}) := (k(\mathbf{x}, \mathbf{x}_1), \dots, k(\mathbf{x}, \mathbf{x}_{t-1}))^\top \in \mathbb{R}^{t-1}$ is the kernel vector, $\mathbf{K} \in \mathbb{R}^{(t-1) \times (t-1)}$ 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 a probability density function (PDF) $p(\cdot \mid \mathcal{D}_{t-1}) = p_t(\cdot)$, a probability $\Pr(\cdot \mid \mathcal{D}_{t-1}) = \Pr_t(\cdot)$, and an expectation $\mathbb{E}[\cdot \mid \mathcal{D}_{t-1}] = \mathbb{E}_t[\cdot]$ for brevity.

Furthermore, for the case of continuous \mathcal{X} , we assume the following smoothness condition for GP:

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 $a \geq 1$ and $b > 0$ such that,

$$\Pr(L_j > c) \leq a \exp\left(-\left(\frac{c}{b}\right)^2\right), \text{ for } j \in [d], \quad (3)$$

where $L_j = \sup_{\mathbf{x} \in \mathcal{X}} \left| \frac{\partial f}{\partial x_j} \right|$ and $[d] = \{1, \dots, d\}$.

This assumption is commonly used (Srinivas et al., 2010; Kandasamy et al., 2018; Takeno et al., 2023; 2024) and holds at least for the squared exponential (SE) and Matérn- ν kernels with $\nu > 2$ (Theorem 5 in Ghosal & Roy, 2006; Srinivas et al., 2010).

Maximum Information Gain: For regret analysis, we will use the quantity called maximum information gain (MIG) (Srinivas et al., 2010; Vakili et al., 2021):

Definition 2.2 (Maximum information gain). Let $f \sim \mathcal{GP}(0, k)$ over $\mathcal{X} \subset [0, r]^d$. Let $A = \{\mathbf{a}_i\}_{i=1}^T$, where $\mathbf{a}_i \in \mathcal{X}$ for all $i \in [T]$. Let $\mathbf{f}_A = (f(\mathbf{a}_i))_{i=1}^T$, $\boldsymbol{\epsilon}_A = (\epsilon_i)_{i=1}^T$, where $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ for all $i \in [T]$, and $\mathbf{y}_A = \mathbf{f}_A + \boldsymbol{\epsilon}_A \in \mathbb{R}^T$. Then, MIG γ_T is defined as follows:

$$\gamma_T := \max_A I(\mathbf{y}_A; \mathbf{f}_A) \text{ such that } |A| = T \text{ and } \mathbf{a}_i \in \mathcal{X} \text{ for all } i \in [T], \quad (4)$$

where I is the Shannon mutual information.

For frequently used kernels, MIG is known to be sublinear (Srinivas et al., 2010; Vakili et al., 2021), for example, $\gamma_T = O((\log T)^{d+1})$ for the SE kernels and $\gamma_T = O(T^{\frac{d}{2\nu+d}} (\log T)^{\frac{2\nu}{2\nu+d}})$ for the Matérn- ν kernels.

Performance Measure: We evaluate the theoretical performance of BO methods by the BCR (Russo & Van Roy, 2014; Kandasamy et al., 2018; Takeno et al., 2023; 2024) defined as

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

where the expectation is taken with all the randomness, that is, $f, (\epsilon_i)_{i \in [T]}$, and the randomness of the algorithm. Our goal is to show the sublinearity of the BCR since it implies that the simple regret $\mathbb{E}[f(\mathbf{x}^*) - f(\hat{\mathbf{x}}_T)] \leq \text{BCR}_T/T \rightarrow 0$ as $T \rightarrow \infty$ (Proposition 8 in Russo & Van Roy, 2018), where $\hat{\mathbf{x}}_T := \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_T(\mathbf{x})$ is a recommendation point at T -th iteration.

GP-based EI algorithm: GP-EI is the widely used AF for BO (Mockus et al., 1978; Jones et al., 1998). As the name suggests, the AF of GP-EI given \mathcal{D}_{t-1} at t -th iteration is defined as follows:

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}(\mathbf{x}), y_{t-1}^{\max}) = \mathbb{E}_t[\max\{f(\mathbf{x}) - y_{t-1}^{\max}, 0\}] \quad (6)$$

$$= \begin{cases} \sigma_{t-1}(\mathbf{x})\tau\left(\frac{\mu_{t-1}(\mathbf{x}) - y_{t-1}^{\max}}{\sigma_{t-1}(\mathbf{x})}\right) & \text{if } \sigma_{t-1}(\mathbf{x}) > 0, \\ \max\{\mu_{t-1}(\mathbf{x}) - y_{t-1}^{\max}, 0\} & \text{if } \sigma_{t-1}(\mathbf{x}) = 0, \end{cases} \quad (7)$$

where $y_{t-1}^{\max} = \max_{i \in [t-1]} y_i$ is the current best observation, $\tau: \mathbb{R} \rightarrow \mathbb{R}^+$ is

$$\tau(c) = c\Phi(c) + \phi(c), \quad (8)$$

and Φ and ϕ are the cumulative distribution function and the PDF of the standard Gaussian distribution. Mainly due to the noise included in y_{t-1}^{\max} , a regret analysis for the original GP-EI is difficult.

3 Literature Review for Regret Analysis of GP-EI

The assumption for the regret analysis in the BO literature is twofold: the Bayesian setting that f follows GPs (Srinivas et al., 2010; Russo & Van Roy, 2014; Takeno et al., 2023; 2024) and the frequentist setting that f belongs to the known reproducing kernel Hilbert space (RKHS) (Srinivas et al., 2010; Chowdhury & Gopalan, 2017; Janz et al., 2020; Iwazaki & Takeno, 2025). Furthermore, the analyzed regret can be categorized as cumulative regret $\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t)$ and simple regret $f(\mathbf{x}^*) - f(\hat{\mathbf{x}}_t)$ with some recommended input at t -th iteration, for example, $\hat{\mathbf{x}}_t := \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_t(\mathbf{x})$. In addition, as we discussed above, we refer to its expected variants as BCR and Bayesian simple regret in the Bayesian setting. Although our main analysis concentrates on the BCR analysis in the Bayesian setting, we review the analyses of GP-EI for both regrets in both settings.

3.1 Regret Analysis for Noise-Free Setting

Seminal works for the regret analysis of GP-EI concentrated on a noise-free setting where observation noise does not exist. Therefore, there is no need to care about the noise. Vazquez & Bect (2010) show that GP-EI asymptotically converges on the optimum for both Bayesian and frequentist settings regarding the stationary and *non-degeneracy* kernels, such as the Matérn kernel family. Grünewälder et al. (2010) proves that the computationally intractable T -step look-ahead variant of GP-EI is near-optimal in the noise-free Bayesian setting under an assumption that the prior mean function and the kernel function are Hölder-continuous. Here, near-optimal means that a regret upper bound matches a regret lower bound except for logarithmic factors. Bull (2011) further shows simple regret upper bounds of GP-EI and modified GP-EI algorithms for the Matérn kernel family in the frequentist setting. In particular, the modified GP-EI is shown to be near-optimal.

3.2 Regret Analysis with Rescaling Posterior Variance

Wang & de Freitas (2014) show the high-probability cumulative regret upper bound in the frequentist setting for a case where additional noise contaminates observations, as with our setting. For this analysis, Wang & de Freitas (2014) propose to use $\mu_{t-1}^{\max} = \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x})$ instead of y_{t-1}^{\max} and to scale $\sigma_{t-1}(\mathbf{x})$ by $\nu_t^{1/2} > 0$ in the AF of GP-EI as follows:

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}), \mu_{t-1}^{\max}) = \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}) \tau \left(\frac{\mu_{t-1}(\mathbf{x}) - \mu_{t-1}^{\max}}{\nu_t^{1/2} \sigma_{t-1}(\mathbf{x})} \right), \quad (9)$$

where $\nu_t = \Theta(\gamma_t)$. Their analysis shows $O(\gamma_T \sqrt{T})^1$ cumulative regret upper bound, which is sublinear, for example, for the Gaussian kernel. However, since $\nu_t \gg 1$ in most cases, the above rescaling strengthens the exploration behavior, and the practical effectiveness can deteriorate. As a result, in the practical scenarios, we must tune ν_t as a hyperparameter. Note that although Wang & de Freitas (2014) further consider a case where hyperparameters of GPs are unknown, we focus on a case where the hyperparameters of GPs are specified beforehand.

Tran-The et al. (2022) analyze $\text{EI}(\mu_{t-1}(\mathbf{x}), \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}), \tilde{\mu}_{t-1}^{\max})$, which avoids an optimization over (continuous) \mathcal{X} by replacing μ_{t-1}^{\max} in (Wang & de Freitas, 2014) as $\tilde{\mu}_{t-1}^{\max} = \max_{i \in [t-1]} \mu_t(\mathbf{x}_i)$, in the frequentist setting. However, Tran-The et al. (2022) only analyze the simple regret $f(\mathbf{x}^*) - f(\hat{\mathbf{x}}_t) = O(\gamma_T / \sqrt{T})$. Therefore, even in the frequentist setting, the cumulative regret upper bound of the GP-EI algorithm with $\tilde{\mu}_{t-1}^{\max}$ has not been shown to our knowledge. Furthermore, their algorithm suffers from the same problem as (Wang & de Freitas, 2014), that is, the rescaling of posterior variance.

Although Wang & de Freitas (2014); Tran-The et al. (2022) discuss only the frequentist setting, their derivation can be extended to the Bayesian setting by carefully applying well-known proof techniques mainly from (Russo & Van Roy, 2014; Kandasamy et al., 2018; Takeno et al., 2023). For completeness, we show high probability cumulative regret upper bounds and BCR upper bounds in the Bayesian setting of $\text{EI}(\mu_{t-1}(\mathbf{x}), \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}), \mu_{t-1}^{\max})$ and $\text{EI}(\mu_{t-1}(\mathbf{x}), \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}), \tilde{\mu}_{t-1}^{\max})$ in Appendices B.2 and B.3, respectively. Since our analysis for finite input domains can apply almost similarly to the frequentist setting, the cumulative regret incurred by $\text{EI}(\mu_{t-1}(\mathbf{x}), \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}), \tilde{\mu}_{t-1}^{\max})$ is $O(\gamma_T \sqrt{T})$ in the frequentist setting (see Appendix B.4 for details).

Hu et al. (2025) recently analyze the GP-EI-based algorithm combining $\text{EI}(\mu_{t-1}(\mathbf{x}), \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}), \tilde{\mu}_{t-1}^{\max})$ and a candidate elimination based on a quantity $\text{EI}(\tilde{\mu}_{t-1}^{\max}, \nu_t^{1/2} \sigma_{t-1}(\mathbf{x}), \mu_{t-1}(\mathbf{x})) / (T - t)$ called evaluation cost. However, although Remark 1 in (Hu et al., 2025) claim that their $O(\gamma_T \sqrt{T})$ cumulative regret upper bound is tighter than the existing result, $O(\gamma_T \sqrt{T})$ cumulative regret upper bound has already been shown in (Wang & de Freitas, 2014) with the reference value μ_{t-1}^{\max} , as described above. Furthermore, even in the case that the reference value is $\tilde{\mu}_{t-1}^{\max}$, the cumulative regret upper bound $O(\gamma_T \sqrt{T})$ can be obtained without any modification to the GP-EI algorithm (see Appendix B). Hence, the importance of the candidate elimination based on the evaluation cost is ambiguous in the theorem.

¹Although Theorem 1 in (Wang & de Freitas, 2014) states $O(\gamma_T^{3/2} \sqrt{T})$, we believe that their proof suggests $O(\gamma_T \sqrt{T})$.

3.3 Other Related Works

Several other studies analyze GP-EI without rescaling of the posterior variance based on (Nguyen et al., 2017). Nguyen et al. (2017) claims that the cumulative regret upper bound for the usual GP-EI with y_t^{\max} is sublinear in the frequentist setting. However, as shortly discussed in Section 3 of (Tran-The et al., 2022), Remark 1 of (Hu et al., 2025), and Section 1 of (Wang et al., 2025), Lemma 7 in (Nguyen et al., 2017), which claims $\sum_{t=1}^T \sigma_t^2(\mathbf{x}) = O(\gamma_T)$ for all $\mathbf{x} \in \mathcal{X}$, is incorrect. Note that the usual upper bound by MIG is given as $\sum_{t=1}^T \sigma_t^2(\mathbf{x}_t) = O(\gamma_T)$ (Srinivas et al., 2010). The incorrect derivation in (Nguyen et al., 2017) is the last equality in the proof of Lemma 7. The counter-examples are shown in Appendix A.

At least to our knowledge, Eq (8) in (Berk et al., 2019), page 37 in (Grosnit et al., 2021), Eq. (9.26) in (Bedi et al., 2022), Lemma B.3 in (Marisu & Pun, 2023), and Lemma 2 in (Zhou et al., 2024) in the regret analyses of GP-EI-based algorithms are also incorrect by using a similar result from Lemma 7 in (Nguyen et al., 2017). In addition, this mistake has been used for regret analyses for other AFs, such as Lemma 6 of (Nguyen et al., 2019), Lemma 12 in (Husain et al., 2023).

Exploration enhanced EI (E³I) (Berk et al., 2019) is the following GP-EI variants:

$$E^3I(\mathbf{x}) = \mathbb{E}_{g_t^* \sim p_t(f(\mathbf{x}^*))} \left[\sigma_t(\mathbf{x}) \tau \left(\frac{\mu_t(\mathbf{x}) - g_t^*}{\sigma_t(\mathbf{x})} \right) \right], \quad (10)$$

where $g_t \sim p_t(f)$ is the sample path from the posterior and $g_t^* = \max_{\mathbf{x} \in \mathcal{X}} g_t(\mathbf{x})$. The expectation regarding $g_t^* \sim p_t(f(\mathbf{x}^*))$ is approximated by Monte Carlo (MC) estimation. However, as we already discussed, their analysis appears incorrect. We will consider GP-EIMS, a similar posterior sampling-based GP-EI algorithm that uses only one posterior sample. In the BCR analyses of GP-EIMS, the probability matching property between g_t^* and $f(\mathbf{x}^*)$, that is, the randomness caused by the algorithm relying on only one sample, plays a key role, as with the analyses of GP-TS (Russo & Van Roy, 2014; Kandasamy et al., 2018; Takeno et al., 2024).

Fu et al. (2024) analyze the GP-EI algorithm without the posterior variance scaling for the bilevel optimization problem. However, the derivation from Eq. (D.8) to Eq. (D.9) in (Fu et al., 2024) is not obvious. In Eq. (D.8) in (Fu et al., 2024), the regret upper bound has the multiplicative term $\frac{1}{\tau(\beta_T) - \beta_T}$ with $\beta_T = \Theta(\log T)$. Then, $\tau(\beta_T) - \beta_T = \beta_T \phi(\beta_T) \left(\frac{1}{\beta_T} - \frac{1 - \Phi(\beta_T)}{\phi(\beta_T)} \right) = O(\beta_T \phi(\beta_T))$ since the mills ratio $\frac{1 - \Phi(c)}{\phi(c)} \rightarrow \frac{1}{c}$ as $c \rightarrow \infty$ (for example, see Eq. (7) in Gasull & Utzet, 2014) and $\frac{1}{c} - \frac{1 - \Phi(c)}{\phi(c)} > 0$ for all $c > 0$. Therefore, we can see $\frac{1}{\tau(\beta_T) - \beta_T} = \Omega(1/(\beta_T \phi(\beta_T)))$. Recalling $\beta_T = \Theta(\log T)$, the term $\frac{1}{\tau(\beta_T) - \beta_T} = \Omega(T^{\log T} / \log T)$, which is not sublinear.

Wang et al. (2025) tackle a theoretical analysis of the usual EI with y_t^{\max} without the posterior variance scaling. However, Wang et al. (2025) analyze a quantity $f(\mathbf{x}^*) - y_T^{\max}$, which is not the established regret definition, and can be negative. Paritularly when the noise variance σ^2 is large compared with $f(\mathbf{x}^*)$, then $f(\mathbf{x}^*) - y_T^{\max}$ can diverge to $-\infty$ rapidly by noise. Therefore, analyzing $f(\mathbf{x}^*) - y_T^{\max}$ is not suggestive in general.

4 GP-EIMS

This section provides the algorithm and the BCR upper bounds of GP-EIMS.

4.1 Algorithm

As with GP-PIMS (Takeno et al., 2024), we use the sampled maximum defined as follows:

$$g_t^* = \max_{\mathbf{x} \in \mathcal{X}} g_t(\mathbf{x}), \quad (11)$$

Algorithm 1 GP-EIMS**Require:** Input space \mathcal{X} , GP prior $\mu = 0$ and k , and initial dataset \mathcal{D}_0

```

1: for  $t = 1, \dots$  do
2:   Fit GP to  $\mathcal{D}_{t-1}$ 
3:   Generate a sample path  $g_t \sim p(f|\mathcal{D}_{t-1})$ 
4:    $g_t^* \leftarrow \max_{\mathbf{x} \in \mathcal{X}} g_t(\mathbf{x})$ 
5:    $\mathbf{x}_t \leftarrow \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}(\mathbf{x}), g_t^*)$ 
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

```

where $g_t \sim p(f|\mathcal{D}_{t-1})$ is the sample path from the posterior. Then, the AF of GP-EIMS is

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}(\mathbf{x}), g_t^*) = \mathbb{E}_{f(\mathbf{x}) \sim p_t(f(\mathbf{x}))} [\max\{f(\mathbf{x}) - g_t^*, 0\} | g_t^*] \quad (12)$$

$$= \sigma_{t-1}(\mathbf{x}) \tau \left(\frac{\mu_{t-1}(\mathbf{x}) - g_t^*}{\sigma_{t-1}(\mathbf{x})} \right). \quad (13)$$

Note that the posterior variance $\sigma_{t-1}^2(\mathbf{x})$ is not rescaled. Finally, an algorithm of GP-EIMS is shown in Algorithm 1.

4.2 Regret Analysis

First, we can obtain a general BCR upper bound for arbitrary BO algorithms, as shown in (Takeno et al., 2024):

Lemma 4.1. *Let*

$$\eta_t := \frac{g_t^* - \mu_{t-1}(\mathbf{x}_t)}{\sigma_{t-1}(\mathbf{x}_t)}. \quad (14)$$

Then, the BCR can be bounded from above as follows:

$$\text{BCR}_T \leq \sqrt{\mathbb{E} \left[\sum_{t=1}^T \eta_t^2 \mathbb{1}\{\eta_t \geq 0\} \right]} \sqrt{C_1 \gamma_T}, \quad (15)$$

where $C_1 := 2/\log(1 + \sigma^{-2})$ and the indicator function $\mathbb{1}\{\eta_t \geq 0\} = 1$ if $\eta_t \geq 0$, and 0 otherwise.

For completeness, we show the proof in Appendix C.1. Hence, we can concentrate on deriving $\sum_{t=1}^T \mathbb{E} [\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}] = o(T^2/\gamma_T)$ based on the property of GP-EIMS so that the resulting BCR upper bounds are sublinear.

For this purpose, the following lemma plays a key role in our analysis. For more detailed proof, see Appendix C.2.

Lemma 4.2. *Fix $\beta > 0$ and $U \leq \max_{\mathbf{x} \in \mathcal{X}} \{\mu_{t-1}(\mathbf{x}) + \beta^{1/2} \sigma_{t-1}(\mathbf{x})\}$. Then, the following inequality holds:*

$$\frac{U - \mu_{t-1}(\mathbf{x}_t)}{\sigma_{t-1}(\mathbf{x}_t)} \leq \sqrt{\log \left(\frac{\sigma^2 + t - 1}{\sigma^2} \right) + \beta + \sqrt{2\pi\beta}}, \quad (16)$$

where $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}(\mathbf{x}), U)$.

Short proof. As with (Wang & de Freitas, 2014), by using the properties of $\tau(\cdot)$ and the posterior variance and the definition of \mathbf{x}_t , we can obtain

$$|\mu_{t-1}(\mathbf{x}_t) - U| \leq \left(\sqrt{\log \left(\frac{\sigma^2 + t - 1}{\sigma^2} \right) - 2 \log(\tau(-\beta^{1/2})) - \log(2\pi)} \right) \sigma_{t-1}(\mathbf{x}_t). \quad (17)$$

However, since the AF is defined by U instead of $\max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x})$ in (Wang & de Freitas, 2014), we must further obtain an upper bound of $-2 \log(\tau(-\beta^{1/2}))$. To obtain the upper bound, we leverage the following useful result:

Lemma 4.3 (Lemma 1 in (Jang, 2011)). *Q-function $1 - \Phi(c)$ can be bounded from above as follows:*

$$1 - \Phi(c) \leq \frac{1}{\sqrt{2\pi}c} \left(1 - \exp\left(-\sqrt{\frac{\pi}{2}}c\right)\right) \exp\left(-\frac{c^2}{2}\right) \quad (18)$$

$$= \frac{1}{c} \left(1 - \exp\left(-\sqrt{\frac{\pi}{2}}c\right)\right) \phi(c), \quad (19)$$

for all $c > 0$.

By using Lemma 4.3 and properties of Φ and ϕ , we can obtain the following inequality:

$$\tau(-\beta^{1/2}) = -\beta^{1/2}\Phi(-\beta^{1/2}) + \phi(-\beta^{1/2}) \quad (20)$$

$$= -\beta^{1/2} \left(1 - \Phi(\beta^{1/2})\right) + \phi(\beta^{1/2}) \quad (21)$$

$$\geq - \left(1 - \exp\left(-\sqrt{\frac{\pi}{2}}\beta^{1/2}\right)\right) \phi(\beta^{1/2}) + \phi(\beta^{1/2}) \quad (22)$$

$$= \exp\left(-\sqrt{\frac{\pi}{2}}\beta^{1/2}\right) \phi(\beta^{1/2}). \quad (23)$$

Hence, we see that $-2 \log(\tau(-\beta^{1/2})) \leq \sqrt{2\pi}\beta - 2 \log(\phi(\beta^{1/2})) = \sqrt{2\pi}\beta + \beta + \log(2\pi)$. Consequently, by combining the results, we obtain the desired result. \square

By leveraging Lemmas 4.1 and 4.2, we show BCR upper bounds for discrete and continuous input domains.

4.2.1 BCR Upper Bound for Finite Input Domain

We combine the following lemma from (Srinivas et al., 2010; Takeno et al., 2023) with Lemma 4.2:

Lemma 4.4 (Lemma 4.1 in (Takeno et al., 2023)). *Suppose that f is a sample path from a GP with zero mean and a predefined kernel k , and \mathcal{X} is finite. Pick $\delta \in (0, 1)$ and $t \geq 1$. Then, for any given \mathcal{D}_{t-1} ,*

$$\Pr_t \left(f(\mathbf{x}) \leq \mu_{t-1}(\mathbf{x}) + \beta^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}), \forall \mathbf{x} \in \mathcal{X} \right) \geq 1 - \delta, \quad (24)$$

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

Since $f(\mathbf{x}^*) | \mathcal{D}_{t-1}$ and $g_t^* | \mathcal{D}_{t-1}$ are identically distributed, this lemma implies that

$$\Pr \left(g_t^* \leq \max_{\mathbf{x} \in \mathcal{X}} \{ \mu_{t-1}(\mathbf{x}) + \beta^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}) \} \right) \geq 1 - \delta. \quad (25)$$

Thus, applying Lemma 4.2 by substituting g_t^* to U , we can derive the following result:

Corollary 4.5. *Assume the same condition as in Lemma 4.4. Then, by running GP-EIMS, the following inequality holds with probability at least $1 - \delta$:*

$$\eta_t \leq \sqrt{\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(\delta) + \sqrt{2\pi}\beta(\delta)}, \quad (26)$$

where $\beta(\delta) = 2 \log(|\mathcal{X}|/(2\delta))$. Furthermore, we obtain

$$\mathbb{E} [\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}] \leq \log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + C_2 + \sqrt{2\pi}C_2, \quad (27)$$

where $C_2 = 2 + 2 \log(|\mathcal{X}|/2)$.

See Appendix C.3 for the proof. The first inequality is a direct consequence of Lemmas 4.2 and 4.4. The second inequality can be derived similarly to the proof of Lemma 3.2 in (Takeno et al., 2024).

Consequently, from Lemma 4.1 and Corollary 4.5, we can obtain the following BCR upper bound:

Theorem 4.6. *Let $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Then, by running GP-EIMS, the BCR can be bounded from above as follows:*

$$\text{BCR}_T \leq \sqrt{C_1 B_T T \gamma_T}, \quad (28)$$

where $B_T = \log\left(\frac{\sigma^2 + T - 1}{\sigma^2}\right) + C_2 + \sqrt{2\pi C_2}$, $C_1 := 2/\log(1 + \sigma^{-2})$, and $C_2 = 2 + 2\log(|\mathcal{X}|/2)$.

See Appendix C.3 for the proof. The BCR upper bound of GP-EIMS for finite input domains is $O(\sqrt{T\gamma_T \log(T|\mathcal{X}|)})$, which is sublinear at least for the SE kernels, linear kernels, and Matérn- ν kernels with $2\nu > d$.

4.2.2 BCR Upper Bound for Continuous Input Domain

For continuous input domains, as with the prior works (for example, Srinivas et al., 2010; Kandasamy et al., 2018; Takeno et al., 2023; 2024), we consider a discretized input set $\mathcal{X}_t \subset \mathcal{X}$, which is a finite set with each dimension equally divided into $m_t \geq 1$. Thus, $|\mathcal{X}_t| = m_t^d$. Furthermore, we define the nearest point of \mathbf{x} in \mathcal{X}_t as $[\mathbf{x}]_t = \arg \min_{\mathbf{x}' \in \mathcal{X}_t} \|\mathbf{x} - \mathbf{x}'\|_1$.

If we consider the AF, $\text{EI}(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}(\mathbf{x}), \tilde{g}_t^*)$, where $\tilde{g}_t^* := \max_{\mathbf{x} \in \mathcal{X}_t} g_t(\mathbf{x})$, a sublinear BCR upper bound can be derived in a similar way to the proof of Theorem 4.6. However, m_t for \mathcal{X}_t must depend on several unknown parameters such as a and b in Assumption 2.1. Thus, m_t must be tuned by users if the computation of \tilde{g}_t^* is required in practice. Hence, using \tilde{g}_t^* means that a new tuning parameter m_t is involved, which is not preferable. Therefore, we consider a BCR upper bound for $\text{EI}(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}(\mathbf{x}), g_t^*)$, which does not require the discretization in the actual algorithm, as with the analyses of GP-PIMS (Takeno et al., 2024).

From Lemma 4.4 and Assumption 2.1, we can obtain the following lemma by appropriately controlling m_t :

Lemma 4.7. *Suppose Assumption 2.1 holds. Pick $\delta \in (0, 1)$ and $t \geq 1$. Then, the following holds:*

$$\Pr_t \left(\forall \mathbf{x} \in \mathcal{X}, f(\mathbf{x}) \leq \mu_{t-1}([\mathbf{x}]_t) + 2\beta_t^{1/2}(\delta/2)\sigma_{t-1}([\mathbf{x}]_t) \right) \geq 1 - \delta \quad (29)$$

where $\beta_t(\delta) = 2d \log(m_t) - 2\log(2\delta)$ and $m_t = \max\{2, \lceil bdr \sqrt{(\sigma^2 + t - 1) \log(2ad)/\sigma^2} \rceil\}$.

See Appendix C.4 for the proof. As with the case of finite input domains, since $f(\mathbf{x}^*) | \mathcal{D}_{t-1}$ and $g_t^* | \mathcal{D}_{t-1}$ are identically distributed, Lemmas 4.2 and 4.7 derive

$$\mathbb{E} [\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}] \leq \log\left(\frac{\sigma^2 + t}{\sigma^2}\right) + C_T + \sqrt{2\pi C_T}, \quad (30)$$

where $C_T = 8(d \log(m_T) + 1)$.

Consequently, the BCR upper bound can be derived:

Theorem 4.8. *Suppose Assumption 2.1 holds and $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$. Then, by running GP-EIMS, the BCR can be bounded from above as follows:*

$$\text{BCR}_T \leq \sqrt{C_1 B_T T \gamma_T}, \quad (31)$$

where $C_1 := 2/\log(1 + \sigma^{-2})$, $B_T = \log\left(\frac{\sigma^2 + T - 1}{\sigma^2}\right) + C_T + \sqrt{2\pi C_T}$, $C_T = 8(d \log(m_T) + 1)$, and $m_T = \max\{2, \lceil bdr \sqrt{(\sigma^2 + T - 1) \log(2ad)/\sigma^2} \rceil\}$.

See Appendix C.4 for the proof. As with Theorem 4.6, $\text{BCR}_T = O(\sqrt{T\gamma_T \log T})$, which is sublinear for widely used kernels, such as the SE kernels, linear kernels, and Matérn- ν kernels with $2\nu > d$.

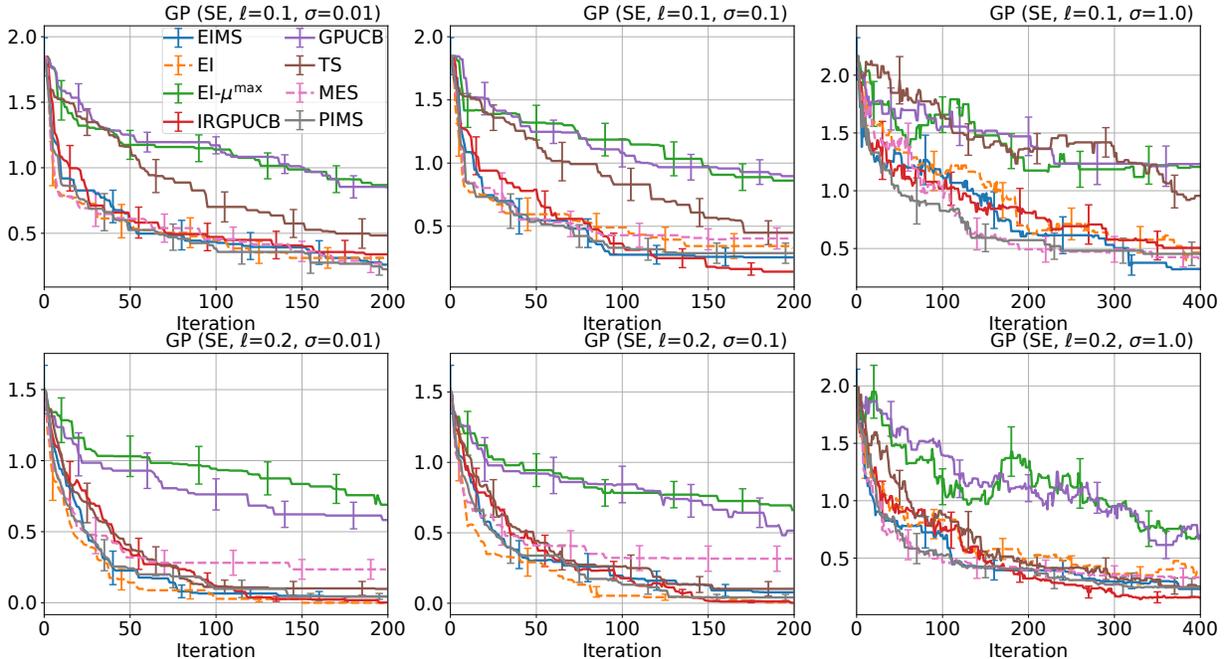


Figure 1: Results of simple regret for synthetic functions generated from a GP defined by the SE kernel. The top and bottom rows show the result on the length scale of the kernel $\ell = 0.1$ and $\ell = 0.2$, respectively. The left, centered, and right columns show the result on the noise standard deviation $\sigma = 0.01, 0.1$, and 1 , respectively.

5 Experiments

We performed numerical experiments using synthetic functions generated from GPs, which match the assumptions for our analysis. We set $\mathcal{X} = \{0.0, 0.1, \dots, 0.9\}^d$ and $d = 4$. Therefore, $|\mathcal{X}| = 10^4$. We employed the SE kernel $k_{\text{SE}}(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|_2^2 / (2\ell^2))$. Furthermore, we actually added observation noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$. We report results changing $\ell \in \{0.1, 0.2\}$ and $\sigma \in \{0.01, 0.1, 1\}$. We fixed the hyperparameters of the GP model, that is ℓ and σ , to the parameters used to generate the synthetic functions and observation noise. We set an initial dataset to data that is closest to 2^d data generated randomly as a Sobol sequence.

As baselines, we employed GP-UCB (Srinivas et al., 2010), improved randomized GP-UCB (IRGP-UCB) (Takeno et al., 2023), which uses the random variable ζ_t instead of β_t in GP-UCB, GP-TS (Russo & Van Roy, 2014), max-value entropy search (MES) (Wang & Jegelka, 2017), GP-PIMS (Takeno et al., 2024), GP-EI (Mockus et al., 1978), and GP-EI using the maximum of the posterior mean (Wang & de Freitas, 2014). We denote the method of (Wang & de Freitas, 2014) as GP-EI- μ^{\max} . We set the hyperparameters of GP-UCB, IRGP-UCB, and GP-EI- μ^{\max} to the theoretically derived parameters for the BCR analyses. Therefore, we set $\beta_t = 2 \log(|\mathcal{X}|t^2 / \sqrt{2\pi} + 1)$ for GP-UCB and $\zeta_t = 2 \log(|\mathcal{X}|/2) + Z_t$ with $Z_t \sim \text{Exp}(\lambda = 1/2)$ for IRGP-UCB (Takeno et al., 2023). Furthermore, we set $\nu_t = 2 \log(|\mathcal{X}|t^2 / \sqrt{2\pi} + 1)$ for GP-EI- μ^{\max} (see Theorem B.4). The number of MC samples in MES was set to 10. We employed random Fourier feature approximations (Rahimi & Recht, 2008) for the posterior sampling in GP-EIMS, GP-TS, MES, and GP-PIMS.

Figures 1 and 2 show the simple regret and cumulative regret for the SE kernels. We measured optimization performance by the simple regret $f(\mathbf{x}^*) - f(\hat{\mathbf{x}}_t)$, where $\hat{\mathbf{x}}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_t(\mathbf{x})$, and the cumulative regret $\sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t)$. We report the mean and standard error over 16 random trials for the generation of the initial dataset, the synthetic functions, the observation noise, and the algorithm’s randomness. The BO methods whose theoretical analysis has not been shown are indicated by dashed lines.

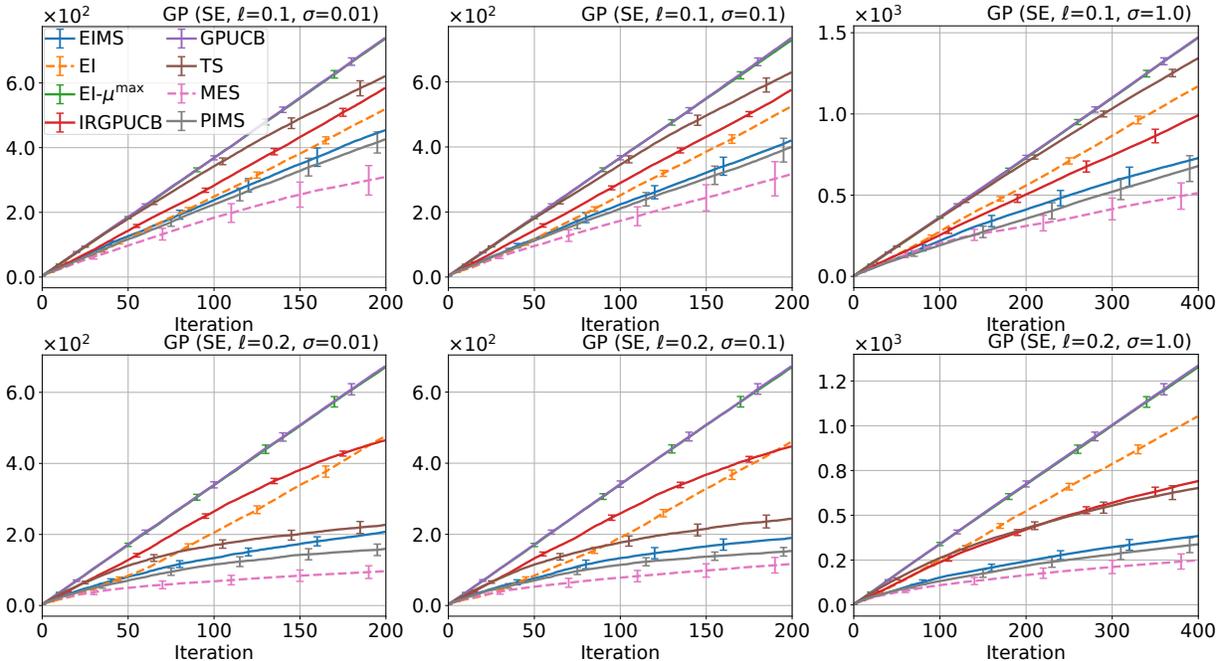


Figure 2: Results of cumulative regret for synthetic functions generated from a GP defined by the SE kernel. The top and bottom rows show the result on the length scale of the kernel $\ell = 0.1$ and $\ell = 0.2$, respectively. The left, centered, and right columns show the result on the noise standard deviation $\sigma = 0.01, 0.1$, and 1 , respectively.

GP-UCB and GP-EI- μ^{\max} consistently inferior to other methods. This is because the theoretical values of β_t and ν_t are too large in practice. Indeed, IRGP-UCB, which uses $\zeta_t = O(\log |\mathcal{X}|)$ instead of $\beta_t = O(\log(|\mathcal{X}|t))$, shows better performance than GP-UCB and GP-EI- μ^{\max} . However, IRGP-UCB slightly deteriorates when $\ell = 0.2$, particularly for the cumulative regret, since ζ_t is still large for the objective function with $\ell = 0.2$ whose local optima are not so many. Note that if ℓ becomes small, then the objective function has large variability and more local optima.

GP-TS shows superior performance when $\ell = 0.2$. However, when $\ell = 0.1$, GP-TS often deteriorates for both simple and cumulative regret. This tendency has been reported as over-exploration of GP-TS (Takeno et al., 2023; 2024; Shahriari et al., 2016).

MES shows superior performance in terms of cumulative regret, although its regret analysis has not been shown, as discussed in (Takeno et al., 2022). However, the simple regret stagnates in the case of $\ell = 0.2$ and $\sigma = 0.01, 0.1$. Therefore, we can interpret that MES is slightly too exploitative in these experiments. On the other hand, it is known that GP-PIMS can be viewed as MES using one MC sample (Takeno et al., 2024). Superior performance of MES in terms of cumulative regret may imply the existence of a regret guarantee of MES.

GP-EI shows superior performance for the simple regret, although its regret analysis for $\sigma > 0$ has not been shown. However, in the experiments for the cumulative regret, GP-EI clearly deteriorates, particularly for the case of $\ell = 0.2$. Since the noisy best observation y_t^{\max} tends to be large with larger noise ϵ , we can expect that GP-EI tends to lean more towards exploration. Since uncertainty sampling that maximizes $\sigma_{t-1}(\mathbf{x})$ achieves the simple regret upper bound that converges to zero (Vakili et al., 2021), it is understandable that GP-EI shows superior performance for the simple regret². On the other hand, we conjecture that the deterioration of GP-EI for the cumulative regret is due to over-exploration. Furthermore, we consider that

²Although Vakili et al. (2021) analyze in the frequentist setting, their analysis immediately suggests a simple regret upper bound in the Bayesian setting.

since the objective function with $\ell = 0.1$ has many local optima, the aggressive exploration by larger y_t^{\max} results in relatively better performance compared with that of $\ell = 0.2$.

GP-EIMS and GP-PIMS show superior performance consistently for both simple and cumulative regret, particularly compared with other theoretically guaranteed methods. Both GP-EIMS and GP-PIMS are based on the posterior sample of g_t^* and show similar results in these experiments, although GP-EIMS chooses a larger $\sigma_{t-1}(\mathbf{x}_t)$ compared with GP-PIMS, due to the formulation of the AF.

We performed the experiments for the Matérn kernels $k_{\text{Mat}}(\mathbf{x}, \mathbf{x}') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\sqrt{2\nu} \frac{\|\mathbf{x} - \mathbf{x}'\|}{\ell} \right)^\nu K_\nu \left(\sqrt{2\nu} \frac{\|\mathbf{x} - \mathbf{x}'\|}{\ell} \right)$, where $\Gamma(\cdot)$ and $K_\nu(\cdot)$ are the gamma function and the modified Bessel function, and several benchmark functions from <https://www.sfu.ca/~ssurjano/optimization.html>. These experiments are shown in Appendix E. In the experiments for the Matérn kernels, we can observe a similar tendency to the experiments for the SE kernel. In the experiments for the benchmark functions, we confirmed that all BO methods perform well in most cases since we employed heuristic parameter settings for GP-UCB, IRGP-UCB, and GP-EI- μ^{\max} . On the other hand, GP-TS, GP-PIMS, and GP-EIMS show good performance without heuristic tunings, and GP-PIMS and GP-EIMS show superior or comparable performance to GP-TS.

6 Conclusion

The regret analyses of GP-EI-based algorithms are fewer than those of theoretically established algorithms, such as GP-UCB and GP-TS. Existing analysis by Wang & de Freitas (2014) requires increasing the posterior variance in the actual algorithm, which can deteriorate the optimization performance. Therefore, we analyzed the posterior sampling-based GP-EI algorithm called GP-EIMS, which does not require rescaling the posterior variance. We derived the sublinear BCR upper bounds of the GP-EIMS algorithm in the Bayesian setting. We finally demonstrated the practical effectiveness of GP-EIMS through numerical experiments.

Limitation and Future Work: Several future directions for this study exist. First, the posterior sampling for continuous \mathcal{X} or discrete \mathcal{X} with huge $|\mathcal{X}|$ requires some approximation, such as random Fourier features, in practice. Theoretical analysis incorporating such approximations (Mutny & Krause, 2018) is interesting. Second, since we show only the BCR upper bounds, whether GP-EIMS (and GP-PIMS) achieves sublinear high-probability cumulative regret bounds is of interest. Third, recent studies (Scarlett, 2018; Iwazaki, 2025) imply that the standard GP-UCB achieves a near-optimal regret upper bound with high probability. Showing whether other BO methods, including GP-EIMS, achieve a near-optimal regret upper bound is vital, as discussed in Section 4 in (Iwazaki, 2025). Lastly, our analysis for GP-EIMS concentrates on the Bayesian setting. As with the analysis of GP-TS in the frequentist setting (Chowdhury & Gopalan, 2017), theoretical analysis in the frequentist setting is also intriguing.

Broader Impact Statement

Our work focuses on theoretical aspects of machine learning methods, and we do not anticipate any negative societal impact.

References

- Yasin Abbasi-yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc., 2011.
- Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected improvements to expected improvement for Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 36, pp. 20577–20612. Curran Associates, Inc., 2023.
- Amrit Singh Bedi, Dheeraj Peddireddy, Vaneet Aggarwal, Brian M. Sadler, and Alec Koppel. Regret and belief complexity trade-off in Gaussian process bandits via information thresholding. *arXiv:2003.10550*, 2022.

- Julian Berk, Vu Nguyen, Sunil Gupta, Santu Rana, and Svetha Venkatesh. Exploration enhanced expected improvement for Bayesian optimization. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD*, pp. 621–637. Springer, 2019.
- Felix Berkenkamp, Andreas Krause, and Angela P Schoellig. Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics. *Machine Learning*, 112(10):3713–3747, 2023.
- Eric Brochu, Vlad M. Cora, and Nando de Freitas. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv.1012.2599*, 2010.
- Adam D Bull. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12(10), 2011.
- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 844–853, 2017.
- Michael T. M. Emmerich, André H. Deutz, and Jan Willem Klinkenberg. Hypervolume-based expected improvement: Monotonicity properties and exact computation. In *2011 IEEE Congress of Evolutionary Computation (CEC)*, pp. 2147–2154, 2011.
- Shi Fu, Fengxiang He, Xinmei Tian, and Dacheng Tao. Convergence of Bayesian bilevel optimization. In *The Twelfth International Conference on Learning Representations*, 2024.
- Jacob Gardner, Matt Kusner, Zhixiang, Kilian Weinberger, and John Cunningham. Bayesian optimization with inequality constraints. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 937–945. PMLR, 2014.
- Armengol Gasull and Frederic Utzet. Approximating mills ratio. *Journal of Mathematical Analysis and Applications*, 420(2):1832–1853, 2014.
- Michael A. Gelbart, Jasper Snoek, and Ryan P. Adams. Bayesian optimization with unknown constraints. In *Proceedings of the 30th Conference on Uncertainty in Artificial Intelligence*, pp. 250–259. AUAI Press, 2014.
- Subhashis Ghosal and Anindya Roy. Posterior consistency of Gaussian process prior for nonparametric binary regression. *The Annals of Statistics*, 34(5):2413 – 2429, 2006.
- Antoine Grosnit, Rasul Tutunov, Alexandre Max Maraval, Ryan-Rhys Griffiths, Alexander I. Cowen-Rivers, Lin Yang, Lin Zhu, Wenlong Lyu, Zhitang Chen, Jun Wang, Jan Peters, and Haitham Bou-Ammar. High-dimensional Bayesian optimisation with variational autoencoders and deep metric learning. *arXiv:2106.03609*, 2021.
- Steffen Grünewälder, Jean-Yves Audibert, Manfred Opper, and John Shawe-Taylor. Regret bounds for Gaussian process bandit problems. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pp. 273–280. PMLR, 2010.
- Shouri Hu, Haowei Wang, Zhongxiang Dai, Bryan Kian Hsiang Low, and Szu Hui Ng. Adjusted expected improvement for cumulative regret minimization in noisy Bayesian optimization. *Journal of Machine Learning Research*, 26(46):1–33, 2025.
- Deng Huang, Theodore T. Allen, William I. Notz, and Richard A. Miler. Sequential kriging optimization using multiple-fidelity evaluations. *Structural and Multidisciplinary Optimization*, 32(5):369–382, 2006.
- Hisham Husain, Vu Nguyen, and Anton van den Hengel. Distributionally robust Bayesian optimization with φ -divergences. In *Advances in Neural Information Processing Systems*, volume 36, pp. 20133–20145. Curran Associates, Inc., 2023.

- Shogo Iwazaki. Improved regret bounds for Gaussian process upper confidence bound in Bayesian optimization. *arXiv:2506.01393*, 2025.
- Shogo Iwazaki and Shion Takeno. Improved regret analysis in Gaussian process bandits: Optimality for noiseless reward, RKHS norm, and non-stationary variance. In *Proceedings of the 41st International Conference on Machine Learning*, Proceedings of Machine Learning Research. PMLR, 2025. To appear.
- Won Mee Jang. A simple upper bound of the Gaussian Q-function with closed-form error bound. *IEEE Communications Letters*, 15(2):157–159, 2011.
- David Janz, David Burt, and Javier Gonzalez. Bandit optimisation of functions in the Matérn kernel RKHS. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 2486–2495, 2020.
- Donald R. Jones. A taxonomy of global optimization methods based on response surfaces. *Journal of global optimization*, 21:345–383, 2001.
- Donald R. Jones, Matthias Schonlau, and William J Welch. Efficient global optimization of expensive black-box functions. *Journal of Global optimization*, 13:455–492, 1998.
- K. Kandasamy, G. Dasarathy, J. Oliva, J. Schneider, and B. Póczos. Gaussian process bandit optimisation with multi-fidelity evaluations. In *Advances in Neural Information Processing Systems 29*, pp. 1000–1008. Curran Associates, Inc., 2016.
- Kirthevasan Kandasamy, Akshay Krishnamurthy, Jeff Schneider, and Barnabas Póczos. Parallelised Bayesian optimisation via Thompson sampling. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, volume 84 of *Proceedings of Machine Learning Research*, pp. 133–142, 2018.
- Harold J. Kushner. A New Method of Locating the Maximum Point of an Arbitrary Multipeak Curve in the Presence of Noise. *Journal of Basic Engineering*, 86(1):97–106, 1964.
- Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained Bayesian optimization with noisy experiments. *Bayesian Anal.*, 14(2):495–519, 2019.
- Daniel James Lizotte. Practical Bayesian optimization. *PhD thesis, University of Alberta*, 2008.
- Godeliva Petrina Marisu and Chi Seng Pun. Bayesian estimation and optimization for learning sequential regularized portfolios. *SIAM Journal on Financial Mathematics*, 14(1):127–157, 2023.
- Jonas Mockus, Vytautas Tiesis, and Antanas Zilinskas. The application of Bayesian methods for seeking the extremum. *Towards Global Optimization*, 2(117-129):2, 1978.
- Mojmir Mutny and Andreas Krause. Efficient high dimensional Bayesian optimization with additivity and quadrature Fourier features. *Advances in Neural Information Processing Systems 31*, pp. 9005–9016, 2018.
- Vu Nguyen, Sunil Gupta, Santu Rana, Cheng Li, and Svetha Venkatesh. Regret for expected improvement over the best-observed value and stopping condition. In *Asian conference on machine learning*, pp. 279–294. PMLR, 2017.
- Vu Nguyen, Sunil Gupta, Santu Rana, Cheng Li, and Svetha Venkatesh. Filtering Bayesian optimization approach in weakly specified search space. *Knowledge and Information Systems*, 60:385–413, 2019.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems 20*, pp. 1177–1184. Curran Associates, Inc., 2008.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.
- Daniel Russo and Benjamin Van Roy. Learning to optimize via posterior sampling. *Mathematics of Operations Research*, 39(4):1221–1243, 2014.

- Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. *Operations Research*, 66(1):230–252, 2018.
- Jonathan Scarlett. Tight regret bounds for Bayesian optimization in one dimension. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 4500–4508. PMLR, 2018.
- Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando De Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016.
- Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical Bayesian optimization of machine learning algorithms. In *Advances in Neural Information Processing Systems 25*, pp. 2951–2959. Curran Associates, Inc., 2012.
- Niranjn Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 1015–1022. Omnipress, 2010.
- Shion Takeno, Tomoyuki Tamura, Kazuki Shitara, and Masayuki Karasuyama. Sequential and parallel constrained max-value entropy search via information lower bound. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 20960–20986, 2022.
- Shion Takeno, Yu Inatsu, and Masayuki Karasuyama. Randomized Gaussian process upper confidence bound with tighter Bayesian regret bounds. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 33490–33515. PMLR, 2023.
- Shion Takeno, Yu Inatsu, Masayuki Karasuyama, and Ichiro Takeuchi. Posterior sampling-based Bayesian optimization with tighter Bayesian regret bounds. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 47510–47534. PMLR, 2024.
- Hung Tran-The, Sunil Gupta, Santu Rana, and Svetha Venkatesh. Regret bounds for expected improvement algorithms in Gaussian process bandit optimization. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 8715–8737. PMLR, 2022.
- Tsuyoshi Ueno, Trevor David Rhone, Zhufeng Hou, Teruyasu Mizoguchi, and Koji Tsuda. COMBO: An efficient Bayesian optimization library for materials science. *Materials discovery*, 4:18–21, 2016.
- Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in Gaussian process bandits. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 82–90, 2021.
- Emmanuel Vazquez and Julien Bect. Convergence properties of the expected improvement algorithm with fixed mean and covariance functions. *Journal of Statistical Planning and Inference*, 140(11):3088–3095, 2010.
- Jingyi Wang, Haowei Wang, Nai-Yuan Chiang, and Cosmin G. Petra. On the convergence rate of noisy Bayesian optimization with expected improvement. *arXiv:2501.09262*, 2025.
- Zi Wang and Stefanie Jegelka. Max-value entropy search for efficient Bayesian optimization. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 3627–3635, 2017.
- Ziyu Wang and Nando de Freitas. Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters. *arXiv:1406.7758*, 2014.
- Han Zhou, Xingchen Ma, and Matthew B Blaschko. A corrected expected improvement acquisition function under noisy observations. In *Proceedings of the 15th Asian Conference on Machine Learning*, volume 222 of *Proceedings of Machine Learning Research*, pp. 1747–1762. PMLR, 2024.

A Technical Issue in (Nguyen et al., 2017)

A counter-example against Lemma 7 in (Nguyen et al., 2017) can be made easily. Let us consider the BO algorithm that always evaluates the same point; that is, $\exists \mathbf{x} \in \mathcal{X}, \forall t \in [T], \mathbf{x}_t = \mathbf{x}$. Then, choosing an input $\tilde{\mathbf{x}}$ such that $k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}) - k(\mathbf{x}, \tilde{\mathbf{x}})^2/k(\mathbf{x}, \mathbf{x}) \geq C$ with some strictly positive constant C , we can confirm $\sigma_t^2(\tilde{\mathbf{x}}) \geq k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}) - k(\mathbf{x}, \tilde{\mathbf{x}})^2/k(\mathbf{x}, \mathbf{x}) \geq C > 0$ for all $t \in [T]$. Therefore, $\sum_{t=1}^T \sigma_t^2(\tilde{\mathbf{x}}) \geq CT = \Omega(T)$, which contradicts Lemma 7 in (Nguyen et al., 2017).

Note that if Lemma 7 in (Nguyen et al., 2017) is true, *any* BO algorithm achieves the converged simple regret upper bound in some settings. Suppose that $\sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}) = O(\gamma_T)$ for all $\mathbf{x} \in \mathcal{X}$ holds. Assume \mathcal{X} is finite and β_T is set so that for all $\mathbf{x} \in \mathcal{X}$, $|\mu_T(\mathbf{x}) - f(\mathbf{x})| \leq \beta_T^{1/2} \sigma_T(\mathbf{x})$ with high probability (for example, as in Srinivas et al., 2010). Then, the following holds with high probability:

$$f(\mathbf{x}^*) - f(\hat{\mathbf{x}}_T) \leq \mu_T(\mathbf{x}^*) + \beta_T^{1/2} \sigma_T(\mathbf{x}^*) - \mu_T(\hat{\mathbf{x}}_T) + \beta_T^{1/2} \sigma_T(\hat{\mathbf{x}}_T) \quad (32)$$

$$\leq 2\beta_T^{1/2} \sigma_T(\tilde{\mathbf{x}}_T) \quad (33)$$

$$\leq \frac{1}{T} 2\beta_T^{1/2} \sum_{t=1}^T \sigma_t(\tilde{\mathbf{x}}_T) \quad (\because \forall \mathbf{x} \in \mathcal{X}, \forall t \in [T], \sigma_t(\mathbf{x}) \geq \sigma_T(\mathbf{x})) \quad (34)$$

$$\stackrel{\text{False}}{=} O\left(\sqrt{\frac{\beta_T \gamma_T}{T}}\right), \quad (35)$$

where $\hat{\mathbf{x}}_T = \arg \max_{\mathbf{x} \in \mathcal{X}} \mu_T(\mathbf{x})$, $\tilde{\mathbf{x}}_T := \arg \max_{\mathbf{x} \in \mathcal{X}} \sigma_T(\mathbf{x})$. This result suggests that (if the equality shown as $\stackrel{\text{False}}{=}$ is true) the simple regret upper bound converges to 0 when $\gamma_T \beta_T = o(T)$ regardless of the BO algorithm. For example, $\gamma_T \beta_T = o(T)$ holds for Gaussian and Matérn- ν kernels with $2\nu > d$ in the Bayesian setting. This result is obviously strange.

B Regret Upper Bounds for GP-EI with Rescaling of Posterior Variance

This section shows the regret upper bounds for GP-EI with rescaling of the posterior variance in the Bayesian setting. The derivation of the regret upper bounds comes mainly from (Wang & de Freitas, 2014) and can be applied to the frequentist setting in almost the same way. For more details on the frequentist setting, we discuss in Appendix B.4. We denote the cumulative regret $R_T := \sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t)$.

B.1 Auxiliary Lemmas

First, to generalize the analysis of (Wang & de Freitas, 2014), we provide the two lemmas modified slightly from (Wang & de Freitas, 2014).

Lemma B.1 (Modified from Lemma 9 of (Wang & de Freitas, 2014)). *Pick $\mathbf{x} \in \mathcal{X}$ and $t \in [T]$. Suppose that $a \in \mathbb{R}$ satisfies the following inequality holds with some $\beta_t^{1/2} > 0$:*

$$|a - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}). \quad (36)$$

Then, the following holds with $b \in \mathbb{R}$ and $\nu > 0$:

$$\max \left\{ (a - b)_+ - \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}), \frac{\tau(-\beta_t^{1/2}/\nu^{1/2})}{\tau(\beta_t^{1/2}/\nu^{1/2})} (a - b)_+ \right\} \quad (37)$$

$$\leq \nu^{1/2} \sigma_{t-1}(\mathbf{x}) \tau \left(\frac{\mu_{t-1}(\mathbf{x}) - b}{\nu^{1/2} \sigma_{t-1}(\mathbf{x})} \right) = \text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2} \sigma_{t-1}(\mathbf{x}), b) \quad (38)$$

$$\leq (a - b)_+ + (\beta_t^{1/2} + \nu^{1/2}) \sigma_{t-1}(\mathbf{x}), \quad (39)$$

where $(\cdot)_+ := \max\{0, \cdot\}$.

Proof. If $\sigma_{t-1}(\mathbf{x}) = 0$, then $a = \mu_{t-1}(\mathbf{x})$ and $\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b) = (a - b)_+$. Therefore, the inequalities are trivial (Note that $\tau(\cdot)$ is a monotonically increasing function). Thus, we assume that $\sigma_{t-1}(\mathbf{x}) > 0$ hereafter. Furthermore, define the variables q and u as follows:

$$q = \frac{a - b}{\sigma_{t-1}(\mathbf{x})}, u = \frac{\mu_{t-1}(\mathbf{x}) - b}{\sigma_{t-1}(\mathbf{x})}. \quad (40)$$

First, we derive the upper bound. From the definition,

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b) = \nu^{1/2}\sigma_{t-1}(\mathbf{x})\tau\left(\frac{u}{\nu^{1/2}}\right). \quad (41)$$

From the assumption, $|u - q| \leq \beta_t^{1/2}$. Since $\tau(\cdot)$ is monotonically increasing and $\tau(c) \leq 1 + c$ for $c > 0$, we obtain

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b) \leq \nu^{1/2}\sigma_{t-1}(\mathbf{x})\tau\left(\frac{(q)_+ + \beta_t^{1/2}}{\nu^{1/2}}\right) \quad (42)$$

$$\leq \nu^{1/2}\sigma_{t-1}(\mathbf{x})\left(\frac{(q)_+ + \beta_t^{1/2}}{\nu^{1/2}} + 1\right) \quad (43)$$

$$\leq (a - b)_+ + (\beta_t^{1/2} + \nu^{1/2})\sigma_{t-1}(\mathbf{x}). \quad (44)$$

Next, we derive the lower bound. If $(a - b)_+ = 0$, then the lower bound is trivial since the GP-EI acquisition function is non-negative. Therefore, assume $(a - b)_+ > 0$. Since $\tau(\cdot)$ is monotonically increasing, we can see that

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b) \geq \nu^{1/2}\sigma_{t-1}(\mathbf{x})\tau\left(\frac{q - \beta_t^{1/2}}{\nu^{1/2}}\right) \quad (\because q - \beta_t^{1/2} \leq u) \quad (45)$$

$$\geq \nu^{1/2}\sigma_{t-1}(\mathbf{x})\left(\frac{q - \beta_t^{1/2}}{\nu^{1/2}}\right) \quad (\because \tau(c) = c + \tau(-c) \geq c) \quad (46)$$

$$= (a - b)_+ - \beta_t^{1/2}\sigma_{t-1}(\mathbf{x}). \quad (47)$$

Furthermore, since $\tau(\cdot)$ is monotonically increasing,

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b) = \nu^{1/2}\sigma_{t-1}(\mathbf{x})\tau\left(\frac{\mu_{t-1}(\mathbf{x}) - a + a - b}{\nu^{1/2}\sigma_{t-1}(\mathbf{x})}\right) \quad (48)$$

$$\geq \nu^{1/2}\sigma_{t-1}(\mathbf{x})\tau\left(\frac{\mu_{t-1}(\mathbf{x}) - a}{\nu^{1/2}\sigma_{t-1}(\mathbf{x})}\right) \quad (\because a - b > 0 \text{ from assumption}) \quad (49)$$

$$\geq \nu^{1/2}\sigma_{t-1}(\mathbf{x})\tau\left(\frac{-\beta_t^{1/2}}{\nu^{1/2}}\right) \quad (\because \mu_{t-1}(\mathbf{x}) - a > -\beta_t^{1/2}\sigma_{t-1}(\mathbf{x})). \quad (50)$$

Hence, we obtain $\frac{\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b)}{\nu^{1/2}\tau(-\beta_t^{1/2}/\nu^{1/2})} \geq \sigma_{t-1}(\mathbf{x})$. By substituting this inequality to

$\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b) \geq (a - b)_+ - \beta_t^{1/2}\sigma_{t-1}(\mathbf{x})$, we can see that

$$\text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2}\sigma_{t-1}(\mathbf{x}), b) \geq \frac{\nu^{1/2}\tau(-\beta_t^{1/2}/\nu^{1/2})}{\beta_t^{1/2} + \nu^{1/2}\tau(-\beta_t^{1/2}/\nu^{1/2})}(a - b)_+ \quad (51)$$

$$= \frac{\nu^{1/2}\tau(-\beta_t^{1/2}/\nu^{1/2})}{\nu^{1/2}\tau(\beta_t^{1/2}/\nu^{1/2})}(a - b)_+. \quad (52)$$

□

Lemma B.2 (Modified from Lemma 10 of (Wang & de Freitas, 2014)). *Assume the kernel function $k(\mathbf{x}, \mathbf{x}') \leq 1$ for all $\mathbf{x}, \mathbf{x}' \in \mathcal{X}$. Let $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \nu^{1/2} \sigma_{t-1}(\mathbf{x}), b)$ with $\min_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x}) \leq b \leq \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x})$ and some $\nu > 0$. Then, the following inequality holds:*

$$b - \mu_{t-1}(\mathbf{x}_t) \leq \sqrt{\log(t-1 + \sigma^2) - \log(\sigma^2)} \nu^{1/2} \sigma_{t-1}(\mathbf{x}_t). \quad (53)$$

Proof. If $b \leq \mu_{t-1}(\mathbf{x}_t)$, then the inequality is trivial. Thus, we assume $b > \mu_{t-1}(\mathbf{x}_t)$ hereafter. From the assumption of b , there exists $\mathbf{z} \in \mathcal{X}$ such that $\mu_{t-1}(\mathbf{z}) \geq b$. Therefore, we obtain

$$\nu^{1/2} \sigma_{t-1}(\mathbf{z}) \tau(0) \leq \nu^{1/2} \sigma_{t-1}(\mathbf{z}) \tau \left(\frac{\mu_{t-1}(\mathbf{z}) - b}{\nu^{1/2} \sigma_{t-1}(\mathbf{z})} \right) \quad (54)$$

$$= \text{EI}(\mu_{t-1}(\mathbf{z}), \nu^{1/2} \sigma_{t-1}(\mathbf{z}), b) \quad (55)$$

$$\leq \text{EI}(\mu_{t-1}(\mathbf{x}_t), \nu^{1/2} \sigma_{t-1}(\mathbf{x}_t), b) \quad (56)$$

$$= \nu^{1/2} \sigma_{t-1}(\mathbf{x}_t) \tau \left(\frac{\mu_{t-1}(\mathbf{x}_t) - b}{\nu^{1/2} \sigma_{t-1}(\mathbf{x}_t)} \right). \quad (57)$$

Then, since we assume $b > \mu_{t-1}(\mathbf{x}_t)$ and $\tau(0) = \frac{1}{\sqrt{2\pi}}$, we can arrange the inequality as follows:

$$\frac{\sigma_{t-1}(\mathbf{z})}{\sqrt{2\pi}} \leq \sigma_{t-1}(\mathbf{x}_t) \phi \left(\frac{\mu_{t-1}(\mathbf{x}_t) - b}{\nu^{1/2} \sigma_{t-1}(\mathbf{x}_t)} \right) \quad (\because \tau(c) < \phi(c) \text{ for } c < 0) \quad (58)$$

$$= \frac{\sigma_{t-1}(\mathbf{x}_t)}{\sqrt{2\pi}} \exp \left(-\frac{1}{2} \left(\frac{\mu_{t-1}(\mathbf{x}_t) - b}{\nu^{1/2} \sigma_{t-1}(\mathbf{x}_t)} \right)^2 \right). \quad (59)$$

Hence, by applying the logarithm, we can see that

$$|\mu_{t-1}(\mathbf{x}_t) - b| \leq \sqrt{2 \log \left(\frac{\sigma_{t-1}(\mathbf{x}_t)}{\sigma_{t-1}(\mathbf{z})} \right)} \nu^{1/2} \sigma_{t-1}(\mathbf{x}_t) \quad (60)$$

$$= \sqrt{\log \left(\frac{\sigma_{t-1}^2(\mathbf{x}_t)}{\sigma_{t-1}^2(\mathbf{z})} \right)} \nu^{1/2} \sigma_{t-1}(\mathbf{x}_t). \quad (61)$$

We know that, for all $\mathbf{x} \in \mathcal{X}$, the posterior variance is bounded as $\sigma^2/(\sigma^2 + t - 1) \leq \sigma_{t-1}^2(\mathbf{x}) \leq 1$ from Lemma D.1 and monotone decreasing property of the posterior variance. By substituting $\sigma_{t-1}^2(\mathbf{x}_t) \leq 1$ and $\sigma^2/(\sigma^2 + t - 1) \leq \sigma_{t-1}^2(\mathbf{z})$, we can obtain the desired result. \square

B.2 Regret Bounds for GP-EI with Maximum Posterior Mean

In this section, although (Wang & de Freitas, 2014) use the parameters $\{\nu_t\}_{t \geq 1}$ to scale the posterior variance, we use $\{\beta_t(\delta)\}_{t \geq 1}$ instead of $\{\nu_t\}_{t \geq 1}$ for simplicity. We believe that this modification does not affect the essential order of the resulting regret upper bound.

Theorem B.3. *Let $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$. Pick $\delta \in (0, 1)$. Assume \mathcal{X} is a finite set or Assumption 2.1 holds.*

(i) *When \mathcal{X} is a finite set, set $\beta_t(\delta) = 2 \log(|\mathcal{X}| t^2 \pi^2 / (6\delta))$.*

(ii) *When Assumption 2.1 holds, set $\beta_t(\delta) = 2 \log(\pi^2 t^2 / (3\delta)) + 2d \log \left(\left[b d t^2 \sqrt{\log(4ad/\delta)} \right] + 1 \right)$.*

If $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \beta_t^{1/2}(\delta) \sigma_{t-1}(\mathbf{x}), \mu_{t-1}^{\max})$, where $\mu_{t-1}^{\max} = \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x})$, then the cumulative regret is bounded from above with probability at least $1 - \delta$ as follows:

$$R_T = O \left(\sqrt{\beta_T(\delta) \gamma_T T \log T} \right). \quad (62)$$

Proof. First, we provide the proof regarding case (i). In this case, from Lemma 5.1 of (Srinivas et al., 2010), the following holds:

$$\Pr\left(\forall \mathbf{x} \in \mathcal{X}, \forall t \in [T], |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x})\right) \geq 1 - \delta. \quad (63)$$

Therefore, we assume $|f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x})$ holds for all $\mathbf{x} \in \mathcal{X}$ and $t \in [T]$ hereafter. We can see that the instantaneous regret r_t can be bounded as follows:

$$r_t = f(\mathbf{x}^*) - f(\mathbf{x}_t) \quad (64)$$

$$= f(\mathbf{x}^*) - \mu_{t-1}^{\max} + \mu_{t-1}^{\max} - f(\mathbf{x}_t) \quad (65)$$

$$\leq I_t(\mathbf{x}^*) + \mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t) + \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t), \quad (66)$$

where $I_t(\mathbf{x}) = (f(\mathbf{x}) - \mu_{t-1}^{\max})_+$. Regarding the term $I_t(\mathbf{x}^*)$, using Lemma B.1 with $a = f(\mathbf{x}^*)$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t(\delta)$, we see that

$$I_t(\mathbf{x}^*) \leq C_3 \text{EI}(\mu_{t-1}(\mathbf{x}^*), \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}^*), \mu_{t-1}^{\max}) \quad (67)$$

$$\leq C_3 \text{EI}(\mu_{t-1}(\mathbf{x}_t), \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t), \mu_{t-1}^{\max}), \quad (\mathbf{x}_t \text{ is the maximizer}) \quad (68)$$

where $C_3 = \frac{\tau(1)}{\tau(-1)}$. Moreover, by applying Lemma B.1 again with $a = \mu_{t-1}(\mathbf{x}_t)$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t(\delta)$, we can obtain

$$I_t(\mathbf{x}^*) \leq C_3 \left((\mu_{t-1}(\mathbf{x}_t) - \mu_{t-1}^{\max})_+ + 2\beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t) \right) \quad (69)$$

$$\leq 2C_3\beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t). \quad (\because \mu_{t-1}^{\max} > \mu_{t-1}(\mathbf{x}_t)). \quad (70)$$

Furthermore, the term $\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)$ is bounded from above by Lemma B.2:

$$\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t) \leq \sqrt{\log(t-1+\sigma^2) - \log(\sigma^2)}\beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t). \quad (71)$$

Hence, aggregating the results, we can obtain

$$r_t \leq \left(2C_3 + \sqrt{\log(t-1+\sigma^2) - \log(\sigma^2)} + 1\right) \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t). \quad (72)$$

Then, the cumulative regret can be bounded from above as follows:

$$R_T = \sum_{t=1}^T r_t \quad (73)$$

$$\leq \sum_{t=1}^T \left(2C_3 + \sqrt{\log(t-1+\sigma^2) - \log(\sigma^2)} + 1\right) \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t) \quad (74)$$

$$\leq \left(2C_3 + \sqrt{\log(T-1+\sigma^2) - \log(\sigma^2)} + 1\right) \sum_{t=1}^T \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t) \quad (75)$$

$$\leq \left(2C_3 + \sqrt{\log(T-1+\sigma^2) - \log(\sigma^2)} + 1\right) \sqrt{\sum_{t=1}^T \beta_t(\delta) \sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}_t)} \quad (\because \text{Cauchy-Schwarz inequality}) \quad (76)$$

$$\leq \left(2C_3 + \sqrt{\log(T-1+\sigma^2) - \log(\sigma^2)} + 1\right) \sqrt{C_1 T \beta_T(\delta) \gamma_T}. \quad (77)$$

In the last inequality, we use $\sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}_t) \leq C_1 \gamma_T$ (Srinivas et al., 2010).

For the case (ii), we consider a finite set \mathcal{X}_t with each dimension evenly divided by $\tau_t = \lceil bdr t^2 \sqrt{\log(4ad/\delta)} \rceil$, such that $|\mathcal{X}_t| = \tau_t^d$. In addition, we denote $[\mathbf{x}]_t = \arg \min_{\mathbf{x}' \in \mathcal{X}_t} \|\mathbf{x} - \mathbf{x}'\|_1$. For this \mathcal{X}_t , using the union

bound and Lemmas 5.5, 5.6, and 5.7 of (Srinivas et al., 2010), the following events simultaneously hold with probability at least $1 - \delta$,

$$\forall t \in [T], \sup_{\mathbf{x} \in \mathcal{X}} |f(\mathbf{x}) - f([\mathbf{x}]_t)| \leq \frac{1}{t^2}, \quad (78)$$

$$\forall t \in [T], \forall \mathbf{x} \in \mathcal{X}_t \cup \{\mathbf{x}_t\}, |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}), \quad (79)$$

since $\beta_t(\delta) \geq 2 \log\left(\frac{(|\mathcal{X}_t|+1)\pi^2 t^2}{3\delta}\right)$. Note that $+1$ in $|\mathcal{X}_t| + 1$ is required to guarantee the confidence bound on \mathbf{x}_t . Therefore, we assume the above two events hold hereafter.

Since the above two events hold, the instantaneous regret can be bounded as

$$r_t = f(\mathbf{x}^*) - f(\mathbf{x}_t) \quad (80)$$

$$= f(\mathbf{x}^*) - f([\mathbf{x}^*]_t) + f([\mathbf{x}^*]_t) - \mu_{t-1}^{\max} + \mu_{t-1}^{\max} - f(\mathbf{x}_t) \quad (81)$$

$$\leq \frac{1}{t^2} + (f([\mathbf{x}^*]_t) - \mu_{t-1}^{\max})_+ + \mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t) + \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t). \quad (82)$$

As with the case (i), using Lemma B.1 at $[\mathbf{x}^*]_t$ with $a = f([\mathbf{x}^*]_t)$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t(\delta)$, we see that

$$(f([\mathbf{x}^*]_t) - \mu_{t-1}^{\max})_+ \leq C_3 \text{EI}(\mu_{t-1}([\mathbf{x}^*]_t), \beta_t^{1/2}(\delta)\sigma_{t-1}([\mathbf{x}^*]_t), \mu_{t-1}^{\max}) \quad (83)$$

$$\leq C_3 \text{EI}(\mu_{t-1}(\mathbf{x}_t), \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t), \mu_{t-1}^{\max}). \quad (\mathbf{x}_t \text{ is the maximizer}) \quad (84)$$

Moreover, using Lemma B.1 again at \mathbf{x}_t with $a = \mu_{t-1}(\mathbf{x}_t)$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t(\delta)$, we can obtain

$$f([\mathbf{x}^*]_t) - \mu_{t-1}^{\max} \leq C_3 \left((\mu_{t-1}(\mathbf{x}_t) - \mu_{t-1}^{\max})_+ + 2\beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t) \right) \quad (85)$$

$$\leq 2C_3\beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t). \quad (86)$$

Futhermore, using Lemma B.2,

$$\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t) \leq \sqrt{\log(t-1+\sigma^2) - \log(\sigma^2)}\beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x}_t). \quad (87)$$

In addition, $\sum_{t=1}^T 1/t^2 \leq \pi^2/6$. Therefore, repeating the same proof as case (i), we obtain

$$R_T \leq \frac{\pi^2}{6} + \left(2C_3 + \sqrt{\log(T-1+\sigma^2) - \log(\sigma^2)} + 1\right) \sqrt{C_1 T \beta_T(\delta) \gamma_T}. \quad (88)$$

□

Theorem B.4. Let $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$. Assume \mathcal{X} is a finite set or Assumption 2.1 holds.

(i) When \mathcal{X} is a finite set, set $\beta_t = 2 \log(|\mathcal{X}|t^2/\sqrt{2\pi} + 1)$.

(ii) When Assumption 2.1 holds, set $\beta_t = 2d \log(\lceil bdr t^2(\sqrt{\log(ad)} + \sqrt{\pi}/2) \rceil) + 2 \log(t^2/\sqrt{2\pi} + 1)$.

If $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \beta_t^{1/2}\sigma_{t-1}(\mathbf{x}), \mu_{t-1}^{\max})$, where $\mu_{t-1}^{\max} = \max_{\mathbf{x} \in \mathcal{X}} \mu_{t-1}(\mathbf{x})$, then the BCR is bounded from above as follows:

$$\text{BCR}_T = O\left(\sqrt{\beta_T \gamma_T T \log T}\right). \quad (89)$$

Proof. First, we consider the case (i).

$$\text{BCR}_T = \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - f(\mathbf{x}_t)] \quad (90)$$

$$= \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - U_t(\mathbf{x}^*) + U_t(\mathbf{x}^*) - \mu_{t-1}^{\max} + \mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)], \quad (91)$$

$$\leq \sum_{t=1}^T \mathbb{E}\left[f(\mathbf{x}^*) - U_t(\mathbf{x}^*) + (U_t(\mathbf{x}^*) - \mu_{t-1}^{\max})_+ + \mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)\right] \quad (92)$$

$$= \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - U_t(\mathbf{x}^*)] + \sum_{t=1}^T \mathbb{E}\left[(U_t(\mathbf{x}^*) - \mu_{t-1}^{\max})_+\right] + \sum_{t=1}^T \mathbb{E}[\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)], \quad (93)$$

where $U_t(\mathbf{x}) = \mu_{t-1}(\mathbf{x}) + \beta_t^{1/2} \sigma_{t-1}(\mathbf{x})$. Then, from the same proof as Theorem B.1 of (Takeno et al., 2023), $\sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - U_t(\mathbf{x}^*)] \leq \pi^2/6$. Note that $\beta_t > 0$ due to the ceil function.

Regarding the term $(U_t(\mathbf{x}^*) - \mu_{t-1}^{\max})_+$, we can apply Lemma B.1 at \mathbf{x}^* with $a = U_t(\mathbf{x}^*)$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t$. Therefore,

$$(U_t(\mathbf{x}^*) - \mu_{t-1}(\mathbf{x}_t))_+ \leq C_3 \text{EI}\left(\mu_{t-1}(\mathbf{x}^*), \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}^*), \mu_{t-1}^{\max}\right) \quad (94)$$

$$\leq C_3 \text{EI}\left(\mu_{t-1}(\mathbf{x}_t), \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t), \mu_{t-1}^{\max}\right), \quad (95)$$

where $C_3 = \frac{\tau(1)}{\tau(-1)}$. Moreover, using Lemma B.1 at \mathbf{x}_t with $a = \mu_{t-1}(\mathbf{x})$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t$, we obtain

$$(U_t(\mathbf{x}^*) - \mu_{t-1}(\mathbf{x}_t))_+ \leq C_3 \left((\mu_{t-1}(\mathbf{x}_t) - \mu_{t-1}^{\max})_+ + 2\beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \right) = 2C_3 \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t). \quad (96)$$

Furthermore, regarding $\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)$, using Lemma B.2,

$$\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t) \leq \sqrt{\log(t-1 + \sigma^2) - \log(\sigma^2)} \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t). \quad (97)$$

Hence, the sum of these terms can be bounded from above as with the proofs of Theorem B.3:

$$\sum_{t=1}^T \mathbb{E}\left[(U_t(\mathbf{x}^*) - \mu_{t-1}^{\max})_+\right] + \sum_{t=1}^T \mathbb{E}[\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)] \leq \left(2C_3 + \sqrt{\log(T-1 + \sigma^2) - \log(\sigma^2)} + 1\right) \sqrt{C_1 T \beta_T \gamma_T}. \quad (98)$$

Consequently, we can obtain the following upper bound:

$$\text{BCR}_T \leq \frac{\pi^2}{6} + \left(2C_3 + \sqrt{\log(T-1 + \sigma^2) - \log(\sigma^2)} + 1\right) \sqrt{C_1 T \beta_T \gamma_T}. \quad (99)$$

Next, for the case (ii),

$$\text{BCR}_T = \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - f(\mathbf{x}_t)] \quad (100)$$

$$= \sum_{t=1}^T \mathbb{E}\left[f(\mathbf{x}^*) - f([\mathbf{x}^*]_t) + f([\mathbf{x}^*]_t) - U_t([\mathbf{x}^*]_t) + U_t([\mathbf{x}^*]_t) - \mu_{t-1}^{\max} + \mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)\right], \quad (101)$$

$$\leq \sum_{t=1}^T \mathbb{E}[f(\mathbf{x}^*) - f([\mathbf{x}^*]_t)] + \sum_{t=1}^T \mathbb{E}[f([\mathbf{x}^*]_t) - U_t([\mathbf{x}^*]_t)] \quad (102)$$

$$+ \sum_{t=1}^T \mathbb{E}\left[(U_t([\mathbf{x}^*]_t) - \mu_{t-1}^{\max})_+\right] + \sum_{t=1}^T \mathbb{E}[\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t)], \quad (103)$$

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

We consider a finite set \mathcal{X}_t with each dimension evenly divided by $\tau_t = \lceil bdr t^2 (\sqrt{\log(ad)} + \sqrt{\pi}/2) \rceil$, such that $|\mathcal{X}_t| = \tau_t^d$. In addition, we denote $[\mathbf{x}]_t = \arg \min_{\mathbf{x}' \in \mathcal{X}_t} \|\mathbf{x} - \mathbf{x}'\|_1$. Then, as with the proof of Theorem B.1 of (Takeno et al., 2023), we can obtain

$$\sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - f([\mathbf{x}^*]_t)] \leq \frac{\pi^2}{6}, \quad (104)$$

$$\sum_{t=1}^T \mathbb{E} [f([\mathbf{x}^*]_t) - U_t([\mathbf{x}^*]_t)] \leq \frac{\pi^2}{6}. \quad (105)$$

As with the case (i), using Lemma B.1 at $[\mathbf{x}^*]_t$ with $a = U_t([\mathbf{x}^*]_t)$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t$, we see that

$$(U_t([\mathbf{x}^*]_t) - \mu_{t-1}^{\max})_+ \leq C_3 \text{EI}(\mu_{t-1}([\mathbf{x}^*]_t), \beta_t^{1/2} \sigma_{t-1}([\mathbf{x}^*]_t), \mu_{t-1}^{\max}) \quad (106)$$

$$\leq C_3 \text{EI}(\mu_{t-1}(\mathbf{x}_t), \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t), \mu_{t-1}^{\max}). \quad (\mathbf{x}_t \text{ is the maximizer}) \quad (107)$$

Moreover, using Lemma B.1 again at \mathbf{x}_t with $a = \mu_{t-1}(\mathbf{x}_t)$, $b = \mu_{t-1}^{\max}$, and $\nu = \beta_t$, we can obtain

$$(U_t([\mathbf{x}^*]_t) - \mu_{t-1}^{\max})_+ \leq C_3 \left((\mu_{t-1}(\mathbf{x}_t) - \mu_{t-1}^{\max})_+ + 2\beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \right) \quad (108)$$

$$\leq 2C_3 \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t). \quad (109)$$

Futhermore, using Lemma B.2,

$$\mu_{t-1}^{\max} - \mu_{t-1}(\mathbf{x}_t) \leq \sqrt{\log(t-1 + \sigma^2) - \log(\sigma^2)} \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t). \quad (110)$$

Therefore, repeating the same proof as Theorem B.3, we obtain

$$\text{BCR}_T \leq \frac{\pi^2}{3} + \left(2C_3 + \sqrt{\log(T-1 + \sigma^2) - \log(\sigma^2)} + 1 \right) \sqrt{C_1 T \beta_T \gamma_T}. \quad (111)$$

□

B.3 Regret Bounds for GP-EI with Maximum Posterior Mean among Evaluated Points

First, we show a variant of Theorem B.3, where the reference value is changed to $\tilde{\mu}_{t-1}^{\max} = \max_{i \in [t-1]} \mu_{t-1}(\mathbf{x}_i)$:

Theorem B.5. *Let $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$. Pick $\delta \in (0, 1)$. Assume \mathcal{X} is a finite set or Assumption 2.1 holds.*

(i) *When \mathcal{X} is a finite set, set $\beta_t(\delta) = 2 \log(|\mathcal{X}| t^2 \pi^2 / (6\delta))$.*

(ii) *When Assumption 2.1 holds, set $\beta_t(\delta) = 2 \log(\pi^2 t^2 / (3\delta)) + 2d \log \left(\left\lceil bdr t^2 \sqrt{\log(4ad/\delta)} \right\rceil + t \right)$.*

If $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \beta_t^{1/2}(\delta) \sigma_{t-1}(\mathbf{x}), \tilde{\mu}_{t-1}^{\max})$, where $\tilde{\mu}_{t-1}^{\max} = \max_{i \in [t-1]} \mu_{t-1}(\mathbf{x}_i)$ for $t \geq 1$ and $\tilde{\mu}_0^{\max} = 0$ for $t = 1$, then the cumulative regret is bounded from above with probability at least $1 - \delta$ as follows:

$$R_T = O \left(\sqrt{\beta_T(\delta) \gamma_T T \log T} \right). \quad (112)$$

Proof. As with the proof of Theorem B.3, we can guarantee the following events with probability at least $1 - \delta$:

$$\forall \mathbf{x} \in \mathcal{X}, |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2}(\delta) \sigma_{t-1}(\mathbf{x}), \quad (113)$$

for the case (i), and

$$\forall t \in [T], \sup_{\mathbf{x} \in \mathcal{X}} |f(\mathbf{x}) - f([\mathbf{x}]_t)| \leq \frac{1}{t^2}, \quad (114)$$

$$\forall t \in [T], \forall \mathbf{x} \in \mathcal{X}_t \cup \{\mathbf{x}_i\}_{i=1}^t, |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2}(\delta) \sigma_{t-1}(\mathbf{x}), \quad (115)$$

$$(116)$$

for the case (ii). Note that, in the case (ii), we additionally consider the bounds on $\{\mathbf{x}_i\}_{i=1}^t$, not only \mathbf{x}_t , using $\beta_t(\delta) = \log\left(\frac{(|\mathcal{X}_t|+t)t^2\pi^2}{3\delta}\right)$ compared with the proof of Theorem B.3. In the remainder of the proof, we assume the above events hold in each case.

By almost the same proof as Theorem B.3, in which we set $a = f(\mathbf{x}_t)$ for a second application of Lemma B.1, we can obtain

$$r_t \leq C_3 \tilde{I}_t(\mathbf{x}_t) + \left(2C_3 + \sqrt{\log(t-1 + \sigma^2) - \log(\sigma^2)} + 1\right) \beta_t^{1/2}(\delta) \sigma_{t-1}(\mathbf{x}_t), \quad \text{for the case (i),} \quad (117)$$

$$r_t \leq \frac{1}{t^2} + C_3 \tilde{I}_t(\mathbf{x}_t) + \left(2C_3 + \sqrt{\log(t-1 + \sigma^2) - \log(\sigma^2)} + 1\right) \beta_t^{1/2}(\delta) \sigma_{t-1}(\mathbf{x}_t), \quad \text{for the case (ii),} \quad (118)$$

where $\tilde{I}_t(\mathbf{x}_t) = (f(\mathbf{x}_t) - \tilde{\mu}_{t-1}^{\max})_+$ and $C_3 = \frac{\tau(1)}{\tau(-1)}$. Thus, the cumulative regret is bounded from above as follows:

$$R_T \leq C_3 \sum_{t=1}^T \tilde{I}_t(\mathbf{x}_t) + \left(2C_3 + \sqrt{\log(T-1 + \sigma^2) - \log(\sigma^2)} + 1\right) \sqrt{C_1 T \beta_T(\delta) \gamma_T}, \quad \text{for the case (i),} \quad (119)$$

$$R_T \leq \frac{\pi^2}{6} + C_3 \sum_{t=1}^T \tilde{I}_t(\mathbf{x}_t) + \left(2C_3 + \sqrt{\log(T-1 + \sigma^2) - \log(\sigma^2)} + 1\right) \sqrt{C_1 T \beta_T(\delta) \gamma_T}, \quad \text{for the case (ii).} \quad (120)$$

Next, we obtain the upper bound of $\sum_{t=1}^T \tilde{I}_t(\mathbf{x}_t)$ by the proof modified slightly from (Tran-The et al., 2022; Hu et al., 2025). First, we can see that

$$\tilde{I}_1(\mathbf{x}_1) \leq (f(\mathbf{x}_1) - 0)_+ \leq \beta_1^{1/2}(\delta), \quad (121)$$

where we use the inequality $|f(\mathbf{x}_1)| \leq \beta_1^{1/2}(\delta) \sqrt{k(\mathbf{x}_1, \mathbf{x}_1)} \leq \beta_1^{1/2}(\delta)$ for both cases (i) and (ii). Then, define an indices $\mathcal{M} = \{m(i) \mid \tilde{I}_{m(i)}(\mathbf{x}_{m(i)}) > 0, m(i) > 1\}$ with ascending order $m(1) < \dots < m(M)$, where $M := |\mathcal{M}|$. In addition, let $m(0) = m(1) - 1 \geq 1$. If $M = 0$, we can easily show $\sum_{t=1}^T \tilde{I}_t(\mathbf{x}_t) = \tilde{I}_1(\mathbf{x}_1) \leq \beta_1^{1/2}(\delta)$. On the other hand, for the case of $M \geq 1$, we can obtain

$$\sum_{t=2}^T \tilde{I}_t(\mathbf{x}_t) = \sum_{i=1}^M f(\mathbf{x}_{m(i)}) - \tilde{\mu}_{m(i)-1}^{\max} \quad (122)$$

$$\leq \sum_{i=1}^M f(\mathbf{x}_{m(i)}) - \mu_{m(i)-1}(\mathbf{x}_{m(i-1)}) \quad (\because \text{Definition of } \tilde{\mu}_{t-1}^{\max}) \quad (123)$$

$$\leq \sum_{i=1}^M f(\mathbf{x}_{m(i)}) - f(\mathbf{x}_{m(i-1)}) + \beta_{m(i)}^{1/2}(\delta) \sigma_{m(i)-1}(\mathbf{x}_{m(i-1)}) \quad (124)$$

$$\leq f(\mathbf{x}_{m(M)}) - f(\mathbf{x}_{m(0)}) + \sum_{i=1}^M \beta_{m(i)}^{1/2}(\delta) \sigma_{m(i)-1}(\mathbf{x}_{m(i-1)}) \quad (125)$$

$$\leq 2\beta_1^{1/2}(\delta) + 2 + \sum_{t=1}^T \beta_t^{1/2}(\delta) \sigma_{t-1}(\mathbf{x}_t) \quad (126)$$

$$\leq 2\beta_1^{1/2}(\delta) + 2 + \sqrt{C_1 T \beta_T(\delta) \gamma_T}, \quad (127)$$

where we use

- in the second inequality it is guaranteed that $\forall t \in [T], \forall \mathbf{x} \in \mathcal{X}, |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2}(\delta)$ for the case (i), and $\forall t \in [T], \forall \mathbf{x} \in \mathcal{X}_t \cup \{\mathbf{x}_i\}_{i=1}^t, |f(\mathbf{x}) - \mu_{t-1}(\mathbf{x})| \leq \beta_t^{1/2}(\delta)\sigma_{t-1}(\mathbf{x})$ for the case (ii) from the assumption;
- in the third inequality, $\sigma_{m(i)-1}(\mathbf{x}_{m(i)-1}) \leq \sigma_{m(i-1)-1}(\mathbf{x}_{m(i-1)})$;
- in the fourth inequality, for all $\mathbf{x} \in \mathcal{X}, |f(\mathbf{x})| \leq \beta_1^{1/2}$ for the case (i), and $|f(\mathbf{x})| \leq |f([\mathbf{x}]_1)| + |f(\mathbf{x}) - f([\mathbf{x}]_1)| \leq \beta_1^{1/2} + 1$ for the case (ii) from the assumption.

Consequently, $\sum_{t=1}^T \tilde{I}_t(\mathbf{x}_t) \leq 3\beta_1^{1/2}(\delta) + 2 + \sqrt{C_1 T \beta_T(\delta) \gamma_T} = O(\sqrt{\beta_T(\delta) \gamma_T T})$, which concludes the proof. \square

Theorem B.6. *Let $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$. Assume \mathcal{X} is a finite set or Assumption 2.1 holds.*

(i) *When \mathcal{X} is a finite set, set $\beta_t = 2 \log(|\mathcal{X}|t^2/\sqrt{2\pi} + 1)$.*

(ii) *When Assumption 2.1 holds, set $\beta_t = 2d \log(\lceil bdr t^2(\sqrt{\log(ad)} + \sqrt{\pi}/2) \rceil) + 2 \log(t^2/\sqrt{2\pi} + 1)$.*

If $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}), \tilde{\mu}_{t-1}^{\max})$, where $\tilde{\mu}_{t-1}^{\max} = \max_{i \in [t-1]} \mu_{t-1}(\mathbf{x}_i)$ for $t \geq 1$ and $\tilde{\mu}_0^{\max} = 0$ for $t = 1$, then the BCR is bounded from above as follows:

$$\text{BCR}_T = O\left(\sqrt{\beta_T \gamma_T T \log T}\right). \quad (128)$$

Proof. By the same proof as Theorem B.4, for both cases (i) and (ii), we can see that

$$\text{BCR}_T = C_3 \sum_{t=1}^T \mathbb{E}[\tilde{I}_t(\mathbf{x}_t)] + O\left(\sqrt{\beta_T \gamma_T T \log T}\right). \quad (129)$$

where $\tilde{I}_t(\mathbf{x}) = (f(\mathbf{x}) - \tilde{\mu}_{t-1}^{\max})_+$ and $C_3 = \tau(1)/\tau(-1)$. Therefore, we here consider the upper bound of $\sum_{t=1}^T \mathbb{E}[\tilde{I}_t(\mathbf{x}_t)]$.

First, we can obtain

$$\mathbb{E}[\tilde{I}_1(\mathbf{x}_1)] = \mathbb{E}[(f(\mathbf{x}_1))_+] = 2\phi(0)\sqrt{k(\mathbf{x}_1, \mathbf{x}_1)} \leq \sqrt{\frac{2}{\pi}}, \quad (130)$$

as $f(\mathbf{x}_1) \sim \mathcal{N}(0, k(\mathbf{x}_1, \mathbf{x}_1))$. Then, define an indices $\mathcal{M} = \{m(i) \mid \tilde{I}_{m(i)}(\mathbf{x}_{m(i)}) > 0, m(i) > 1\}$ with ascending order $m(1) < \dots < m(M)$, where $M := |\mathcal{M}|$. In addition, let $m(0) = m(1) - 1 \geq 1$. If $M = 0$, we can easily show $\sum_{t=1}^T \tilde{I}_t(\mathbf{x}_t) = \tilde{I}_1(\mathbf{x}_1) \leq \sqrt{\frac{2}{\pi}}$. Thus, we consider the case of $M \geq 1$. For any \mathcal{M} under $M \geq 1$, we can

obtain

$$\sum_{t=2}^T \tilde{I}_t(\mathbf{x}_t) = \sum_{i=1}^M f(\mathbf{x}_{m(i)}) - \tilde{\mu}_{m(i)-1}^{\max} \quad (131)$$

$$\leq \sum_{i=1}^M f(\mathbf{x}_{m(i)}) - \mu_{m(i)-1}(\mathbf{x}_{m(i-1)}) \quad (132)$$

$$= \sum_{i=1}^M f(\mathbf{x}_{m(i)}) - f(\mathbf{x}_{m(i-1)}) + f(\mathbf{x}_{m(i-1)}) - U_{m(i)}(\mathbf{x}_{m(i-1)}) + \beta_{m(i)}^{1/2} \sigma_{m(i)-1}(\mathbf{x}_{m(i-1)}) \quad (133)$$

$$\leq f(\mathbf{x}_{m(M)}) - f(\mathbf{x}_{m(0)}) + \sum_{i=1}^M f(\mathbf{x}_{m(i-1)}) - U_{m(i)}(\mathbf{x}_{m(i-1)}) + \sum_{i=1}^M \beta_{m(i)}^{1/2} \sigma_{m(i)-1}(\mathbf{x}_{m(i-1)}) \quad (134)$$

$$\leq \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) + \sum_{i=1}^M (f(\mathbf{x}_{m(i-1)}) - U_{m(i)}(\mathbf{x}_{m(i-1)}))_+ + \sum_{t=1}^T \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \quad (135)$$

$$\leq \max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) + \sum_{t=2}^T \sum_{j=1}^{t-1} (f(\mathbf{x}_j) - U_t(\mathbf{x}_j))_+ + \sum_{t=1}^T \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t), \quad (136)$$

where $U_t(\mathbf{x}) = \mu_{t-1}(\mathbf{x}) + \beta_t^{1/2} \sigma_{t-1}(\mathbf{x})$. In the above transformation;

- in the first inequality, we use the definition of $\tilde{\mu}_{m(i)-1}^{\max}$;
- in the second inequality, we use $\sigma_{m(i)-1}(\mathbf{x}_{m(i-1)}) \leq \sigma_{m(i-1)-1}(\mathbf{x}_{m(i-1)})$.

From Lemma D.2,

$$\mathbb{E} \left[\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \right] \leq 2\sqrt{2 \log(|\mathcal{X}|/2) + 2}, \quad \text{for the case (i),} \quad (137)$$

$$\mathbb{E} \left[\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) - \min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \right] \leq 2\sqrt{2d \log(\lceil bdr(\log(ad) + \sqrt{\pi}/2) \rceil) + 2} - 2 \log 2 + 2, \quad \text{for the case (ii).} \quad (138)$$

In addition, as with the proof of Theorem B.1 of (Takeno et al., 2023), we can obtain

$$\mathbb{E} \left[\sum_{t=2}^T \sum_{j=1}^{t-1} (f(\mathbf{x}_j) - U_t(\mathbf{x}_j))_+ \right] \leq \sum_{t=1}^T \sum_{j=1}^{t-1} \frac{1}{t^2} \quad (139)$$

$$\leq \sum_{t=1}^T \frac{1}{t} \quad (140)$$

$$\leq 1 + \log T. \quad (141)$$

Furthermore, we know that $\sum_{t=1}^T \beta_t^{1/2} \sigma_{t-1}(\mathbf{x}_t) \leq \sqrt{C_1 \beta_T T \gamma_T}$. Consequently, we can observe $\sum_{t=1}^T \mathbb{E}[\tilde{I}_t(\mathbf{x}_t)] = O(\sqrt{C_1 \beta_T T \gamma_T})$. \square

B.4 Discussion on Modification to Frequentist Setting

We focus on the analysis in the Bayesian setting. However, our proof approach for finite input domains can apply to the frequentist setting since, in general, (i) $\sup_{\mathbf{x} \in \mathcal{X}} |f(\mathbf{x})| \leq B$ with some constant B from the assumption and (ii) the confidence bounds for f hold for all $\mathbf{x} \in \mathcal{X}$ even if \mathcal{X} is an infinite input domain (for example, Srinivas et al., 2010; Chowdhury & Gopalan, 2017). Therefore, since β_t is set as $O(\sqrt{\gamma_T})$ to make confidence bounds in the frequentist setting, our proof suggests that GP-EI with $\tilde{\mu}_t^{1/2}$ achieves $O(\sqrt{\beta_T \gamma_T T \log T}) = O(\gamma_T \sqrt{T \log T})$ cumulative regret upper bound with high probability without any other

modification to the GP-EI algorithm like (Hu et al., 2025). To our knowledge, this has not been revealed by the existing studies.

Note that, in the frequentist setting, an extension to the expected regret upper bound from the high-probability regret bound is trivial under some conditions, such as $\sup_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \leq B$ and $\inf_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x}) \geq -B$, which is assumed in most existing studies (for example, Srinivas et al., 2010; Chowdhury & Gopalan, 2017). If we obtain the high probability regret upper bound as $R_T = O(g(T) \log(1/\delta))$ with some increasing sublinear function $g(T) = o(T/\log(T))$, by setting $\delta = 1/T$, we see that

$$\mathbb{E}[R_T] \leq g(T) \log(1/\delta) + \delta \sum_{t=1}^T f(\mathbf{x}^*) - f(\mathbf{x}_t) \leq g(T) \log(T) + 2B, \quad (142)$$

which is sublinear. Similar derivation is often performed in the literature of the bandit algorithms, for example, discussed in (Abbasi-yadkori et al., 2011).

C Proofs for GP-EIMS

Here, we show the proofs of the main paper's theorems, lemmas, and corollaries.

C.1 Proof of Lemma 4.1

Lemma 4.1. *Let*

$$\eta_t := \frac{g_t^* - \mu_{t-1}(\mathbf{x}_t)}{\sigma_{t-1}(\mathbf{x}_t)}. \quad (143)$$

Then, the BCR incurred by GP-EIMS can be bounded from above as follows:

$$\text{BCR}_T \leq \sqrt{\mathbb{E} \left[\sum_{t=1}^T \eta_t^2 \mathbb{1}\{\eta_t \geq 0\} \right]} \sqrt{C_1 \gamma_T}, \quad (144)$$

where $C_1 := 2/\log(1 + \sigma^{-2})$ and the indicator function $\mathbb{1}\{\eta_t \geq 0\} = 1$ if $\eta_t \geq 0$, and 0 otherwise.

Proof. From the tower property and linearity of expectation, we obtain

$$\text{BCR}_T = \sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - f(\mathbf{x}_t)] \quad (145)$$

$$= \sum_{t=1}^T \mathbb{E} [f(\mathbf{x}^*) - g_t^* + g_t^* - f(\mathbf{x}_t)] \quad (146)$$

$$= \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_t [f(\mathbf{x}^*) - g_t^*] + \mathbb{E}_t [g_t^* - f(\mathbf{x}_t)]] \quad (147)$$

$$\stackrel{(a)}{=} \sum_{t=1}^T \mathbb{E}_{\mathcal{D}_{t-1}} [\mathbb{E}_t [g_t^* - f(\mathbf{x}_t)]] \quad (148)$$

$$\stackrel{(b)}{=} \sum_{t=1}^T \mathbb{E} [g_t^* - \mu_{t-1}(\mathbf{x}_t)]. \quad (149)$$

For the equality (a), we use the fact $\mathbb{E}_t[f(\mathbf{x}^*)] = \mathbb{E}_t[g_t^*]$ since $f(\mathbf{x}^*) \mid \mathcal{D}_{t-1}$ and $g_t^* \mid \mathcal{D}_{t-1}$ are identically distributed. For the equality (b), we use the tower property of expectation and $\mathbb{E}_t[f(\mathbf{x}_t)] = \mathbb{E}_t[\mu_{t-1}(\mathbf{x}_t)]$, which holds since \mathbf{x}_t defined via g_t and f are independent. Then, we define the random variable η_t as follows:

$$\eta_t := \frac{g_t^* - \mu_{t-1}(\mathbf{x}_t)}{\sigma_{t-1}(\mathbf{x}_t)}. \quad (150)$$

Therefore, since $\sigma_{t-1}(\mathbf{x}) > 0$ due to $\sigma^2 > 0$, we obtain

$$\text{BCR}_T = \sum_{t=1}^T \mathbb{E} [\eta_t \sigma_{t-1}(\mathbf{x}_t)]. \quad (151)$$

Furthermore, from the linearity of expectation, we can see that

$$\text{BCR}_T = \mathbb{E} \left[\sum_{t=1}^T \eta_t \sigma_{t-1}(\mathbf{x}_t) \right] \quad (152)$$

$$\leq \mathbb{E} \left[\sum_{t=1}^T \eta_t \mathbb{1}\{\eta_t \geq 0\} \sigma_{t-1}(\mathbf{x}_t) \right] \quad (\because \sigma_{t-1}(\mathbf{x}_t) > 0) \quad (153)$$

$$\leq \mathbb{E} \left[\sqrt{\sum_{t=1}^T \eta_t^2 \mathbb{1}\{\eta_t \geq 0\}} \sqrt{\sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}_t)} \right] \quad (\because \text{Cauchy-Schwarz inequality}) \quad (154)$$

$$\leq \mathbb{E} \left[\sqrt{\sum_{t=1}^T \eta_t^2 \mathbb{1}\{\eta_t \geq 0\}} \right] \sqrt{C_1 \gamma_T}, \quad (155)$$

$$\leq \sqrt{\mathbb{E} \left[\sum_{t=1}^T \eta_t^2 \mathbb{1}\{\eta_t \geq 0\} \right]} \sqrt{C_1 \gamma_T}, \quad (\because \text{Jensen's inequality}) \quad (156)$$

where $C_1 := 2/\log(1 + \sigma^{-2})$ and the indicator function $\mathbb{1}\{\eta_t \geq 0\} = 1$ if $\eta_t \geq 0$, and 0 otherwise. For the third inequality, we use $\sum_{t=1}^T \sigma_{t-1}^2(\mathbf{x}_t) \leq C_1 \gamma_T$ (Lemma 5.2 in Srinivas et al., 2010). \square

C.2 Proof of Lemma 4.2

Lemma 4.2. Fix $\beta > 0$ and $U \leq \max_{\mathbf{x} \in \mathcal{X}} \{\mu_{t-1}(\mathbf{x}) + \beta^{1/2} \sigma_{t-1}(\mathbf{x})\}$. Then, the following inequality holds:

$$\frac{U - \mu_{t-1}(\mathbf{x}_t)}{\sigma_{t-1}(\mathbf{x}_t)} \leq \sqrt{\log \left(\frac{\sigma^2 + t - 1}{\sigma^2} \right) + \beta + \sqrt{2\pi\beta}}, \quad (157)$$

where $\mathbf{x}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \text{EI}(\mu_{t-1}(\mathbf{x}), \sigma_{t-1}(\mathbf{x}), U)$.

Proof. Let $\mathbf{z} \in \mathcal{X}$ be the input that satisfies $U \leq \mu_{t-1}(\mathbf{z}) + \beta^{1/2} \sigma_{t-1}(\mathbf{z})$, that is $\frac{U - \mu_{t-1}(\mathbf{z})}{\sigma_{t-1}(\mathbf{z})} \leq \beta^{1/2}$. If $U \leq \mu_{t-1}(\mathbf{x}_t)$, it is obvious due to $\frac{U - \mu_{t-1}(\mathbf{x}_t)}{\sigma_{t-1}(\mathbf{x}_t)} \leq 0$. Thus, assume $U > \mu_{t-1}(\mathbf{x}_t)$ hereafter. Then, from the assumption and the monotonic increasing property of $\tau(c) = c\Phi(c) + \phi(c)$,

$$\text{EI}(\mu_{t-1}(\mathbf{z}), \sigma_{t-1}(\mathbf{z}), U) = \sigma_t(\mathbf{z}) \tau \left(\frac{\mu_{t-1}(\mathbf{z}) - U}{\sigma_{t-1}(\mathbf{z})} \right) \quad (158)$$

$$\geq \sigma_t(\mathbf{z}) \tau(-\beta^{1/2}). \quad (159)$$

In addition, due to the definition of \mathbf{x}_t , we can see that

$$\text{EI}(\mu_{t-1}(\mathbf{z}), \sigma_{t-1}(\mathbf{z}), U) \leq \text{EI}(\mu_{t-1}(\mathbf{x}_t), \sigma_{t-1}(\mathbf{x}_t), U) \quad (160)$$

Therefore, we obtain that

$$\sigma_t(\mathbf{z}) \tau(-\beta^{1/2}) \leq \text{EI}(\mu_{t-1}(\mathbf{x}_t), \sigma_{t-1}(\mathbf{x}_t), U) \quad (161)$$

$$= \sigma_{t-1}(\mathbf{x}_t) \tau \left(\frac{\mu_{t-1}(\mathbf{x}_t) - U}{\sigma_{t-1}(\mathbf{x}_t)} \right). \quad (162)$$

From the assumption $U > \mu_{t-1}(\mathbf{x}_t)$, we can see that $\frac{\mu_{t-1}(\mathbf{x}_t) - U}{\sigma_{t-1}(\mathbf{x}_t)} < 0$. Then, since $\tau(c) \leq \phi(c)$ if $c < 0$, we further obtain

$$\sigma_t(\mathbf{z})\tau(-\beta^{1/2}) \leq \sigma_{t-1}(\mathbf{x}_t)\phi\left(\frac{\mu_{t-1}(\mathbf{x}_t) - U}{\sigma_{t-1}(\mathbf{x}_t)}\right) \quad (163)$$

$$= \frac{\sigma_{t-1}(\mathbf{x}_t)}{\sqrt{2\pi}} \exp\left(-\frac{(\mu_{t-1}(\mathbf{x}_t) - U)^2}{2\sigma_{t-1}^2(\mathbf{x}_t)}\right). \quad (164)$$

Hence, by transforming and applying the logarithm and the square root, we can obtain that

$$|\mu_{t-1}(\mathbf{x}_t) - U| \leq \sqrt{2 \log\left(\frac{\sigma_{t-1}(\mathbf{x}_t)}{\sqrt{2\pi}\sigma_t(\mathbf{z})\tau(-\beta^{1/2})}\right)} \sigma_{t-1}(\mathbf{x}_t). \quad (165)$$

Due to the monotonicity of $\sigma_{t-1}(\mathbf{x}) \geq \sigma_t(\mathbf{x})$, we see that $\sigma_{t-1}(\mathbf{x}_t) \leq 1$. In addition, from Lemma D.1, $\sigma_{t-1}^2(\mathbf{z}) \geq \frac{\sigma^2}{\sigma^2 + t - 1}$. Therefore,

$$|\mu_{t-1}(\mathbf{x}_t) - U| \leq \left(\sqrt{\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right)} - 2 \log(\tau(-\beta^{1/2})) - \log(2\pi)\right) \sigma_{t-1}(\mathbf{x}_t). \quad (166)$$

Then, for $\tau(-\beta^{1/2})$, we can bound as follows:

$$\tau(-\beta^{1/2}) = -\beta^{1/2}\Phi(-\beta^{1/2}) + \phi(-\beta^{1/2}) \quad (167)$$

$$= -\beta^{1/2}\left(1 - \Phi(\beta^{1/2})\right) + \phi(\beta^{1/2}) \quad (168)$$

$$\geq -\left(1 - \exp\left(-\sqrt{\frac{\pi}{2}}\beta^{1/2}\right)\right)\phi(\beta^{1/2}) + \phi(\beta^{1/2}) \quad (169)$$

$$= \exp\left(-\sqrt{\frac{\pi}{2}}\beta^{1/2}\right)\phi(\beta^{1/2}), \quad (170)$$

where we use Lemma 4.3. Thus, we obtain

$$-2 \log(\tau(-\beta^{1/2})) \leq -2\left(-\sqrt{\frac{\pi}{2}}\beta^{1/2} - \log\sqrt{2\pi} - \frac{\beta}{2}\right) \quad (171)$$

$$= \sqrt{2\pi\beta} + \log(2\pi) + \beta. \quad (172)$$

By combining the results, we obtain the desired result. \square

C.3 Proofs for Sec. 4.2.1

First, we show the proof of the following corollary:

Corollary 4.5. *Assume the same condition as in Lemma 4.4. Then, by running GP-EIMS, the following inequality holds with probability at least $1 - \delta$:*

$$\eta_t \leq \sqrt{\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(\delta) + \sqrt{2\pi\beta(\delta)}}, \quad (173)$$

where $\beta(\delta) = 2 \log(|\mathcal{X}|/(2\delta))$. Furthermore, we obtain

$$\mathbb{E}[\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}] \leq \log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + C_2 + \sqrt{2\pi C_2}, \quad (174)$$

where $C_2 = 2 + 2 \log(|\mathcal{X}|/2)$.

Proof. From Lemma 4.4, we see that

$$\Pr\left(g_t^* \leq \mu_{t-1}(\mathbf{z}_t) + \beta^{1/2}(\delta)\sigma_{t-1}(\mathbf{z}_t)\right) \geq 1 - \delta, \quad (175)$$

where $\mathbf{z}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} g_t(\mathbf{x})$. Therefore, as a direct consequence of Lemma 4.2, we obtain that

$$\eta_t \leq \sqrt{\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(\delta) + \sqrt{2\pi\beta(\delta)}}, \quad (176)$$

which holds with probability at least $1 - \delta$. Furthermore, since the right-hand side is positive, the following also holds:

$$\Pr\left(\eta_t^2 \mathbb{1}\{\eta_t \geq 0\} \leq \log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(\delta) + \sqrt{2\pi\beta(\delta)}\right) \geq 1 - \delta \quad (177)$$

$$\Leftrightarrow F\left(\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(\delta) + \sqrt{2\pi\beta(\delta)}\right) \geq 1 - \delta, \quad (178)$$

where $F(\cdot)$ is the cumulative distribution function of $\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}$. By applying F^{-1} , which is the generalized inverse function of F and monotonically increasing, we obtain that

$$\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(\delta) + \sqrt{2\pi\beta(\delta)} \geq F^{-1}(1 - \delta). \quad (179)$$

By substituting $Z \sim \text{Uni}(0, 1)$ to δ and taking the expectation with respect to Z ,

$$\mathbb{E}_Z \left[\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(Z) + \sqrt{2\pi\beta(Z)} \right] \geq \mathbb{E}_Z [F^{-1}(1 - Z)] = \mathbb{E}_Z [F^{-1}(Z)] = \mathbb{E} [\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}]. \quad (180)$$

In the last equality, we use the fact that $\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}$ and $F^{-1}(Z)$ are identically distributed. Furthermore, the left-hand side can be transformed as

$$\mathbb{E}_Z \left[\log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \beta(Z) + \sqrt{2\pi\beta(Z)} \right] = \log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \mathbb{E}_Z[\beta(Z)] + \mathbb{E}_Z[\sqrt{2\pi\beta(Z)}] \quad (181)$$

$$\leq \log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + \mathbb{E}_Z[\beta(Z)] + \sqrt{2\pi\mathbb{E}_Z[\beta(Z)]}, \quad (182)$$

where the last inequality is obtained by Jensen's inequality. In addition, we can derive

$$\mathbb{E}_Z[\beta(Z)] = 2 \log(|\mathcal{X}|/2) + \mathbb{E}_Z[2 \log(1/Z)] = 2 \log(|\mathcal{X}|/2) + 2, \quad (183)$$

since $\log(1/Z)$ follows the exponential distribution. Consequently, we obtain that

$$\mathbb{E} [\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}] \leq \log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + C_2 + \sqrt{2\pi C_2}, \quad (184)$$

where $C_2 = \mathbb{E}_Z[\beta(Z)] = 2 \log(|\mathcal{X}|/2) + 2$. □

Then, we provide the detailed proof of Theorem 4.6:

Theorem 4.6. *Let $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$, and \mathcal{X} be finite. Then, by running GP-EIMS, the BCR can be bounded from above as follows:*

$$\text{BCR}_T \leq \sqrt{C_1 B_T T \gamma_T}, \quad (185)$$

where $B_T = \log\left(\frac{\sigma^2 + T - 1}{\sigma^2}\right) + C_2 + \sqrt{2\pi C_2}$, $C_1 := 2/\log(1 + \sigma^{-2})$, and $C_2 = 2 + 2 \log(|\mathcal{X}|/2)$.

Proof. From Lemma 4.1 and Corollary 4.5, we see that

$$\text{BCR}_T \leq \sqrt{\sum_{t=1}^T \mathbb{E} [\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}]} \sqrt{C_1 \gamma_T} \quad (186)$$

$$\leq \sqrt{C_1 B_T T \gamma_T}. \quad (187)$$

□

C.4 Proofs for Sec. 4.2.2

First we show the proof of Lemma 4.7:

Lemma 4.7. *Suppose Assumption 2.1 holds. Pick $\delta \in (0, 1)$ and $t \geq 1$. Then, the following holds:*

$$\Pr_t \left(\forall \mathbf{x} \in \mathcal{X}, f(\mathbf{x}) \leq \mu_{t-1}([\mathbf{x}]_t) + 2\beta_t^{1/2}(\delta/2)\sigma_{t-1}([\mathbf{x}]_t) \right) \geq 1 - \delta \quad (188)$$

where $\beta_t(\delta) = 2d \log(m_t) - 2 \log(2\delta)$ and $m_t = \max\{2, \lceil bdr \sqrt{(\sigma^2 + t - 1) \log(2ad)/\sigma^2} \rceil\}$.

Proof. Under Assumption 2.1, using the union bound with respect to $j \in [d]$, we see

$$\Pr \left(\max_{j \in [d]} L_j > L(\delta) \right) \leq \delta, \quad (189)$$

where $L(\delta) := b\sqrt{\log(ad/\delta)}$. Therefore, the following inequality holds with probability at least $1 - \delta$:

$$\Pr(\forall \mathbf{x} \in \mathcal{X}, f(\mathbf{x}) - f([\mathbf{x}]_t) \leq L(\delta)\|\mathbf{x} - [\mathbf{x}]_t\|_1) \geq 1 - \delta. \quad (190)$$

Hence, by using the union bound, the following upper bound holds with probability at least $1 - \delta$:

$$f(\mathbf{x}) - \mu_{t-1}([\mathbf{x}]_t) = f(\mathbf{x}) - f([\mathbf{x}]_t) + f([\mathbf{x}]_t) - \mu_{t-1}([\mathbf{x}]_t) \quad (191)$$

$$\leq L_{\delta/2}\|\mathbf{x} - [\mathbf{x}]_t\|_1 + \beta_t^{1/2}(\delta/2)\sigma_{t-1}([\mathbf{x}]_t). \quad (192)$$

In addition, from the construction of \mathcal{X}_t ,

$$\forall \mathbf{x} \in \mathcal{X}, \|\mathbf{x} - [\mathbf{x}]_t\|_1 \leq \frac{dr}{m_t}. \quad (193)$$

Therefore, we see that

$$L_{\delta/2}\|\mathbf{x} - [\mathbf{x}]_t\|_1 \leq \frac{L_{\delta/2}dr}{m_t} \quad (194)$$

$$= \frac{bdr\sqrt{\log(2ad/\delta)}}{m_t} \quad (195)$$

$$\leq \sqrt{\frac{\sigma^2}{\sigma^2 + t - 1}} \sqrt{1 - \frac{\log(\delta)}{\log(2ad)}} \quad (\because m_t \geq bdr\sqrt{(\sigma^2 + t - 1)\log(2ad)/\sigma^2}) \quad (196)$$

$$\leq \sqrt{\frac{\sigma^2}{\sigma^2 + t - 1}} \sqrt{1 - 2\log(\delta)}. \quad (\because \log(2ad) \geq \log 2 > 0.69) \quad (197)$$

On the other hand,

$$\beta_t^{1/2}(\delta/2) = \sqrt{2d \log(m_t) - 2 \log(\delta)} \quad (198)$$

$$\geq \sqrt{2d \log(2) - 2 \log(\delta)}. \quad (\because m_t \geq 2) \quad (199)$$

Moreover, from Lemma D.1, $\sigma_{t-1}^2(x) \geq \sigma^2/(\sigma^2 + t - 1)$ for all $\mathbf{x} \in \mathcal{X}$. Therefore, we obtain that

$$L_{\delta/2}\|\mathbf{x} - [\mathbf{x}]_t\|_1 \leq \beta_t^{1/2}(\delta/2)\sigma_{t-1}([\mathbf{x}]_t). \quad (200)$$

Finally, we derive the desired result:

$$f(\mathbf{x}) - \mu_{t-1}([\mathbf{x}]_t) \leq 2\beta_t^{1/2}(\delta/2)\sigma_{t-1}([\mathbf{x}]_t). \quad (201)$$

□

Next, we show the BCR upper bound for continuous input domains:

Theorem 4.8. *Suppose Assumption 2.1 holds and $f \sim \mathcal{GP}(0, k)$, where the kernel function k is normalized as $k(\mathbf{x}, \mathbf{x}) \leq 1$. Then, by running GP-EIMS, the BCR can be bounded from above as follows:*

$$\text{BCR}_T \leq \sqrt{C_1 B_T T \gamma_T}, \quad (202)$$

where $C_1 := 2/\log(1 + \sigma^{-2})$, $B_T = \log\left(\frac{\sigma^2 + T - 1}{\sigma^2}\right) + C_T + \sqrt{2\pi C_T}$, $C_T = 8(d \log(m_T) + 1)$, and $m_T = \max\{2, \lceil bdr \sqrt{(\sigma^2 + T - 1) \log(2ad)/\sigma^2} \rceil\}$.

Proof. As with the proof of Corollary 4.5, using Lemmas 4.2 and 4.7, we can derive that

$$\mathbb{E}[\eta_t^2 \mathbb{1}\{\eta_t \geq 0\}] \leq \log\left(\frac{\sigma^2 + t - 1}{\sigma^2}\right) + C_t + \sqrt{2\pi C_t}, \quad (203)$$

where $C_t = 8(d \log(m_t) + 1)$. Then, as with the proof of Theorem 4.6, we can obtain that

$$\text{BCR}_T \leq \sqrt{\sum_{t=1}^T \mathbb{E}[\eta_t^2 \mathbb{1}\{\eta_t > 0\}]} \sqrt{C_1 \gamma_T} \quad (204)$$

$$\leq \sqrt{C_1 B_T T \gamma_T}. \quad (205)$$

□

D Auxiliary Lemmas

Lemma D.1 (Lemma 4.2 in (Takeno et al., 2024)). *Let k be a kernel s.t. $k(\mathbf{x}, \mathbf{x}) \leq 1$. Then, the posterior variance is lower bounded as,*

$$\sigma_t^2(\mathbf{x}) \geq \frac{\sigma^2}{\sigma^2 + t}, \quad (206)$$

for all $\mathbf{x} \in \mathcal{X}$ and for all $t \geq 0$.

Lemma D.2. *Let $f \sim \mathcal{GP}(0, k)$, where k is a kernel s.t. $k(\mathbf{x}, \mathbf{x}) \leq 1$. If \mathcal{X} is finite, $\mathbb{E}[\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})] = \mathbb{E}[\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})] \leq \sqrt{2 \log(|\mathcal{X}|/2) + 2}$. If Assumption 2.1 holds, $\mathbb{E}[\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})] = \mathbb{E}[\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})] \leq \sqrt{2d \log(\lceil bdr(\log(ad) + \sqrt{\pi}/2) \rceil) + 2} - 2 \log 2 + 1$.*

Proof. First, since f follows a centered GP, $\mathbb{E}[\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})] = \mathbb{E}[\min_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})]$. For a finite input domain \mathcal{X} , from Lemma 4.2 of (Takeno et al., 2023), we can obtain

$$\mathbb{E}\left[\max_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})\right] \leq \mathbb{E}[\zeta^{1/2}] \leq \sqrt{\mathbb{E}[\zeta]} = \sqrt{2 \log(|\mathcal{X}|/2) + 2}, \quad (207)$$

where $\zeta = 2 \log(|\mathcal{X}|/2) + 2Z$ and $Z \sim \text{Exp}(\lambda = 1)$.

For an infinite input domain, we consider a finite set \mathcal{X}_t with each dimension evenly divided by $\tau_t = \lceil bdr(\log(ad) + \sqrt{\pi}/2) \rceil$, such that $|\mathcal{X}_t| = \tau_t^d$. In addition, we denote $[\mathbf{x}]_t = \arg \min_{\mathbf{x}' \in \mathcal{X}_t} \|\mathbf{x} - \mathbf{x}'\|_1$. Then, by Lemma H.2 of (Takeno et al., 2023), we see that

$$\mathbb{E}[f(\mathbf{x}^*) - f([\mathbf{x}^*]_t)] \leq 1. \quad (208)$$

Furthermore, as with the case of a finite input domain,

$$\mathbb{E}\left[\max_{\mathbf{x} \in \mathcal{X}_t} f(\mathbf{x})\right] \leq \sqrt{2 \log(|\mathcal{X}_t|/2) + 2} = \sqrt{2d \log(\lceil bdr(\log(ad) + \sqrt{\pi}/2) \rceil) + 2} - 2 \log 2, \quad (209)$$

which concludes the proof. □

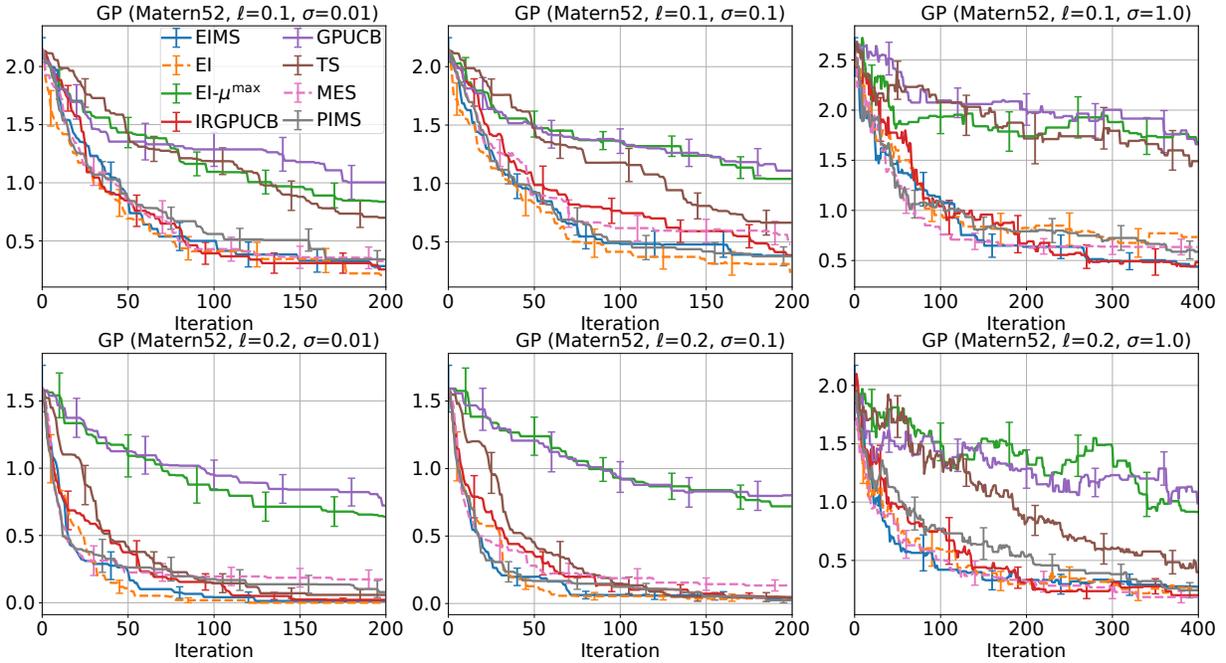


Figure 3: Results of simple regret for synthetic functions generated from GP defined by the Matérn kernel with $\nu = 5/2$. The top and bottom rows show the result on the length scale of the kernel $\ell = 0.1$ and $\ell = 0.2$, respectively. The left, centered, and right columns show the result on the noise standard deviation $\sigma = 0.01, 0.1$, and 1 , respectively.

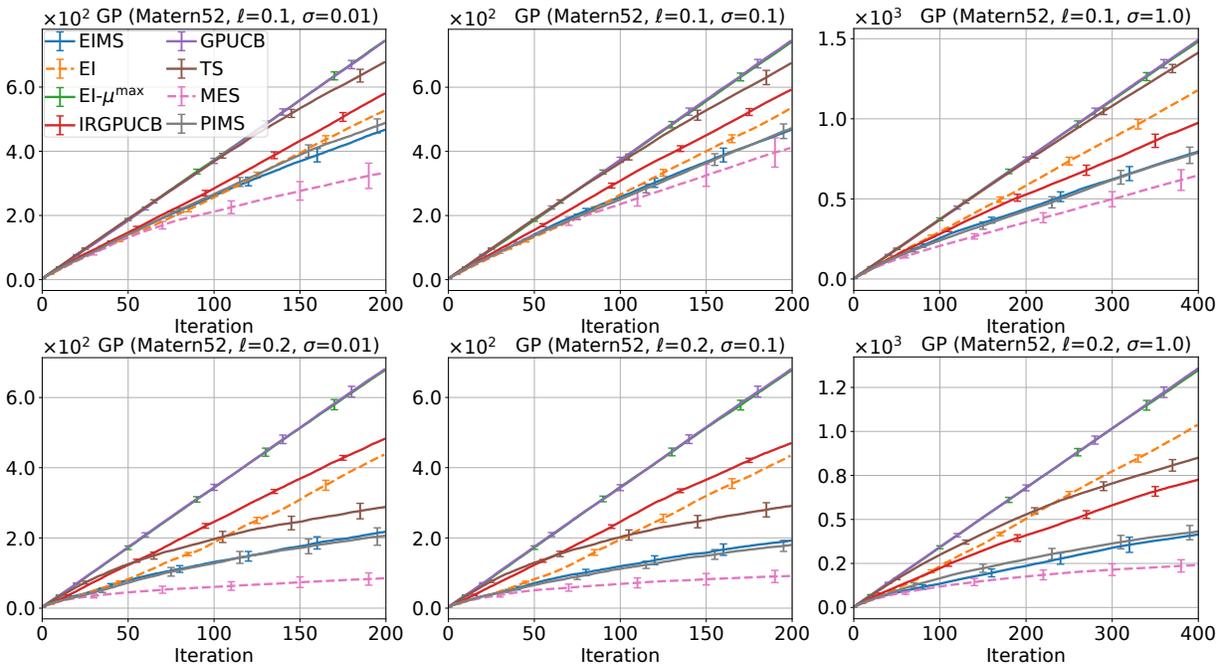


Figure 4: Results of cumulative regret for synthetic functions generated from GP defined by the Matérn kernel with $\nu = 5/2$. The top and bottom rows show the result on the length scale of the kernel $\ell = 0.1$ and $\ell = 0.2$, respectively. The left, centered, and right columns show the result on the noise standard deviation $\sigma = 0.01, 0.1$, and 1 , respectively.

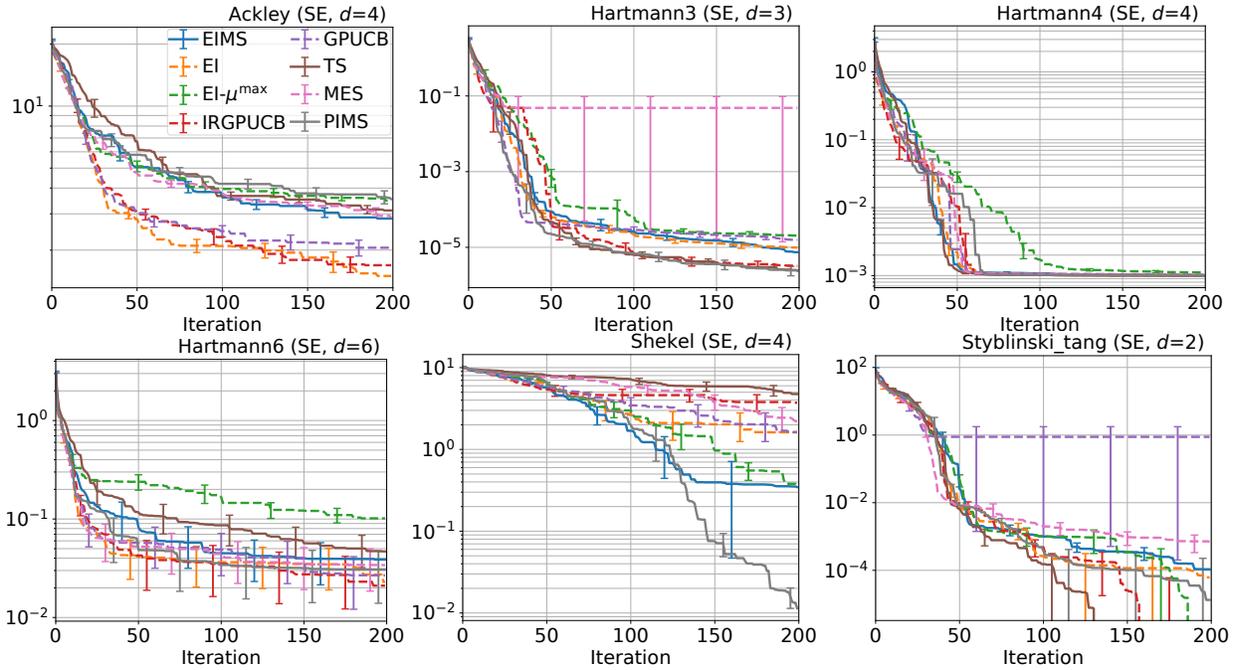


Figure 5: Results of simple regret for the benchmark functions. The BO methods with the theoretical guarantee, that is, GP-EIMS, GP-PIMS, and GP-TS, are indicated by solid lines.

E Additional Experiments

E.1 Synthetic Function Experiments with Matérn Kernel

Figures 3 and 4 show the simple regret and cumulative regret for the Matérn kernel with $\nu = 5/2$. The experimental settings are the same as those in the main paper, except for the kernel function. Also in these experiments, we can observe a similar tendency to the results of the SE kernel function.

E.2 Benchmark Function Experiments

Figure 5 shows the results of the simple regret for several benchmark functions. We employed the SE kernel function with automatic relevance determination (Rasmussen & Williams, 2005). We performed hyperparameter selection for the GP model every five iterations. In these experiments, we did not add actual noise. Since $\arg \max_{\mathbf{x} \in \mathcal{X}} \mu_t(\mathbf{x})$ is unstable due to the hyperparameter selection, we define the simple regret using $\hat{\mathbf{x}}_t = \arg \max_{\mathbf{x} \in \mathcal{X}} \{\mu_t(\mathbf{x}) - 2\sigma_t(\mathbf{x})\}$. We employed the heuristic hyperparameters for the AFs of GP-UCB, IRGP-UCB, and GP-EI- μ^{\max} as $\beta_t = 0.2d \log(2t)$ as with (Kandasamy et al., 2016), $\nu_t = 0.2d \log(2t)$, and $\zeta_t = \max\{0.2d \log(2t) - 2, 0\} + Z_t$, where $Z_t \sim \text{Exp}(\lambda = 1/2)$. Therefore, GP-UCB, IRGP-UCB, and GP-EI- μ^{\max} are no longer theoretically guaranteed and indicated by dashed lines. Other settings are the same as those of the synthetic function experiments.

Except for the Ackley and Shekel functions, almost all BO methods perform well. For the Ackley function, GP-UCB, IRGP-UCB, and GP-EI show superior performance. This is because these BO methods are more exploitative than other BO methods in this experiment, and exploitative AFs work well for the Ackley function. However, for the Shekel function, more explorative GP-PIMS and GP-EIMS show superior performance. Therefore, we believe that these differences in performance come from the exploitation-exploration tradeoff. On the other hand, regarding the BO methods with the theoretical guarantee without heuristic tunings, we consider that GP-PIMS and GP-EIMS are superior or comparable to GP-TS.