
Regret Bounds for Expected Improvement Algorithms in Gaussian Process Bandit Optimization

Hung Tran-The

Sunil Gupta

Santu Rana

Svetha Venkatesh

Applied Artificial Intelligence Institute, Deakin University, Geelong, Australia

Abstract

The expected improvement (EI) algorithm is one of the most popular strategies for optimization under uncertainty due to its simplicity and efficiency. Despite its popularity, the theoretical aspects of this algorithm have not been properly analyzed. In particular, whether in the noisy setting, the EI strategy with a standard incumbent converges is still an open question of the Gaussian process bandit optimization problem. We aim to answer this question by proposing a variant of EI with a standard incumbent defined via the GP predictive mean. We prove that our algorithm converges, and achieves a cumulative regret bound of $\mathcal{O}(\gamma_T \sqrt{T})$, where γ_T is the maximum information gain between T observations and the Gaussian process model. Based on this variant of EI, we further propose an algorithm called Improved GP-EI that converges faster than previous counterparts. In particular, our proposed variants of EI do not require the knowledge of the RKHS norm and the noise’s sub-Gaussianity parameter as in previous works. Empirical validation in our paper demonstrates the effectiveness of our algorithms compared to several baselines.

1 Introduction

The problem of sequentially optimizing a black-box function based on bandit feedback has recently attracted a great deal of attention and finds application in robotics (Lizotte et al., 2007; Martinez-cantin et al., 2007), environmental monitoring (Marchant and Ramos, 2012), automatic machine learning (Bergstra

et al., 2011; Snoek et al., 2012; Hoffman et al., 2014) and reinforcement learning (Wilson et al., 2014; Balakrishnan et al., 2020). Under this model, the goal is to design a sequential algorithm in a search space \mathcal{X} , i.e., a sequence x_1, x_2, \dots, x_T such that at iteration T , the algorithm returns a state with the highest possible value. For this problem, a widely used performance measure is the *cumulative regret* R_T , which is given by $R_T = \sum_{t=1}^T \sup_{x \in \mathcal{X}} f(x) - f(x_t^+)$, where x_t^+ is the point reported by the algorithm at iteration t .

In order to make this problem tractable, one must make smoothness assumptions on the function. A versatile means for doing this is to model the function as a Gaussian process (GP) which captures the smoothness properties through a suitably chosen kernel. For GP-based algorithms, there are two settings: *Bayesian* (De Freitas et al., 2012; Tran-The et al., 2021) and *non-Bayesian* (Scarlett et al., 2017). In Bayesian setting, the function is assumed to be sampled from a GP while in the non-Bayesian setting, the function is treated as fixed and unknown, and assumed to lie in a reproducing kernel Hilbert space (RKHS). Under these assumptions, the optimization problem is usually called the *Gaussian process bandit optimization* whereas the optimization in Bayesian setting is referred to as *Bayesian optimization*. In this paper, we focus on the non-Bayesian setting, i.e. Gaussian process bandit optimization.

The Expected Improvement (EI) (Moćkus, 1975) is one of the most widely used strategy to optimize black-box functions due to its simplicity and ability to handle uncertainty (e.g., works of Osborne (2010); Wilson et al. (2014); Qin et al. (2017); Malkomes and Garnett (2018); Nguyen and Osborne (2020)). Unlike other popular strategies, upper confidence bound (UCB) and Thompson sampling (TS), EI is a greedy improvement-based strategy, which samples the next point offering the greater expected improvement over the current incumbent. Formally, $\alpha^{EI}(x) = \mathbb{E}[\max\{0, f(x) - \xi\}]$, where ξ is the incumbent to be defined. In the noiseless setting, ξ is defined as the current best observation so far, f_t^+ . Given \mathcal{D}_t which is the set of sampled points up to iteration t , f_t^+ at iteration t is computed as

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

$f_t^+ = \max_{1 \leq i \leq t} f(x_i)$. However, in the noisy setting, such a choice is not clear due to the noise. As an alternative, ξ is typically defined as either the current best value of the GP predictive mean, formally $\xi = \mu_t^+$ which is computed as $\mu_t^+ = \max_{1 \leq i \leq t} \mu_t(x_i)$, or the current best observation value (containing the noise), formally $\xi = y_t^+$. Although these incumbents can be easily computed, they make the theoretical analysis harder.

A key challenge of analyzing EI-based algorithms comes from its improvement function involving nonlinear, non-convex term unlike UCB and TS. This causes the difficulty of the analysis of EI. In the noisy setting, another challenge comes from the fact that the incumbents such as $\xi = \mu_t^+$ or $\xi = y_t^+$ do not have monotonicity property like the function f_t^+ ($f_{t+1}^+ \geq f_t^+$). This is one of crucial properties to reach the convergence in the noise-free setting (Bull, 2011). These reasons explain why convergence properties of GP-EI are not been well studied especially in the noisy setting.

Wang and de Freitas (2014) studied GP-EI when the lower and upper bounds of hyper-parameters of GP are known. However, to guarantee the convergence, their non-peer reviewed work uses an alternative choice of the incumbent as the maximum of the GP predictive mean $\xi = \max_{x \in \mathcal{X}} \mu_t(x)$. As a result, their GP-EI algorithm requires an additional optimization step to approximate $\mu_t(x)$ at each iteration which is computationally expensive compared to the use of standard incumbents μ_t^+ or y_t^+ especially when the search space is large or unbounded (Tran-The et al., 2020).

Nguyen et al. (2017) proposed a “weak” version of GP-EI which uses y^+ as the incumbent. This incumbent is easily computed, however their version of GP-EI needs to use an assumption that values of the variance function at all sampled points are not allowed to exceed a lower bound κ . Due to this, their GP-EI regret upper bound depends on κ and this bound quickly explodes as $\kappa \rightarrow \infty$. As a result, their analysis does not solve the traditional EI algorithm as we consider in this paper. Therefore, *the natural question of whether GP-EI with a standard setting of incumbent (e.g., $\xi = \mu^+$ or $\xi = y^+$) converges, and if true then what convergence rate GP-EI can reach are open problems?* In this paper, we provide an affirmative answer to these questions. Our main contributions are as follows:

- We propose a variant of GP-EI for Gaussian process bandit optimization. This algorithm uses a standard incumbent $\xi = \max\{\mu_{t-1}(x_i)\}_{x_i \in \mathcal{D}_{t-1}}$ at iteration t , where \mathcal{D}_{t-1} is the set of sampled points up to iteration t . Our algorithm enjoys a cumulative regret bound of $\mathcal{O}(\gamma_T \sqrt{T})$, where γ_T is the maximum information gain between T ob-

servations and the GP model. To our knowledge, this is the first GP-EI algorithm using a standard incumbent with theoretical guarantees.

- Based on our above algorithm, we propose an efficient variant of GP-EI called Improved-GP-EI of which f lies in RKHS equipped by a Matérn- ν kernel. Improved-GP-EI can achieve regret $\tilde{\mathcal{O}}(T^{\frac{d(2d+3)+2\nu}{d(2d+4)+4\nu}})$ for every $\nu > 1$ and $d \geq 1$. By using a searching partitioning strategy, Improved-GP-EI avoids the quick growth of the global scale of variance functions. In particular, it does not require the knowledge of the RKHS norm and the measurement noise’s sub-Gaussianity parameter. These parameters are required by most previous algorithms for theoretical guarantees, but are usually unknown in real applications.
- We demonstrate the practical effectiveness of Improved-GP-EI against GP-EI and π -GP-UCB on various synthetic functions.

2 Preliminaries

We consider a global optimisation problem whose goal is to maximise $f(x)$ subject to $x \in \mathcal{X} \subset \mathbb{R}^d$, where d is the number of dimensions and f is an expensive blackbox function that can only be evaluated point-wise. The performance of a global optimisation algorithm is typically evaluated using the cumulative regret which we have defined in section Introduction.

2.1 Regularity Assumptions

We assume that f lives in a RKHS of functions $\mathcal{X} \rightarrow \mathbb{R}$ with positive semi-definite kernel function k . This H_k space is defined as the Hilbert space of functions on \mathcal{X} equipped with an inner product $\langle \cdot, \cdot \rangle_k$ obeying the reproducing property: $f(x) = \langle f, k(x, \cdot) \rangle_k$ for all $f \in H_k(\mathcal{X})$. The RKHS norm $\|f\|_k = \sqrt{\langle f, f \rangle_k}$ is a measure of smoothness of f , with respect to the kernel function k , and satisfies: $f \in H_k(\mathcal{X})$ if and only if $\|f\|_k < \infty$. We assume that the RKHS norm of the unknown target function is bounded by $\|f\|_k \leq B$. Two common kernels that satisfy bounded variance property are Squared Exponential (SE) and Matérn, defined as

$$k_{\text{SE}}(x, x') = \exp\left(-\frac{\|x - x'\|_2^2}{2l^2}\right), \tag{1}$$

$$k_{\text{Matérn}}(x, x') = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\|x - x'\|_2}{l}\right)^\nu \mathcal{B}_\nu\left(\frac{\|x - x'\|_2}{l}\right), \tag{2}$$

where Γ denotes the Gamma function, \mathcal{B}_ν denotes the modified Bessel function of the second kind, ν is a

parameter controlling the smoothness of the function and l is the lengthscale of the kernel. Important special cases of ν include $\nu = \frac{1}{2}$ that corresponds to the exponential kernel and $\nu \rightarrow \infty$ that corresponds to the square exponential (SE) kernel. The Matérn kernel is of particular practical significance, since it offers a more suitable set of assumptions for the modeling and optimisation of physical quantities (Stein, 1999).

2.2 Gaussian process bandit optimization

Gaussian process bandit optimization proceeds sequentially in an iterative fashion. At each iteration, a surrogate model is used to probabilistically model $f(\mathbf{x})$. Gaussian process (GP) (Rasmussen and Williams, 2005) is a popular choice for the surrogate model as it offers a prior over a large class of functions and its posterior and predictive distributions are tractable. Formally, we assume $f(x) \sim \mathcal{GP}(m(x), \omega^2 k(x, x'))$ a prior distribution where $m(x)$ is the mean function and $\omega^2 k(x, x')$ is the covariance function in which $k(x, x')$ is a kernel function associated with the RKHS H_k in which f is assumed to have norm at most B , and $\omega > 0$ is a parameter to capture the global scale of variation of function f . Without loss of generality, we assume that $m(x) = 0$. Given a set of observations $\mathcal{D}_{1:t} = \{x_i, y_i\}_{i=1}^t$, the predictive distribution can be derived as $P(f_{t+1} | \mathcal{D}_{1:t}, x) = \mathcal{N}(\mu_{t+1}(x), \omega^2 \sigma_{t+1}^2(x))$, where

$$\mu_{t+1}(x) = k_t(x)^T [K_t + \lambda I]^{-1} y_{1:t},$$

$$\sigma_{t+1}^2(x) = k(x, x) - k_t(x)^T (K_t + \lambda I)^{-1} k_t(x),$$

where we define $k_t(x) = [k(x, x_1), \dots, k(x, x_t)]^T$, $K_t = [k(x_i, x_j)]_{1 \leq i, j \leq t}$, $y_{1:t} = [y_1, \dots, y_t]$ and λ as variance of the measurement noise.

We assume that kernel function k is fixed and known and, without loss of generality, the variance of k is bounded as $k(x, x) \leq 1$. These assumptions are similar as in (Srinivas et al., 2012; Chowdhury et al., 2017; Janz et al., 2020). However, unlike these works, we do not require the knowledge of the sub-Gaussianity parameter R and upper bound B on the RKHS norm of f .

In our setting, we note that the parameters ω and λ are possibly time-dependent. They can be set specific to an algorithm as in many previous works (e.g., (Bull, 2011), (Agrawal and Goyal, 2013), (Wang and de Freitas, 2014), (Chowdhury et al., 2017)).

2.3 Expected Improvement

An acquisition function is used to suggest the point x_{t+1} where the function should be next evaluated. The acquisition step uses the predictive mean and the predictive variance from the surrogate model to balance

Algorithm 1 GP-EI Algorithm

Input: Prior $\mathcal{GP}(0, k)$

- 1: **for** $t = 1$ to T **do**
- 2: Set $\omega_t = \sqrt{\gamma_{t-1} + 1 + \ln(\frac{1}{\delta})}$
- 3: Choose the next sampling point: $x_t = \operatorname{argmax}_{x \in \mathcal{X}} \rho(\mu_{t-1}(x) - \mu_t^+, \omega_t \sigma_{t-1}(x))$
- 4: Observe $y_t = f(x_t) + \epsilon_t$
- 5: Update Gaussian process
- 6: **end for**

Output: Report points $x_t^+ = \operatorname{argmax}_{1 \leq i \leq t} \mu_{t-1}(x_i)$ for all $1 \leq t \leq T$

the exploration of the search space and exploitation of current promising region. Some examples of acquisition functions include GP-EI (Bull, 2011), GP-UCB (Srinivas et al., 2012), GP-TS (Chowdhury et al., 2017), and entropy based methods e.g., PES (Hernández-Lobato et al., 2014).

In the noisy case, the function is evaluated as $y_t = f(x_t) + \epsilon_t$, which is a noisy version of the function value at x_t . We assume that the noise sequence $\{\epsilon_t\}_{t=1}^\infty$ is conditionally R -sub-Gaussian for a fixed constant $R \geq 0$, i.e., $\forall t \geq 0, \forall \lambda \in \mathbb{R}, \mathbb{E}[e^{\lambda \epsilon_t} | \mathcal{F}_{t-1}] \leq \exp(\frac{\lambda^2 R^2}{2})$, where \mathcal{F}_{t-1} is the σ -algebra generated by the random variables. This is a mild assumption on the noise and is standard in the BO literature (Chowdhury et al., 2017) and also in bandit literature (Abbasi-yadkori et al., 2011).

We let $\mathcal{D}_t = \{x_1, \dots, x_t\}$ denote the set of chosen points to be evaluated up to iteration t . The noise in the evaluation of the incumbent causes it to be brittle. A standard choice of the incumbent in this setting is the best value of the GP mean function so far $\mu_t^+ = \max\{\mu_{t-1}(x_i)\}_{x_i \in \mathcal{D}_t}$. We note that the maximization is only over the observed points, not the complete input space \mathcal{X} . For this choice, EI is written in closed form as:

$$\alpha_t^{EI}(x) = \mathbb{E}[\max\{0, f(x) - \mu_t^+\} | \mathcal{D}_t] \quad (3)$$

$$= \rho(\mu_{t-1}(x) - \mu_t^+, \omega \sigma_t(x)), \quad (4)$$

where $\rho(u, v)$ with two arguments u and v is defined as

$$\rho(u, v) = \begin{cases} u\Phi(\frac{u}{v}) + v\phi(\frac{u}{v}), & \text{if } v > 0, \\ \max\{0, u\}, & \text{if } v = 0, \end{cases} \quad (5)$$

and Φ and ϕ are the standard normal distribution and density functions respectively.

3 Gaussian Process Expected Improvement (GP-EI) Algorithm

Our GP-EI algorithm is represented in Algorithm 1. 1. The time-varying scale parameter $\omega_t =$

$\sqrt{\gamma_{t-1} + 1 + \ln(1/\delta)}$ is used to control the exploration of the algorithm for guaranteeing the convergence. Here δ is a free parameter in $(0, 1)$ and γ_t is the maximum information gain at time t which is defined as $\gamma_t = \max_{A \in \mathcal{X}: |A|=t} I(y_A; f_A)$, where $I(y_A; f_A)$ denotes the mutual information between $f_A = [f(x)]_{x \in A}$ and $y_A = f_A + \epsilon_A$ and $\epsilon_A \in \mathcal{N}(0, \lambda \omega_t^2 I)$. The algorithm choose a point $x_t = \operatorname{argmax}_{x \in \mathcal{X}} \alpha_t^{EI}(x)$ which is computed by Eq (5) to sample. After T iterations, the algorithm returns points $x_t^+ = \operatorname{argmax}_{1 \leq i \leq t} \mu_{t-1}(x_i)$ for every $1 \leq t \leq T$.

We note that the time-varying scale of variation of function f has been usually utilized by many previous works e.g. (Chowdhury et al., 2017) which analyzes the GP-TS algorithm, (Agrawal and Goyal, 2013) which analyzes a Thompson Sampling algorithm but for the contextual bandit problem, and (Wang and de Freitas, 2014) which analyzes the EI algorithm. For example, in the setting of (Chowdhury et al., 2017), they used a Gaussian process with mean 0 and variance $v_t^2 k(\cdot, \cdot)$, where v_t is a global scale parameter of the variance which is allowed to vary with time. For their GP-Thompson sampling, the time-varying scale parameter is set as $v_t = B + R\sqrt{2(\gamma_{t-1} + 1 + \ln(2/\delta))}$ (See section 3.2 therein).

Comparison with related algorithms. Most of previous algorithms for Gaussian process bandit optimization such as GP-UCB Srinivas et al. (2012), Improved-GP-UCB (Chowdhury et al., 2017), π -GP-UCB (Janz et al., 2020), and GP-TS (Chowdhury et al., 2017) require to know exactly the sub-Gaussianity parameter R and upper bound B on the RKHS norm of f so that theoretical convergence guarantees hold. However, these parameters are often unknown in real applications. To overcome this issue, as an example, Berkenkamp et al. (2019) proposed to learn unknown B by starting from an initial guess B_0 and then scale up the norm bound over time. As a result, B is replaced by a time-varying function $B_0 b(t) g(t)^d$ in their Theorem 1, where functions $b(t)$ and $g(t)$ are designed heuristically. As a result, this causes an additional $\mathcal{O}(b(t)g(t)^{3d/2})$ factor in the regret. Unlike these algorithms, our GP-EI algorithm can avoid the need to specify or learn such parameters. This is because our algorithm uses $\omega_t = \sqrt{\gamma_t + 1 + \ln(1/\delta)}$ which is independent of B and R .

3.1 Theoretical Result

Importantly, we achieve a cumulative regret bound for the proposed GP-EI algorithm, denoted by $R_T = \sum_{t=1}^T (f(x^*) - f(x_t^+))$ as follows:

Theorem 1. *Pick $\delta \in (0, 1)$. Then with probability*

at least $1 - \delta$, the cumulative regret of Algorithm 1 is bounded as:

$$R_T = \mathcal{O}(\gamma_T \sqrt{T}).$$

The complete form of $R_T = \mathcal{O}(\beta_T \sqrt{T \gamma_T})$, where $\beta_T = B + R\sqrt{2(\gamma_{t-1} + 1 + \ln(1/\delta))}$ is provided in Supplementary Material. Here we remove the influence of constants B, R for simplicity. The regret bound of our GP-EI algorithm is same as that of Improved-GP-UCB (Chowdhury et al., 2017) but improve GP-TS (Chowdhury et al., 2017) by a factor $\sqrt{d \ln(dT)}$. For SE kernels, R_T is sublinear on T . However, for Matérn kernels, this proposed algorithm still has some limitations. First, its regret bound is not always sublinear in T . Vakili et al. (2021) currently provides a new bound for γ_T as $\gamma_T = \tilde{\mathcal{O}}(T^{\frac{d}{2\nu+d}})$. By this, $2\nu > d$ is required so that our proposed GP-EI algorithm obtains a sublinear regret. Second, γ_t of the proposed GP-EI algorithm grows quickly with t . In practice, it can cause unnecessary explorations. These motivate us to propose a new variant of the GP-EI algorithm in the next section.

3.2 The Improved-GP-EI Algorithm

In this section, we propose a new variant of GP-EI, called Improved-GP-EI, inspired from the π -GP-UCB algorithm (Janz et al., 2020). Improved-GP-EI uses a global scale $\omega_T = \sqrt{\ln(T) \ln \ln(T)}$ growing only polylogarithmically with T . Here we assume that T is known as in (Janz et al., 2020) and we use the same ω_T at all iterations from 1 to T . Improved-GP-EI is an adaptation of π -GP-UCB algorithm (Janz et al., 2020) to GP-EI. The key difference is that (1) our Improved GP-EI uses EI acquisition function instead of UCB, and (2) Improved GP-EI use a new time-varying scale (in T) parameter ω_T instead of a constant like (Janz et al., 2020). Now we start to describe the underlying idea of Improved GP-EI.

Improved-GP-EI algorithm At each iteration t , the algorithm constructs a cover (a set of hypercubes) of domain \mathcal{X} , \mathcal{A}_t , and selects a point x_t to be evaluated in the next iteration by taking a maximizer of the GP-EI constructed independently on each cover element. The cover \mathcal{A}_t is constructed by induction starting from the initial cover \mathcal{A}_1 . Similar to π -GP-UCB, we set $b = \frac{d+1}{d+2\nu}$ and $q = \frac{d(d+1)}{d(d+2)+2\nu}$. For a hypercube $A \in \mathcal{X}$, we will use ρ_A to denote its diameter.

Let \mathcal{A}_1 be any set of closed hypercubes of cardinality at most $\mathcal{O}(T^q)$ overlapping at edges only and covering the domain \mathcal{X} . At each iteration t , we build a GP, select the next point x_t and then construct a new cover \mathcal{A}_{t+1} as below:

- **Update GP:** Fit an independent GP on each cover element $A \in \mathcal{A}_t$, *using only the data within* A . We define the subset of \mathcal{D}_t in A as $\mathcal{D}_t^A = \{x_i^A \in \mathcal{D}_t | x_i^A \in A\}$. For \mathcal{D}_t^A of cardinality N , we define the kernel k as $k_t^A = [k(x_1^A, x), \dots, k(x_N^A, x)]^T$ and the kernel matrix as $K_t^A = [k(x, x')]_{x, x' \in \mathcal{D}_t^A}$. We use $y_{1:t}^A$ to denote the observations corresponding to points in \mathcal{D}_t^A . For a regularisation parameter $\lambda > 0$, we define the Gaussian process on $A \in \mathcal{X}$ by mean,

$$\mu_t^A = k_t^A(x)(K_t^A + \lambda I)^{-1}y_{1:t}^A,$$

and predictive standard deviation,

$$\sigma_t^A(x) = \sqrt{k(x, x) - k_t^A(x)^T(K_t^A + \lambda I)^{-1}k_t^A(x)}.$$

- **Next point selection:** We select the next point to evaluate as $x_t = \operatorname{argmax}_{A \in \mathcal{A}_t: x \in A} \alpha_t^{EI} = \operatorname{argmax}_{A \in \mathcal{A}_t: x \in A} \rho(\mu_{t-1}(x) - \mu_{t-1}^{A+}, \omega_T \sigma_{t-1}^A(x))$, where μ_{t-1}^{A+} is defined as $\mu_{t-1}^{A+} = \max_{x_i \in A \cap \mathcal{D}_{t-1}} \mu_{t-1}^A(x_i)$, and function ρ is defined by Eq(5).
- **Build \mathcal{A}_{t+1} :** Split any element $A \in \mathcal{A}_t$ for which $\rho_A^{-1/b} < |\mathcal{D}_{t+1}^A| + 1$ along the middle of each side, resulting in 2^d new hypercubes. Let \mathcal{A}_{t+1} be the set of the newly created hypercubes and the elements of \mathcal{A}_t that were not split. (See (Janz et al., 2020) for details.)

We now present the regret bound for our Improved GP-EI algorithm.

Theorem 2. *Pick $\delta \in (0, 1)$. Let $H_k(\mathcal{X})$ be the RKHS of a Matérn kernel k with parameter $\nu > 1$. Then with probability at least $1 - \delta$, the cumulative regret of Improved GP-EI has the following rate:*

$$R_T = \tilde{O}\left(T^{\frac{d(2d+3)+2\nu}{d(2d+4)+4\nu}}\right).$$

The regret bound of Improved GP-EI is sublinear in T for every $\nu > 1$, thus improves over that of our above proposed GP-EI algorithm which requires $2\nu \geq d$. Furthermore, we use $\omega_T = \sqrt{\ln(T)\ln\ln(T)}$ which grows only poly-logarithmically with T . This property is useful in practice as it avoids unnecessary explorations. Finally, we note that our Improved GP-EI achieves the same regret rate as π -GP-UCB (Janz et al., 2020) on the regret, however it does not require to know parameters B, R like the work of Janz et al. (2020).

3.3 Proof Sketch for Theorem 1

In this section, we provide the proof sketch for Theorem 1. A complete proof of Theorem 1 is provided in Ap-

pendix B. To bound the cumulative regret R_T , we proceed to bound instantaneous regrets $r_t = f(x^*) - f(x_t^+)$. The proof involves two steps as follows.

Upper bounding the instantaneous regret $r_t = f(x^*) - f(x_t^+)$: We break down r_t into two terms as follows:

$$r_t = f(x^*) - f(x_t^+) = \underbrace{f(x^*) - \mu_t^+}_{\text{Term 1}} + \underbrace{\mu_t^+ - f(x_t^+)}_{\text{Term 2}}$$

Set $I_t = \max\{0, f(x_{t+1}) - \mu_t^+\}$. We upper bound Term 1 through the following lemma:

Lemma 1. *Pick $\delta \in (0, 1)$. Then with probability at least $1 - \delta$ we have*

$$f(x^*) - \mu_t^+ \leq \frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)}(I_t + (\beta_t + \omega_t)\sigma_t(x_{t+1})),$$

where given any $z \in \mathbb{R}$, the function $\tau(z)$ is defined as $\tau(z) = z\Phi(z) + \phi(z)$, where Φ and ϕ are the standard normal distribution and density functions respectively.

We upper bound Term 2 through the following lemma:

Lemma 2. *Pick a $\delta \in (0, 1)$. Then with probability $1 - \delta$ we have*

$$\mu_t^+ - f(x_t^+) \leq \frac{\beta_t}{\omega_t}(\sqrt{2\pi}(\beta_t + \omega_t)\sigma_t(x_{t+1}) + \sqrt{2\pi}I_t).$$

By using $\omega_t = \sqrt{\gamma_{t-1} + 1 + \ln(\frac{1}{\delta})}$, we obtain $\frac{\beta_t}{\omega_t} \leq B + \sqrt{2}$ and there exists a constant $C > 0$ such that $\frac{\tau(\frac{\beta_t}{\omega_t})}{\tau(-\frac{\beta_t}{\omega_t})} \leq C$ for every t . Combining these results, we obtain an upper bound for the regret r_t as $r_t \leq (B + C + 2)(I_t + 2\beta_t\sigma_t(x_{t+1}))$, where $C > 0$ is constant.

To achieve Lemma 1, we adapt several results of (Bull, 2011) in noise-free setting to our the noisy setting, and to achieve Lemma 2, we exploit additionally properties of the function $\tau(z)$ and the points $x_t^+ = \operatorname{argmax}_{1 \leq i \leq t} \mu_{t-1}(x_i)$. We note that in the noise-free setting, we can use $\xi = f^+$ as the incumbent in the form of the expected improvement $\alpha^{EI}(x) = \mathbb{E}[\max\{0, f(x) - \xi\}]$, where f_t^+ is the current best observed function value so far. In the noisy setting, the function values can not be observed due to noises. Using $\xi = \mu_t^+$ as a replacement allows to compute easily the incumbent but also causes the difficulty in the theoretical analysis. While Bull (2011) leverages the *monotonicity* of $f_t^+ = \max_{1 \leq i \leq t} f(x_i)$ to derive directly an upper bound for the regret r_t , this is very challenging in our setting because the values μ_i^+ with $1 \leq i \leq t$ have no monotonicity property. To overcome this, we propose another manner to upper bound r_T by

seeking to upper bound the sum of $\sum_{t=1}^T r_t$. We obtain $\sum_{t=1}^T r_t \leq (C + B + 2)(\sum_{t=1}^T I_t + \sum_{t=0}^{T-1} \beta_t \sigma_t(x_{t+1}))$.

While upper bounding $\sum_{t=0}^{T-1} \sigma_t(x_{t+1})$ can be achieved via the maximum information gain like previous works in the noisy setting (Srinivas et al., 2012; Chowdhury et al., 2017), upper bounding the sum $\sum_{t=0}^{T-1} \max\{0, f(x_{t+1}) - \mu_t^+\}$ is the *key challenge* in our regret analysis. We overcome this difficulty by exploiting the *monotonicity of variance functions* (Vivarelli, 1998) which shows that $\sigma_t(x) \geq \sigma_{t'}(x)$ if $t \leq t'$. To our knowledge, we exploit for the first time this property for Gaussian process bandit optimization problem.

We achieve an upper bound for $\sum_{t=0}^{T-1} \max\{0, f(x_{t+1}) - \mu_t^+\}$ as in the following important lemma.

Lemma 3. *Pick a $\delta \in (0, 1)$. Then with probability at least $1 - \delta$ we have that*

$$\sum_{t=0}^{T-1} \max\{0, f(x_{t+1}) - \mu_t^+\} \leq \mathcal{O}(\beta_T \sqrt{T\gamma_T}),$$

where $\beta_T = B + R\sqrt{2(\gamma_{T-1} + 1 + \ln(1/\delta))}$.

To prove this lemma, we need two auxiliary lemmas from the literature.

Lemma 4 (Theorem 2 of Chowdhury et al. (2017)). *Pick $\delta \in (0, 1)$. We define $\beta_t = B + R\sqrt{2(\gamma_{t-1} + 1 + \ln(1/\delta))}$ for every $1 \leq t \leq T$. Then $\mathbb{P}(\forall 1 \leq t \leq T, \forall x \in \mathcal{X}, |f(x) - \mu_{t-1}(x)| \leq \beta_t \sigma_{t-1}(x)) \geq 1 - \delta$.*

Lemma 5 (Lemma 5 of De Freitas et al. (2012)). *When $f \in H_k(\mathcal{X})$, then for every $x, y \in \mathcal{X}$, we have $|f(x) - f(y)| \leq BL\|x - y\|_1$, where L is the Lipschitz constant in $H_k(\mathcal{X})$.*

Proof of Lemma 3 Set $S_T = \sum_{t=0}^{T-1} I_t$. There are three cases to be considered:

Case 1 $S_T = 0$. This happens when for every t : $f(x_{t+1}) - \mu_t^+ \leq 0$.

Case 2 There exists a unique index $1 \leq t' \leq T$ such that $f(x_{t'+1}) - \mu_{t'}^+ > 0$. It follows that $S_T = f(x_{t'+1}) - \mu_{t'}^+$. In this case, we have that

$$\begin{aligned} S_T &= f(x_{t'+1}) - \mu_{t'}^+ \\ &\leq f(x_{t'+1}) - (f(x') - \beta_{t'+1} \sigma_{t'}(x')) \\ &\leq f(x_{t'+1}) - f(x') + \beta_{t'+1} \sigma_{t'}(x') \\ &\leq BL\|x_{t'+1} - x'\|_1 + \beta_{t'+1} \\ &= \mathcal{O}(\beta_T), \end{aligned}$$

where in the last inequality, we use Lemma 5, the inequality $\beta_{t'+1} \leq \beta_T$, and the fact that $\sigma_{t'}^2(x) \leq 1$. Finally, because the domain \mathcal{X} is bounded, $\|x_{t'+1} - x'\|_1$ is bounded.

Case 3 There are $0 \leq t_1 < t_2, \dots, < t_l \leq T - 1$ where $l \geq 2$ such that $f(x_{t_i+1}) \geq \mu_{t_i}^+$. Thus, we have

$$\begin{aligned} \sum_{t=0}^{T-1} I_t &= \sum_{t=1}^T \max\{0, f(x_{t+1}) - \mu_t^+\} \\ &= \sum_{i=1}^l (f(x_{t_i+1}) - \mu_{t_i}^+) \\ &\leq \sum_{i=1}^l (\beta_{t_i+1} \sigma_{t_i}(x_{t_i+1}) + \mu_{t_i}(x_{t_i+1}) - \mu_{t_i}^+) \\ &\leq \underbrace{\sum_{i=1}^l \beta_{t_i+1} \sigma_{t_i}(x_{t_i+1})}_{\text{Term 5}} + \underbrace{\sum_{i=1}^l (\mu_{t_i}(x_{t_i+1}) - \mu_{t_i}^+)}_{\text{Term 6}} \end{aligned}$$

Bound Term 5

$$\begin{aligned} \sum_{i=1}^l \beta_{t_i+1} \sigma_{t_i}(x_{t_i+1}) &\leq \sum_{t=0}^{T-1} \beta_{t+1} \sigma_t(x_{t+1}) \\ &\leq \beta_T \sum_{t=0}^{T-1} \sigma_t(x_{t+1}) \end{aligned}$$

Bound Term 6 Set $M_1 = \sum_{i=1}^l (\mu_{t_i}(x_{t_i+1}) - \mu_{t_i}^+)$.

$$\begin{aligned} M_1 &= \sum_{i=1}^{l-1} (\mu_{t_{i-1}}(x_{t_{i-1}+1}) - \mu_{t_i}^+) + \mu_{t_l}(x_{t_l+1}) - \mu_{t_l}^+ \\ &\leq \underbrace{\sum_{i=1}^{l-1} (\mu_{t_{i-1}}(x_{t_{i-1}+1}) - \mu_{t_i}(x_{t_{i-1}+1}))}_{\text{Term 7}} \\ &\quad + \underbrace{\mu_{t_l}(x_{t_l+1}) - \mu_{t_l}^+}_{\text{Term 8}} \end{aligned}$$

Bound Term 7 Set $M_3 = \sum_{i=1}^{l-1} (\mu_{t_{i-1}}(x_{t_{i-1}+1}) - \mu_{t_i}(x_{t_{i-1}+1}))$ for simplicity. We go to bound M_3 .

$$\begin{aligned} M_3 &\leq \sum_{i=1}^{l-1} (f(x_{t_{i-1}+1}) + \beta_{t_{i-1}+1} \sigma_{t_{i-1}}(x_{t_{i-1}+1})) \\ &\quad - (f(x_{t_{i-1}+1}) - \beta_{t_i+1} \sigma_{t_i}(x_{t_{i-1}+1})) \\ &= \sum_{i=1}^{l-1} \beta_{t_{i-1}+1} \sigma_{t_{i-1}}(x_{t_{i-1}+1}) + \beta_{t_i+1} \sigma_{t_i}(x_{t_{i-1}+1}) \\ &\leq \sum_{i=1}^{l-1} (\beta_{t_{i-1}+1} + \beta_{t_i+1}) \sigma_{t_{i-1}}(x_{t_{i-1}+1}) \\ &\leq 2\beta_T \sum_{i=1}^{l-1} \sigma_{t_{i-1}}(x_{t_{i-1}+1}) \\ &\leq 2\beta_T \sum_{i=0}^{T-1} \sigma_i(x_{i+1}), \end{aligned}$$

where in the first inequality, we use Lemma 4: $\mu_{t_{i-1}}(x_{t_{i-1}+1}) \leq f(x_{t_{i-1}+1}) + \beta_{t_{i-1}}\sigma_{t_{i-1}}(x_{t_{i-1}+1})$; $\mu_{t_i}(x_{t_{i-1}+1}) \geq f(x_{t_{i-1}+1}) - \beta_{t_i}\sigma_{t_i}(x_{t_{i-1}+1})$. In the second inequality, we use the fact that $f(x_{t_{i-1}+1}) \leq f(x_{t_i+1})$. In the third inequality, we use the *decreasing monotonicity of variance functions* ((Vivarelli, 1998) and (Chowdhury et al., 2017), see Section F). Here, we use

$$\sigma_{t_i}(x_{t_{i-1}+1}) \leq \sigma_{t_{i-1}}(x_{t_{i-1}+1}),$$

because $t_i > t_{i-1}$. This step is crucial to bound M_3 . Without this step, M_3 may be bounded by two sums: $\sum_{i=1}^{l-1} \sigma_{t_{i-1}}(x_{t_{i-1}+1})$ and $\sum_{i=1}^{l-1} \sigma_{t_i}(x_{t_{i-1}+1})$. While the first term can be bounded in terms of the information gain, bounding the second is challenging, and was what led to an error in Lemma 7 of Nguyen et al. (2017).

Bound Term 8 Set $M_2 = \mu_{t_i}(x_{t_i+1}) - \mu_{t_1}^+$. We have

$$\begin{aligned} M_2 &\leq f(x_{t_i+1}) + \beta_{t_i+1}\sigma_{t_i}(x_{t_i+1}) \\ &\quad - (f(x_{t_1} - \beta_{t_1}\sigma_{t_1}(x_{t_1}))) \\ &\leq f(x_{t_i+1}) - f(x_{t_1}) + \beta_{t_i+1}\sigma_{t_i}(x_{t_i+1}) \\ &\quad + \beta_{t_1+1}\sigma_{t_1}(x_{t_1}) \\ &\leq f(x_{t_i+1}) - f(x_{t_1}) + \beta_{t_i+1} + \beta_{t_1} \\ &\leq BL\|x_{t_i+1} - x_{t_1}\|_1 + 2\beta_T \\ &\leq \mathcal{O}(\beta_T) \end{aligned}$$

The argument to achieve the bound for Term 8 is similar to Case 2.

For every x_i , where $1 \leq i \leq T-1$, Lemma 4 holds with probability $1 - \delta$. Therefore, Lemma 4 holds with probability at least $1 - \delta$ for all x_i , where $1 \leq i \leq T-1$. Combining Term 5, Term 7, Term 8, with probability $1 - T\delta$ we have

$$\sum_{t=0}^{T-1} \max\{0, f(x_{t+1}) - \mu_t^+\} \leq \mathcal{O}(\beta_T \sum_{i=1}^{T-1} \sigma_i(x_{i+1})).$$

On the other hand, following Lemma 4 of Chowdhury et al. (2017), we have $\sum_{i=1}^{T-1} \sigma_i(x_{i+1}) \leq \sqrt{4(T+2)\gamma_T}$. Thus,

$$\sum_{t=0}^{T-1} \max\{0, f(x_{t+1}) - \mu_t^+\} = \mathcal{O}(\beta_T \sqrt{T\gamma_T}).$$

Thus, for all cases, Lemma 3 holds.

3.4 Proof Sketch for Theorem 2

Theorem 2 holds using a non-trivial combination of the proof techniques as above and the technical results for π -GP-UCB Janz et al. (2020). A complete proof of Theorem 2 is provided in Appendix C of the Supplementary Material.

4 Related works

A recent review of GP-EI can be found in (Zhan and Xing, 2020).

In the noise-free setting, the most notable work in this sub-line is the one of Bull (2011) which showed that GP-EI can obtain an $\mathcal{O}(T^{-1/d})$ upper bound on the simple regret.

In noisy setting, Srinivas et al. (2012) first introduced GP-UCB for both Bayesian and non-Bayesian settings. Valko et al. (2013) introduced KernelUCB for the case of a finite-armed bandit which can be extended to a continuum-armed bandit via a discretization argument. Another notable work for GP-UCB is Improved GP-UCB (Chowdhury et al., 2017), which offers a slightly faster convergence rate. A limitation of these algorithms is that the proposed regrets are sublinear only if $2\nu > d(d+1)$. To address this limitation, (Janz et al., 2020) recently introduced π -GP-UCB which is built upon Improved GP-UCB with guarantees that π GP-UCB has a sublinear regret for every $\nu > 1$ and $d \geq 1$. GP-TS is another approach for Gaussian process bandit optimization. This was introduced by Chowdhury et al. (2017) by extending the Thompson sampling algorithm in finite-armed bandits to continuum-armed bandits. Scarlett et al. (2017) provided lower bounds for Gaussian process bandit optimization for both simple and cumulative regret. Recently Vivarelli (1998) provided new upper bounds for GP-UCB and GP-TS as well as the new bounds for the maximum information gain γ_t .

A limitation of GP-UCB and GP-TS algorithms is that the explorations are maintained via an upper confidence bound which requires to know several parameters e.g. the bound on the function RKHS norm, sub-Gaussianity level of the measurement noise. These parameters are usually unknown in practice. Consequently, these parameters are often set in a heuristic manner (Berkenkamp et al., 2017; Bogunovic et al., 2018; Janz et al., 2020; Wachi and Sui, 2020). Our GP-EI based algorithms avoid this limitation.

EI has also been studied in finite-armed bandit setting. Ryz (2016) studied EI for the problem of best-arm identification. Later, Qin et al. (2017) proposed an improvement of this algorithm on computational efficiency. However, these results and analysis techniques do not apply to our settings of Gaussian processes and the RKHS norm.

5 Experiments

While the main focus of this paper is performing theoretical analysis for GP-EI in noiseless and noisy settings, we have also proposed a new algorithm termed as




Figure 1: Comparison of methods for Hartmann3, Shekel, Hartmann6 and Ackley functions.

“Improved-GP-EI” and also provided a variant of GP-EI which we will call Modified-GP-EI in noisy setting. In this section we have performed a comparison of our Improved-GP-EI as well as Modified-GP-EI against GP-EI with fixed ω and π -GP-UCB proposed by Janz et al. (2020), which has the tightest regret so far in the noisy setting. To compare the sample-efficiency of all the algorithms, we used four functions in noisy setting: Hartmann3 ($d = 3$), Shekel ($d = 4$), Hartmann6 ($d = 6$) and Ackley ($d = 10$). We use the simple regret as evaluation metric for performance measurement of the methods. More precisely, for baselines GP-EI and π -GP-UCB, we use $\log_{10}(f(x^*) - \min_{1 \leq i \leq t} f(x_i))$. For our proposed methods, we use $\log_{10}(f(x^*) - \min_{1 \leq i \leq t} f(x_i^+))$, where $x_t^+ = \operatorname{argmax}_{1 \leq i \leq t} \mu_{t-1}(x_i)$ as defined in our algorithms. We note that we use the log only for more separated visualisation of the graph lines, which otherwise look closely.

All implementations are in Python 3.6. For each test function, we repeat the experiments 15 times. We plot the mean and a confidence bound of one standard deviation across all the runs. We used Matérn kernel with $\nu = 2.5$ and the length scale $l = 0.2$. Set $T = 100$, and $\delta = 0.05$. ω_T of Modified-GP-EI and Improved-GP-EI is set following Theorem 1 and Theorem 2. For GP-EI, we use $\omega = 1$ by default.

As seen from Figure 1, the Modified-GP-EI performs similar to GP-EI. This is expected as both methods only differ by ω_T , which is used for theoretical guarantee. π -GP-UCB is the most inefficient (except Shekel), probably because it requires to know the RKHS norm and sub-Gaussianity parameters which are unknown in practice. Following (Janz et al., 2020), we used $B = 1$, $R = 1$. In contrast, Improved-GP-EI does not need to know such hyper-parameters and clearly outperforms π -GP-UCB.

We note that our Improved-GP-EI results in a significantly more scalable algorithm. The cover construction (i.e. partitioning the space) of Improved-GP-EI permits it to perform the inverse of the kernel K_t on a subset of the data to reduce the computations from $O(t^3)$ to $O(\sum_{i=1}^p t_i^3)$ where $t = \sum_{i=1}^p t_i$ and p is the number of hypercubes in the partition. E.g. for Hartmann3, the

average runtime per iteration of Improved-GP-EI is 1.21 mins compared to GP-EI’s 1.37 min. This difference gets larger for larger horizons.

5.1 Synthetic Test Functions

In this part, we benchmark on synthetic functions in RKHS spaces. The first function is built in the RKHS space equipped with a Matérn kernel with $\nu = 2.5$ and $l = 0.2$. The second function is built in the RKHS space equipped with a SE kernel with $l = 1$. We construct each function f in a five-dimensional space by sampling 500 points $\hat{x}_1, \dots, \hat{x}_m$, uniformly on $[0, 1]^5$, and $\hat{a}_1, \dots, \hat{a}_m$ each independent uniform on $[-1, 1]$ and defining $f(x) = \sum_{i=1}^m \hat{a}_i k(\hat{x}_i, x)$ for all $x \in \mathcal{X}$ and k is a kernel. The RKHS norm of this function is computed as $\|f\|_k^2 = \sum_{i,j=1}^m \hat{a}_i \hat{a}_j k(\hat{x}_j, \hat{x}_i)$. This norm is unknown for all the algorithms. Since the baseline π -GP-UCB requires to know this norm and the sub-Gaussianity parameter R , we set $B = 1$ and $R = 1$ (in a heuristic manner) for π -GP-UCB. Otherwise, our proposed Improved-GP-EI algorithm as well as Modified-GP-EI and GP-EI do not require to know these parameters. We used $\lambda = 0.01$ for the noise setting. As seen from




Figure 2: Comparison of methods for the functions generated from the RKHS space. The left corresponds to the Matérn kernel ($\nu = 2.5$, $l = 0.2$). The right corresponds to the Squared Exponential (SE) kernel ($l = 1.0$). The SE kernel is considered as a special case of the Matérn kernel. We estimate the algorithms based on the function values at each iteration.

Figure 2, for both the Matérn kernel and the SE kernel, Improved-GP-EI outperforms all the other algorithms

while Modified-GP-EI and GP-EI show competitive performance. Improved-GP-EI performs better than Modified-GP-EI because it avoids using a large global scale variance parameter ω_T . Improved-GP-EI uses ω_T growing only poly-logarithmically with T . π -GP-UCB is outperformed by the EI based algorithms which are parameter-free. This is also observed in (Chowdhury et al., 2017). We agree with Chowdhury et al. (2017) that the UCB-based algorithms are somewhat less robust on the choice of kernel than EI-based algorithms.

6 Conclusion

We have demonstrated that GP-EI can converge with standard incumbent defined as the current best value of the GP predictive mean. Further we have proposed a variant of GP-EI, called Improved GP-EI which converges for every $\nu > 1$ and every $d \geq 1$. The empirical results have demonstrated the effectiveness of our proposed Improved GP-EI.

Acknowledgements

This research was partially funded by the Australian Government through the Australian Research Council (ARC). Prof Venkatesh is the recipient of an ARC Australian Laureate Fellowship (FL170100006).

References

On the convergence rates of expected improvement methods. *Operations Research*, 64(6):1515–1528, 2016. Winner, INFORMS Simulation Society Outstanding Publication Award, 2017.

Yasin Abbasi-yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24, pages 2312–2320. Curran Associates, Inc., 2011. URL <https://proceedings.neurips.cc/paper/2011/file/e1d5be1c7f2f456670de3d53c7b54f4a-Paper.pdf>.

Shipra Agrawal and Navin Goyal. Thompson sampling for contextual bandits with linear payoffs. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 127–135, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.

Sreejith Balakrishnan, Quoc Phong Nguyen, Bryan Kian Hsiang Low, and Harold Soh. Efficient exploration of reward functions in inverse reinforcement learning via bayesian optimization, 2020.

James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 24, pages 2546–2554. Curran Associates, Inc., 2011. URL <https://proceedings.neurips.cc/paper/2011/file/86e8f7ab32cfd12577bc2619bc635690-Paper.pdf>.

Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, pages 908–918. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/766ebcd59621e305170616ba3d3dac32-Paper.pdf>.

Felix Berkenkamp, Angela P. Schoellig, and Andreas Krause. No-regret bayesian optimization with unknown hyperparameters. *Journal of Machine Learning Research*, 20(50):1–24, 2019.

Ilija Bogunovic, Jonathan Scarlett, Stefanie Jegelka, and Volkan Cevher. Adversarially robust optimization with gaussian processes. NIPS’18, page 5765–5775, Red Hook, NY, USA, 2018. Curran Associates Inc.

Adam D. Bull. Convergence rates of efficient global optimization algorithms. *J. Mach. Learn. Res.*, 12: 2879–2904, November 2011. ISSN 1532-4435.

Sayak Ray Chowdhury, Aditya Gopalan, and abc. On kernelized multi-armed bandits. ICML’17, page 844–853. JMLR.org, 2017.

Nando De Freitas, Alex J. Smola, and Masrour Zoghi. Exponential regret bounds for gaussian process bandits with deterministic observations. In *Proceedings of the 29th International Conference on International Conference on Machine Learning*, ICML’12, page 955–962, Madison, WI, USA, 2012. Omnipress. ISBN 9781450312851.

José Miguel Hernández-Lobato, Matthew W Hoffman, and Zoubin Ghahramani. Predictive entropy search for efficient global optimization of black-box functions. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.

Matthew Hoffman, Bobak Shahriari, and Nando Freitas. On correlation and budget constraints in model-based bandit optimization with application to automatic machine learning. In Samuel Kaski and

- Jukka Corander, editors, *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, volume 33 of *Proceedings of Machine Learning Research*, pages 365–374, Reykjavik, Iceland, 22–25 Apr 2014. PMLR. URL <http://proceedings.mlr.press/v33/hoffman14.html>.
- David Janz, David Burt, and Javier Gonzalez. Bandit optimisation of functions in the matern kernel rkhs. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2486–2495. PMLR, 26–28 Aug 2020.
- Daniel Lizotte, Tao Wang, Michael Bowling, and Dale Schuurmans. Automatic gait optimization with gaussian process regression. *IJCAI’07*, page 944–949, San Francisco, CA, USA, 2007. Morgan Kaufmann Publishers Inc.
- Gustavo Malkomes and Roman Garnett. Automating bayesian optimization with bayesian optimization. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/2b64c2f19d868305aa8bbc2d72902cc5-Paper.pdf>.
- R. Marchant and F. Ramos. Bayesian optimisation for intelligent environmental monitoring. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2242–2249, 2012. doi: 10.1109/IROS.2012.6385653.
- Ruben Martinez-cantin, Nando de Freitas, Arnaud Doucet, and José A. Castellanos. Active policy learning for robot planning and exploration under uncertainty. In *IN PROCEEDINGS OF ROBOTICS: SCIENCE AND SYSTEMS*, 2007.
- J. Močkus. On bayesian methods for seeking the extremum. In G. I. Marchuk, editor, *Optimization Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, pages 400–404, Berlin, Heidelberg, 1975. Springer Berlin Heidelberg. ISBN 978-3-540-37497-8.
- Vu Nguyen and Michael A. Osborne. Knowing the what but not the where in Bayesian optimization. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7317–7326. PMLR, 13–18 Jul 2020.
- Vu Nguyen, Sunil Gupta, Santu Rana, Cheng Li, and Svetha Venkatesh. Regret for expected improvement over the best-observed value and stopping condition. In Min-Ling Zhang and Yung-Kyun Noh, editors, *Proceedings of the Ninth Asian Conference on Machine Learning*, volume 77 of *Proceedings of Machine Learning Research*, pages 279–294. PMLR, 15–17 Nov 2017.
- Michael A. Osborne. Bayesian gaussian processes for sequential prediction, optimisation and quadrature. 2010.
- Chao Qin, Diego Klabjan, and Daniel Russo. Improving the expected improvement algorithm. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 5387–5397, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005. ISBN 026218253X.
- Jonathan Scarlett, Ilija Bogunovic, and Volkan Cevher. Lower bounds on regret for noisy Gaussian process bandit optimization. In Satyen Kale and Ohad Shamir, editors, *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 1723–1742, Amsterdam, Netherlands, 07–10 Jul 2017. PMLR.
- Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 2, NIPS’12*, page 2951–2959, Red Hook, NY, USA, 2012. Curran Associates Inc.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias W. Seeger. Information-theoretic regret bounds for gaussian process optimization in the bandit setting. *IEEE Trans. Inf. Theor.*, 58(5):3250–3265, May 2012. ISSN 0018-9448. doi: 10.1109/TIT.2011.2182033. URL <http://dx.doi.org/10.1109/TIT.2011.2182033>.
- Michael L. Stein. *Interpolation of spatial data*. Springer Series in Statistics. Springer-Verlag, New York, 1999. ISBN 0-387-98629-4. doi: 10.1007/978-1-4612-1494-6. URL <http://dx.doi.org/10.1007/978-1-4612-1494-6>. Some theory for Kriging.
- Hung Tran-The, Sunil Gupta, Santu Rana, Huong Ha, and Svetha Venkatesh. Sub-linear regret bounds for bayesian optimisation in unknown search spaces. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 16271–16281. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/bb073f2855d769be5bf191f6378f7150-Paper.pdf>.

- Hung Tran-The, Sunil Gupta, Santu Rana, and Svetha Venkatesh. Bayesian optimistic optimisation with exponentially decaying regret. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10390–10400. PMLR, 18–24 Jul 2021.
- Sattar Vakili, Kia Khezeli, and Victor Picheny. On information gain and regret bounds in gaussian process bandits. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 82–90. PMLR, 13–15 Apr 2021.
- Michal Valko, Nathan Korda, Rémi Munos, Ilias Flaounas, and Nello Cristianini. Finite-time analysis of kernelised contextual bandits. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, UAI’13, page 654–663, Arlington, Virginia, USA, 2013. AUAI Press.
- Francesco Vivarelli. Studies on the generalisation of gaussian processes and bayesian neural networks. In *PhD thesis*. Aston University, 1998.
- Akifumi Wachi and Yanan Sui. Safe reinforcement learning in constrained Markov decision processes. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 9797–9806, Virtual, 13–18 Jul 2020. PMLR.
- Ziyu Wang and Nando de Freitas. Theoretical analysis of bayesian optimisation with unknown gaussian process hyper-parameters, 2014.
- Aaron Wilson, Alan Fern, and Prasad Tadepalli. Using trajectory data to improve bayesian optimization for reinforcement learning. *Journal of Machine Learning Research*, 15(8):253–282, 2014. URL <http://jmlr.org/papers/v15/wilson14a.html>.
- Dawei Zhan and Huanlai Xing. Expected improvement for expensive optimization: a review. *J. Glob. Optim.*, 78(3):507–544, 2020. doi: 10.1007/s10898-020-00923-x. URL <https://doi.org/10.1007/s10898-020-00923-x>.

Supplementary Material: Regret Bounds for Expected Improvement Algorithms in Gaussian Process Bandit Optimization

Before providing the theoretical results and the additional experiments, we summarize some important notations used in our proofs in Table 2 and in section A we first explain the derivation of the EI acquisition function in Section 2.3. Next, in Section B and C we provide the proof for Theorem 1 and Theorem 2 in the main paper, respectively.

Table 1: A summary of the notations used in the paper.

Common	Definition
d	the number of dimensions of the search space \mathcal{X}
x^*	$\operatorname{argmax}_{x \in \mathcal{X}} f(x)$, an optimal
B	the upper bound of the RKHS norm
R	the sub-Gaussianity parameter of the noise
ν	the parameter controlling the smoothness of the function
l	the lengthscale of the kernel
$\alpha_t^{EI}(x)$	the EI acquisition function at iteration t
ξ	the incumbent of EI
$\Phi(z)$	the standard normal distribution function
$\phi(z)$	the density function
$\tau(z)$	$z\Phi(z) + \phi(z)$ which is a increasing function
Section B	
ω_t	the replacement of ω in the noisy setting
x_t^+	$x_t^+ = \operatorname{argmax}\{\mu_{t-1}(x_i)\}_{x_i \in \mathcal{D}_{t-1}}$
μ_t^+	the best GP predictive mean, $\mu_t^+ = \max\{\mu_{t-1}(x_i)\}_{x_i \in \mathcal{D}_t}$
γ_t	the maximum information gain up to t iterations
Section C	
b	the constant defined as $b = \frac{d+1}{d+2\nu}$
q	the constant defined as $q = \frac{d(d+1)}{d(d+2)+2\nu}$
ρ_A	the diameter of the hypercube A
\mathcal{D}_t^A	the subset of \mathcal{D}_t in A
K_t^A	the kernel matrix, defined as $K_t^A = [k(x, x')]_{x, x' \in \mathcal{D}_t^A}$
γ_t^A	the information gain for A , defined as $\gamma_t^A = \frac{1}{2} \log I + \lambda^{-1} K_t^A $ (See (Janz et al., 2020))
\mathcal{A}_t	the set of the newly created hypercubes at iteration t and the hypercubes of \mathcal{A}_{t-1} that were not split

A Derivation of the Expected Improvement in Section 2.3.

Set $I_t(x) = \max\{0, f(x) - \mu_t^+\}$. $I_t(x)$ is positive when the prediction is higher than the best value known thus far. Otherwise, $I_t(x)$ is set to zero. The new query point is found by maximizing the expected improvement:

$$x = \operatorname{argmax}_x \mathbb{E}(I_t(x)).$$

As defined in Section 3.2, for every $x \in \mathcal{X}$, the posterior distribution of $f(x)$, at iteration t , is $\mathcal{N}(\mu_{t-1}(x), \omega^2 \sigma_{t-1}^2(x))$. Therefore, the likelihood of improvement I_t on a normal posterior distribution characterized by $\mu_{t-1}(x), \omega \sigma_{t-1}(x)$ can be computed from the normal density function:

$$\frac{1}{\sqrt{2\pi\omega\sigma_t(x)}} \exp\left(-\frac{(\mu_t(x) - f(x) - I_t(x))^2}{2\omega^2\sigma_t^2(x)}\right).$$

The expected improvement is the integral over this function:

$$\begin{aligned}
 \mathbb{E}(I_t) &= \int_{I_t=0}^{I_t=\infty} I_t \frac{1}{\sqrt{2\pi}\omega\sigma_t(x)} \exp\left(-\frac{(\mu_{t-1}(x) - f(x) - I_t(x))^2}{2\omega^2\sigma_t^2(x)}\right) dt \\
 &= \omega\sigma_t(x) \left[\frac{\mu_{t-1}(x) - \mu_t^+}{\omega\sigma_{t-1}(x)} \Phi\left(\frac{\mu_{t-1}(x) - \mu_t^+}{\omega\sigma_{t-1}(x)}\right) + \phi\left(\frac{\mu_{t-1}(x) - \mu_t^+}{\omega\sigma_{t-1}(x)}\right) \right] \\
 &= (\mu_{t-1}(x) - \mu_t^+) \Phi\left(\frac{\mu_{t-1}(x) - \mu_t^+}{\omega\sigma_{t-1}(x)}\right) + \omega\sigma_{t-1}(x) \phi\left(\frac{\mu_{t-1}(x) - \mu_t^+}{\omega\sigma_{t-1}(x)}\right)
 \end{aligned}$$

Setting $u = \mu_{t-1}(x) - \mu_t^+$ and $v = \omega\sigma_{t-1}(x)$, and $\rho(u, v) = u\Phi(\frac{u}{v}) + v\phi(\frac{u}{v})$, we obtain the formula of EI as :

$$\alpha_t^{EI}(x) = \rho(u, v) = \rho(\mu_{t-1}(x) - \mu_t^+, \omega\sigma_{t-1}(x)).$$

B Proof of Theorem 1

Instead of upper bounding directly the simple regret, we will seek to upper bound the sum $\sum_{t=1}^T r_t$ to exploit the results from the maximum information gain (γ_T) in the noisy setting. The proof for Theorem 1 involves two steps.

- Upper bounding the instantaneous regret $r_t = f(x^*) - f(x_t^+)$.
- Upper bounding the sum $\sum_{t=1}^T r_t$

B.1 Upper bounding the instantaneous regret $r_t = f(x^*) - f(x_t^+)$:

To obtain a bound on r_t , we break down r_t into two terms as follows:

$$\begin{aligned}
 r_t &= f(x^*) - f(x_t^+) \\
 &= \underbrace{f(x^*) - \mu_t^+}_{\text{Term 1}} + \underbrace{\mu_t^+ - f(x_t^+)}_{\text{Term 2}}
 \end{aligned}$$

Upper Bounding Term 1. First, we provide the lower bound and upper bound for the acquisition function $\alpha_t^{EI}(x)$ in the noisy setting. These bounds are similar to those in the noiseless setting (See Lemma 4).

Lemma 6 (Based on Lemma 9 of (Wang and de Freitas, 2014)). *Pick $\delta \in (0, 1)$. For $x \in \mathcal{X}$, $t \in \mathbb{N}$, set $I_t(x) = \max\{0, f(x) - \mu_t^+\}$. Then with probability at least $1 - \delta$ we have*

$$I_t(x) - \beta_t\sigma_{t-1}(x) \leq \alpha_t^{EI}(x) \leq I_t(x) + (\beta_t + \omega_t)\sigma_{t-1}(x).$$

Proof. If $\sigma_{t-1}(x) = 0$ then $\alpha_t^{EI}(x) = I_t(x)$, which makes the result trivial. We now assume that $\sigma_{t-1}(x) > 0$. Set $q = \frac{f(x) - \mu_t^+}{\sigma_{t-1}(x)}$ and $u = \frac{\mu_{t-1}(x) - \mu_t^+}{\sigma_{t-1}(x)}$. Then we have that

$$\alpha_t^{EI}(x) = \omega_t\sigma_{t-1}(x)\tau\left(\frac{u}{\omega_t}\right).$$

By Lemma 11, we have that $|u - q| \leq \beta_t$ with probability $1 - \delta$. Set $\tau(z) = z\Phi(z) + \phi(z)$. As $\tau'(z) = \Phi(z) \in [0, 1]$, τ is non-decreasing and $\tau(z) \leq 1 + z$ for $z > 0$. Hence,

$$\begin{aligned}
 \alpha_t^{EI}(x) &\leq \omega_t\sigma_{t-1}(x)\tau\left(\frac{\max\{0, q\} + \beta_t}{\omega_t}\right) \\
 &\leq \omega_t\sigma_{t-1}(x)\left(\frac{\max\{0, q\} + \beta_t}{\omega_t} + 1\right) \\
 &= I_t(x) + (\beta_t + \omega_t)\sigma_{t-1}(x)
 \end{aligned}$$

If $I_t(x) = 0$ then the lower bound is trivial as $\alpha_t^{EI}(x)$ is non-negative. Thus suppose $I_t(x) > 0$. Since $\alpha_t^{EI}(x) \geq 0$ and $\tau(z) \geq 0$ for all z , and $\tau(z) = z + \tau(-z) \geq z$. Therefore,

$$\begin{aligned}\alpha_t^{EI}(x) &\geq \omega_t \sigma_{t-1}(x) \tau\left(\frac{q - \beta_t}{\omega_t}\right) \\ &\geq \omega_t \sigma_{t-1}(x) \left(\frac{q - \beta_t}{\omega_t}\right) \\ &= I_t(x) - \beta_t \sigma_{t-1}(x)\end{aligned}$$

■

Now we will use the results from Lemma 6 to upper bound Term 1 as in the follow lemma.

Lemma 7. *Pick $\delta \in (0, 0.5)$. Then with probability at least $1 - 2\delta$ we have*

$$f(x^*) - \mu_t^+ \leq \frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} (\max\{0, f(x_t) - \mu_{t-1}^+\} + (\beta_t + \omega_t)\sigma_{t-1}(x_t)).$$

Proof. If $\sigma_{t-1}(x^*) = 0$ then by definition of $\alpha_t^{EI}(x)$, we have $\alpha_t^{EI}(x^*) = I_t(x^*)$. We have

$$\begin{aligned}I_t(x^*) &= \alpha_t^{EI}(x^*) \\ &\leq \alpha_t^{EI}(x_t) \\ &\leq \max\{0, f(x_t) - \mu_t^+\} + (\beta_t + \omega_t)\sigma_{t-1}(x_t) \\ &\leq \frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} (\max\{0, f(x_t) - \mu_{t-1}^+\} + (\beta_t + \omega_t)\sigma_{t-1}(x_t)),\end{aligned}$$

where in the first inequality, we use the definition $\alpha_t^{EI}(x) = \max_{x \in \mathcal{X}} \alpha_t^{EI}(x)$. In the second inequality, we use Lemma 6. The third inequality holds since $\frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} \geq \frac{\tau(0)}{\tau(0)} = 1$ due to the fact that function $\tau(z)$ is an increasing function. Thus, the lemma holds with probability $1 - \delta$.

We now consider $\sigma_{t-1}(x^*) > 0$. If $f(x^*) < \mu_t^+$ then the lemma will be trivial. We now consider $f(x^*) \geq \mu_t^+$. By Lemma 11, $\mu_{t-1}(x^*) - f(x^*) \geq -\beta_t \sigma_{t-1}(x^*)$ with probability $1 - \delta$. Combining with the fact that $f(x^*) \geq \mu_t^+$, we have that $\frac{\mu_{t-1}(x^*) - \mu_t^+}{\omega_t \sigma_{t-1}(x^*)} \geq \frac{-\beta_t}{\omega_t}$ with probability $1 - \delta$. On the other hand, following the derivation of the acquisition function EI , we have $\alpha_t^{EI}(x^*) = \omega_t \sigma_{t-1}(x^*) \tau\left(\frac{\mu_{t-1}(x^*) - \mu_t^+}{\omega_t \sigma_{t-1}(x^*)}\right)$. Therefore, $\alpha_t^{EI}(x^*) \geq \omega_t \sigma_{t-1}(x^*) \tau\left(-\beta_t/\omega_t\right)$ with probability $1 - \delta$ due to the fact that $\tau(z)$ is an increasing function.

By combining inequalities $\alpha_t^{EI}(x^*) \geq \omega_t \sigma_{t-1}(x^*) \tau\left(-\frac{\beta_t}{\omega_t}\right)$, $\alpha_t^{EI}(x^*) \geq I_t(x^*) - \beta_t \sigma_{t-1}(x^*)$ which is proven in Lemma 6, we obtain $\frac{\beta_t}{(\omega_t \tau\left(-\frac{\beta_t}{\omega_t}\right))} \alpha_t^{EI}(x^*) + \alpha_t^{EI}(x^*) \geq I_t(x^*)$. Now, using the fact that $\tau(z) = z + \tau(-z)$ for $z = \frac{\beta_t}{\omega_t}$, we obtain

$$I_t(x^*) \leq \frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} EI_t(x^*) \quad (6)$$

This inequality (6) holds with probability $1 - \delta$. Finally, we achieve

$$\begin{aligned}f(x^*) - \mu_t^+ &\leq I_t(x^*) \\ &\leq \frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} EI_t(x^*) \\ &\leq \frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} EI_t(x_t) \\ &\leq \frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} (\max\{0, f(x_t) - \mu_t^+\} + (\beta_t + \omega_t)\sigma_{t-1}(x_t)),\end{aligned}$$

where the first inequality holds by the definition of the function I_t . The second one comes from (6). The third one holds by the property of the chosen point $x_t = \operatorname{argmax}_x \alpha_t^{EI}(x)$. The final inequality hold due to Lemma 6. ■

Upper bounding Term 2. Bounding Term 2 is our important result. This is represented in the following lemma.

Lemma 8. *Pick a $\delta \in (0, 1)$. Then with probability $1 - \delta$ we have*

$$\mu_t^+ - f(x_t^+) \leq \frac{\beta_t}{\omega_t} (\sqrt{2\pi}(\beta_t + \omega_t)\sigma_{t-1}(x_t) + \sqrt{2\pi}\max\{0, (f(x_t) - \mu_t^+)\}).$$

Proof. By the definition of our GP-EI algorithm, $x_t^+ = \operatorname{argmax}_{x \in \mathcal{D}_t} \mu_{t-1}(x)$. It implies that $\mu_t(x_t^+) = \mu_t^+$. Hence,

$$\begin{aligned} \alpha_t^{EI}(x_t^+) &= \omega_t \sigma_{t-1}(x_t^+) \tau\left(\frac{\mu_t(x_t^+) - \mu_t^+}{\omega_t \sigma_{t-1}(x_t^+)}\right) \\ &= \omega_t \sigma_{t-1}(x_t^+) \tau(0) \\ &= \frac{1}{\sqrt{2\pi}} \omega_t \sigma_{t-1}(x_t^+) \end{aligned}$$

Combining this with the fact that $\alpha_t^{EI}(x_t) = \max_{x \in \mathcal{X}} \alpha_t^{EI}(x)$, we have

$$\begin{aligned} \frac{1}{\sqrt{2\pi}} \omega_t \sigma_{t-1}(x_t^+) &= \alpha_t^{EI}(x_t^+) \\ &\leq \alpha_t^{EI}(x_t) \\ &\leq \max\{0, f(x_t) - \mu_t^+\} + (\beta_t + \omega_t)\sigma_{t-1}(x_t) \end{aligned}$$

where the inequality holds due to Lemma 6.

Thus,

$$\omega_t \sigma_{t-1}(x_t^+) \leq (\sqrt{2\pi}(\beta_t + \omega_t)\sigma_{t-1}(x_{t+1}) + \sqrt{2\pi}\max\{0, (f(x_t) - \mu_t^+)\}) \quad (7)$$

Finally, we can upper bound $\mu_t^+ - f(x_t^+)$ with probability $1 - \delta$ based on (7) as follows:

$$\begin{aligned} \mu_t^+ - f(x_t^+) &\leq \beta_t \sigma_{t-1}(x_t^+) \\ &\leq \frac{\beta_t}{\omega_t} (\sqrt{2\pi}(\beta_t + \omega_t)\sigma_{t-1}(x_t) + \sqrt{2\pi}\max\{0, (f(x_t) - \mu_t^+)\}) \end{aligned}$$

where in the first inequality, $\mu_t^+ - f(x_t^+) \leq \beta_t \sigma_{t-1}(x_t^+)$ with probability $1 - \delta$ due to Lemma 11. $f(x_t') - f(x_t^+) \leq 0$ because by definition $x_t' \in \mathcal{D}_t$ and $x_t^+ = \operatorname{argmax}_{x_i \in \mathcal{D}_t} f(x_i)$. The second inequality holds with probability $1 - \delta$ due to Eq(7). Thus, the lemma holds with probability $1 - \delta$. ■

Upper bounding the instantaneous regret $r_t = f(x^*) - f(x_t^+)$. Combining Term 1 (in Lemma 7) and Term 2 (in Lemma 8), with high probability we have

$$r_t \leq \left(\frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} + \sqrt{2\pi} \frac{\beta_t}{\omega_t} \right) [\max\{0, f(x_t) - \mu_t^+\} + (\beta_t + \omega_t)\sigma_{t-1}(x_t)]$$

.

Now we make this bound simple. By using the assumption of Theorem 2, $\omega_t = \sqrt{\gamma_{t-1} + 1 + \ln(\frac{1}{\delta})}$, we can bound

the ratio $\frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)}$ as follows:

Lemma 9. *There exists $C > 0$ and $\frac{\tau\left(\frac{\beta_t}{\omega_t}\right)}{\tau\left(-\frac{\beta_t}{\omega_t}\right)} \leq C$ for every $T \geq 1$ and $1 \leq t \leq T$.*

Proof. By definition, $\beta_t = B + \sqrt{2(\gamma_{t-1} + 1 + \ln(1/\delta))}$. Hence, $\frac{\beta_t}{\omega_t} \leq B + \sqrt{2}$. Since the function $\tau(z)$ is non-decreasing, we have that $\tau(\frac{\beta_t}{\omega_t}) \leq \tau(B + \sqrt{2})$, and $\tau(-\frac{\beta_t}{\omega_t}) \geq \tau(-B - \sqrt{2})$. Thus,

$$\frac{\tau(\frac{\beta_t}{\omega_t})}{\tau(-\frac{\beta_t}{\omega_t})} \leq \frac{\tau(B + \sqrt{2})}{\tau(-B - \sqrt{2})}.$$

Setting $C = \frac{\tau(B + \sqrt{2})}{\tau(-B - \sqrt{2})}$ which is a constant, we have that $\frac{\tau(\frac{\beta_t}{\omega_t})}{\tau(-\frac{\beta_t}{\omega_t})} \leq C$ for every $T \geq 1$ and $1 \leq t \leq T$. The lemma holds. \blacksquare

Thus, we obtain an upper bound on r_t as follows:

Lemma 10. *There exist constant $C > 0$ such that*

$$r_t \leq (C + \sqrt{2\pi}(B + \sqrt{2}))(I_t + (\beta_t + \omega_t)\sigma_{t-1}(x_t)),$$

where $I_t = \max\{0, f(x_t) - \mu_t^+\}$.

Proof.

$$\begin{aligned} r_t &\leq \left(\frac{\tau(\frac{\beta_t}{\omega_t})}{\tau(-\frac{\beta_t}{\omega_t})} + \sqrt{2\pi} \frac{\beta_t}{\omega_t} \right) [\max\{0, f(x_t) - \mu_t^+\} + (\beta_t + \omega_t)\sigma_{t-1}(x_t)] \\ &\leq (C + \sqrt{2\pi}(B + \sqrt{2}))(I_t + (\beta_t + \omega_t)\sigma_{t-1}(x_t)), \end{aligned}$$

\square

B.2 Upper bounding the sum $\sum_{t=1}^T r_t$

Using Lemma 10, we obtain an upper bound for $\sum_{t=1}^T r_t$ as follows:

$$\sum_{t=1}^T r_t \leq (C + \sqrt{2\pi}B + 2\sqrt{\pi}) \underbrace{\sum_{t=1}^{T-1} I_t}_{\text{Term 3}} + (C + \sqrt{2\pi}B + 2\sqrt{\pi}) \underbrace{\sum_{t=0}^{T-1} (\beta_t + \omega_t)\sigma_t(x_{t+1})}_{\text{Term 4}}.$$

To upper bound the sum, we will go to upper bound Term 3 and Term 4. While Term 4 can be bounded via the maximum information gain similar as in the existing works of GP-UCB and GP-TS (Srinivas et al., 2012; Chowdhury et al., 2017), bounding Term 3 is the key step in our proof for the acquisition function GP-EI.

Before providing upper bounds for Term 3 and Term 4, we restate some important results from existing works in the following section.

B.2.1 Auxiliary Lemmas

Lemma 11 (Theorem 2 of (Chowdhury et al., 2017)). *Pick $\delta \in (0, 1)$. Fix a horizon $T > 1$. We define $\beta_t = B + R\sqrt{2(\gamma_{t-1} + 1 + \ln(1/\delta))}$ for every $1 \leq t \leq T$. Then*

$$\mathbb{P}(\forall 1 \leq t \leq T, \forall x \in \mathcal{X}, |f(x) - \mu_{t-1}(x)| \leq \beta_t \sigma_{t-1}(x)) \geq 1 - \delta.$$

Lemma 12 (Lemma 3 of (Chowdhury et al., 2017)). *The information gain for the points selected can be expressed in terms of the predictive variances as follows:*

$$I(y_T; f_T) = \frac{1}{2} \sum_{t=1}^T \ln(1 + \lambda^{-1} \sigma_{t-1}(x_t)).$$

Lemma 13 (Lemma 5 of (De Freitas et al., 2012)). *When $f \in H_k(\mathcal{X})$, then for every $x, y \in \mathcal{X}$, we have*

$$|f(x) - f(y)| \leq BL\|x - y\|_1,$$

where L is a Lipschitz constant in $H_k(\mathcal{X})$.

Proof. Lemma 5 of (De Freitas et al., 2012) holds for functions in an RKHS equipped by a kernel k . The Lipschitz constant is defined as

$$\sup_{x \in \mathcal{X}} \partial_x \partial_{x'} k(x - x')|_{x'=x}$$

, so the second derivative on k is needed. For SE kernel, it always holds. For Matérn kernel, to our best understanding, if we have $\nu \geq 1/2$, we can ensure the differentiability. \square

We now are ready to bound Term 3 and Term 4.

B.2.2 Upper Bounding Term 3

Lemma 14. *Pick $\delta \in (0, 1)$. Then with probability at least $1 - \delta$ we have that*

$$\sum_{t=1}^T I_t = \mathcal{O}(\beta_T \sqrt{T\gamma_T}).$$

Proof. Set $S_T = \sum_{t=0}^{T-1} I_t$. There are three cases to be considered:

Case 1 $S_T = 0$. This happens when for every t : $f(x_{t+1}) - \mu_t^+ \leq 0$.

Case 2 There exists a unique index $1 \leq t' \leq T$ such that $f(x_{t'+1}) - \mu_{t'}^+ > 0$. It follows that $S_T = f(x_{t'+1}) - \mu_{t'}^+$. In this case, we have that

$$\begin{aligned} S_T &= f(x_{t'+1}) - \mu_{t'}^+ \\ &\leq f(x_{t'+1}) - (f(x') - \beta_{t'+1}\sigma_{t'}(x')) \\ &\leq f(x_{t'+1}) - f(x') + \beta_{t'+1}\sigma_{t'}(x') \\ &\leq BL\|x_{t'+1} - x'\|_1 + \beta_{t'+1} \\ &= \mathcal{O}(\beta_T), \end{aligned}$$

where in the last inequality, we use Lemma 5, the inequality $\beta_{t'+1} \leq \beta_T$, and the fact that $\sigma'_t(x) \leq 1$. Finally, because the domain \mathcal{X} is bounded, $\|x_{t'+1} - x'\|_1$ is bounded.

Case 3 There are $0 \leq t_1 < t_2, \dots, < t_l \leq T - 1$ where $l \geq 2$ such that $f(x_{t_i+1}) \geq \mu_{t_i}^+$. Thus, we have

$$\begin{aligned} \sum_{t=0}^{T-1} I_t &= \sum_{t=1}^T \max\{0, f(x_{t+1}) - \mu_t^+\} \\ &= \sum_{i=1}^l (f(x_{t_i+1}) - \mu_{t_i}^+) \\ &\leq \sum_{i=1}^l (\beta_{t_i+1}\sigma_{t_i}(x_{t_i+1}) + \mu_{t_i}(x_{t_i+1}) - \mu_{t_i}^+) \\ &\leq \underbrace{\sum_{i=1}^l \beta_{t_i+1}\sigma_{t_i}(x_{t_i+1})}_{\text{Term 5}} + \underbrace{\sum_{i=1}^l (\mu_{t_i}(x_{t_i+1}) - \mu_{t_i}^+)}_{\text{Term 6}} \end{aligned}$$

Bound Term 5

$$\begin{aligned} \sum_{i=1}^l \beta_{t_i+1} \sigma_{t_i}(x_{t_i+1}) &\leq \sum_{t=0}^{T-1} \beta_{t+1} \sigma_t(x_{t+1}) \\ &\leq \beta_T \sum_{t=0}^{T-1} \sigma_t(x_{t+1}) \end{aligned}$$

Bound Term 6 Set $M_1 = \sum_{i=1}^l (\mu_{t_i}(x_{t_i+1}) - \mu_{t_i}^+)$.

$$\begin{aligned} M_1 &= \mu_{t_l}(x_{t_l+1}) - \mu_{t_1}^+ + \sum_{i=1}^{l-1} (\mu_{t_{i-1}}(x_{t_{i-1}+1}) - \mu_{t_i}^+) \\ &\leq \underbrace{\mu_{t_l}(x_{t_l}) - \mu_{t_1}^+}_{\text{Term 7}} + \underbrace{\sum_{i=1}^{l-1} (\mu_{t_{i-1}}(x_{t_{i-1}+1}) - \mu_{t_i}(x_{t_{i-1}+1}))}_{\text{Term 8}} \end{aligned}$$

Bound Term 7 Set $M_2 = \mu_{t_l}(x_{t_l+1}) - \mu_{t_1}^+$. We have

$$\begin{aligned} M_2 &\leq f(x_{t_l+1}) + \beta_{t_l+1} \sigma_{t_l}(x_{t_l+1}) - (f(x_{t_1}) - \beta_{t_1} \sigma_{t_1}(x_{t_1})) \\ &\leq f(x_{t_l+1}) - f(x_{t_1}) + \beta_{t_l+1} \sigma_{t_l}(x_{t_l+1}) + \beta_{t_1+1} \sigma_{t_1}(x_{t_1}) \\ &\leq f(x_{t_l+1}) - f(x_{t_1}) + \beta_{t_l+1} + \beta_{t_1} \\ &\leq BL \|x_{t_l+1} - x_{t_1}\|_1 + 2\beta_T \\ &\leq \mathcal{O}(\beta_T) \end{aligned}$$

The argument to achieve the bound for Term 7 is similar to Case 2.

Bound Term 8 Set $M_3 = \sum_{i=1}^{l-1} (\mu_{t_{i-1}}(x_{t_{i-1}+1}) - \mu_{t_i}(x_{t_{i-1}+1}))$ for simplicity. We go to bound M .

$$\begin{aligned} M_3 &\leq \sum_{i=1}^{l-1} (f(x_{t_{i-1}+1}) + \beta_{t_{i-1}+1} \sigma_{t_{i-1}}(x_{t_{i-1}+1}) \\ &\quad - (f(x_{t_{i-1}+1}) - \beta_{t_i+1} \sigma_{t_i}(x_{t_{i-1}+1}))) \\ &= \sum_{i=1}^{l-1} \beta_{t_{i-1}+1} \sigma_{t_{i-1}}(x_{t_{i-1}+1}) + \beta_{t_i+1} \sigma_{t_i}(x_{t_{i-1}+1}) \\ &\leq \sum_{i=1}^{l-1} (\beta_{t_{i-1}+1} + \beta_{t_i+1}) \sigma_{t_{i-1}}(x_{t_{i-1}+1}) \\ &\leq 2\beta_T \sum_{i=1}^{l-1} \sigma_{t_{i-1}}(x_{t_{i-1}+1}) \\ &\leq 2\beta_T \sum_{i=0}^{T-1} \sigma_i(x_{i+1}), \end{aligned}$$

where in the first inequality, we use Lemma 4: $\mu_{t_{i-1}}(x_{t_{i-1}+1}) \leq f(x_{t_{i-1}+1}) + \beta_{t_{i-1}} \sigma_{t_{i-1}}(x_{t_{i-1}+1})$; $\mu_{t_i}(x_{t_{i-1}+1}) \geq f(x_{t_{i-1}+1}) - \beta_{t_i} \sigma_{t_i}(x_{t_{i-1}+1})$. In the second inequality, we use the fact that $f(x_{t_{i-1}+1}) \leq f(x_{t_i+1})$. In the third inequality, we use the *decreasing monotonicity of variance functions* ((Vivarelli, 1998) and (Chowdhury et al., 2017), see Section F). Here, we use

$$\sigma_{t_i}(x_{t_{i-1}+1}) \leq \sigma_{t_{i-1}}(x_{t_{i-1}+1}),$$

because $t_i > t_{i-1}$. This step is crucial to bound M_3 . Without this step, M_3 may be bounded by two sums: $\sum_{i=1}^{l-1} \sigma_{t_{i-1}}(x_{t_{i-1}+1})$ and $\sum_{i=1}^{l-1} \sigma_{t_i}(x_{t_{i-1}+1})$. While the first term can be bounded in terms of the information gain, bounding the second term is very challenging.

For every x_i , where $1 \leq i \leq T - 1$, Lemma 4 holds with probability $1 - \delta$. Therefore, Lemma 4 holds with probability at least $1 - \delta$ for all x_i , where $1 \leq i \leq T - 1$. Combining Term 5, Term 7, Term 8, with probability $1 - T\delta$ we have

$$\sum_{t=0}^{T-1} \max\{0, f(x_{t+1}) - \mu_t^+\} \leq \mathcal{O}(\beta_T \sum_{i=1}^{T-1} \sigma_i(x_{i+1})).$$

On the other hand, following Lemma 4 of Chowdhury et al. (2017), we have $\sum_{i=1}^{T-1} \sigma_i(x_{i+1}) \leq \sqrt{4(T+2)\gamma_T}$. Thus,

$$\sum_{t=0}^{T-1} \max\{0, f(x_{t+1}) - \mu_t^+\} = \mathcal{O}(\beta_T \sqrt{T\gamma_T}).$$

■

B.2.3 Upper Bounding Term 4

Lemma 15. *Let x_1, \dots, x_t be the points selected by Algorithm 2. The sum of predictive standard deviation at those points can be expressed in terms of the maximum information gain. More precisely,*

$$\sum_{t=0}^{T-1} \sigma_t(x_{t+1}) \leq \sqrt{4(T+1)\gamma_T}.$$

Proof. By Cauchy-Schwartz inequality, $\sum_{t=0}^{T-1} \sigma_t(x_{t+1}) \leq \sqrt{T \sum_{t=1}^{T-1} \sigma_t^2(x_{t+1})}$. By assumption, $0 \leq k(x, x) \leq 1$. It implies that $0 \leq \sigma_{t-1}^2(x) \leq 1$ for all $x \in \mathcal{X}$. Following Lemma 4 of (Chowdhury et al., 2017), we get $\sum_{t=1}^T \sigma_{t-1}(x_t) \leq \sqrt{4T(1+2/T)\gamma_T} \leq \sqrt{4(T+2)\gamma_T}$. Note that we here use $\lambda = 1 + 2/T$ as the setting of (Chowdhury et al., 2017).

Thus, we have that $\sum_{t=0}^{T-1} \sigma_t(x_{t+1}) \leq \sqrt{4(T+2)\gamma_T}$ ■

Combining Lemma 14 and Lemma 15, we obtain an upper bound for the cumulative regret $\sum_{t=1}^T r_t$ as follows:

Theorem 3. *Pick $\delta \in (0, 1)$. Then with probability at least $1 - \delta$, the cumulative regret of Algorithm 1 is bounded as:*

$$R_T = \mathcal{O}(\gamma_T \sqrt{T}).$$

C Proof of Theorem 2

In this section, we provide a complete proof for Theorem 2 in the main paper. Here we use a non-trivial combination of the proof techniques as above and the technical results for π -GP-UCB (Janz et al., 2020). A key difference from (Janz et al., 2020) lies in steps from equation 16 to 21 to obtain an upper bound for R_T^+ which is easy in (Janz et al., 2020).

To obtain an upper bound the sum of $R_T = \sum_{t=1}^T r_t$, where $r_t = f(x^*) - f(x_t^+)$, we denote $\tilde{\mathcal{A}}_T = \cup_{t \leq T} \mathcal{A}_t$, the set of all cover elements created until time T , and define the initial time for an element $A \in \tilde{\mathcal{A}}_T$, as $\phi(A) = \min\{t : A \in \mathcal{A}_t\}$, and the terminal time as $\phi'(A) = \max\{t : A \in \mathcal{A}_t\}$.

Unlike the analysis of the sum $R_T = \sum_{t=1}^T r_t$ in the proof of Theorem 1, we here analyze $R_T = \sum_{t=1}^T r_t$ into

groups based on the partitioning of the searching space.

$$R_T^+ = \sum_{t=1}^T r_t \quad (8)$$

$$= \sum_{t=1}^T \sum_{A \in \mathcal{A}_t} \mathbf{1}\{x_t \in A\} r_t \quad (9)$$

$$= \sum_{A \in \tilde{\mathcal{A}}_T} \sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t \quad (10)$$

The next step is to estimate each component $\sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t$. Given a set $A \in \mathcal{A}_T$, we consider the set $\mathcal{D}_T \cap A$. Without loss of generality, assume that $\mathcal{D}_T \cap A = \{x_1^A, \dots, x_{n(A)}^A\}$ with $n(A) = |\mathcal{D}_T \cap A|$. We assume that this order is the real order of time that x_i^A is selected by the algorithm. It means that x_1^A corresponds to time t_1, \dots , and $x_{n(A)}^A$ corresponds to time $t_{n(A)}$ where $1 \leq t_1 < t_2 < \dots < t_{n(A)} \leq T$. In addition, for every $1 \leq i \leq n(A)$, $i \leq t_i$. Let $x_i^{A+} = \operatorname{argmax}_{x_j^A \in \mathcal{D}_T \cap A, 1 \leq j \leq i} f(x_j^A)$ and let $r_i^A = f(x^*) - f(x_i^{A+})$ as the instantaneous regret at iteration i for the sampled points in A . We have that

$$\sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t = \sum_{i=1}^{n(A)} r_{t_i} \quad (11)$$

$$= \sum_{i=1}^{n(A)} (f(x^*) - f(x_{t_i}^+)) \quad (12)$$

$$\leq \sum_{i=1}^{n(A)} (f(x^*) - f(x_i^{A+})) \quad (13)$$

$$= \sum_{i=1}^{n(A)} r_i^A, \quad (14)$$

where in Eq (12), we use the fact the $f(x_{t_i}^+) \geq f(x_i^{A+})$. Indeed, we have that $x_i^{A+} = \operatorname{argmax}_{x_j^A \in \mathcal{D}_T \cap A, 1 \leq j \leq i} f(x_j^A)$ which is defined as above. $x_{t_i}^+$ corresponds to x_i^A in set A . Therefore, $\{x_1^A, \dots, x_i^A\}$ is the subset of $\{x_1, x_2, \dots, x_{t_i}\}$. It implies that $f(x_{t_i}^+) \geq f(x_i^{A+})$.

By Lemma 21 which we will provide in section C.1, we have that

$$\sum_{i=1}^{n(A)} r_i^A \leq \tilde{O}(\sqrt{n(A)\gamma_{\phi'(A)}^A}), \quad (15)$$

where $\gamma_{\phi'(A)}^A$ is the information gain of A at iteration $\phi'(A) = \max\{t : A \in \mathcal{A}_t\}$ (See Lemma 17 in Section C.1).

Thus, $\sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t \leq \tilde{O}(\sqrt{n(A)\gamma_{\phi'(A)}^A})$. It is equivalent that

$$\frac{(\sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t)^2}{n(A)} \leq \tilde{O}(\gamma_{\phi'(A)}^A) \quad (16)$$

Since the inequality at Eq(16) holds for every $A \in \tilde{\mathcal{A}}_T$, we get

$$\sum_{A \in \tilde{\mathcal{A}}_T} \frac{(\sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t)^2}{n(A)} \leq \sum_{A \in \tilde{\mathcal{A}}_T} \tilde{O}(\gamma_{\phi'(A)}^A) \quad (17)$$

On the other hand, applying the Cauchy–Schwarz inequality, we have that

$$\left(\sum_{A \in \tilde{\mathcal{A}}_T} n(A) \left(\sum_{A \in \tilde{\mathcal{A}}_T} \frac{(\sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t)^2}{n(A)} \right) \right) \geq \left(\sum_{A \in \tilde{\mathcal{A}}_T} \sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t \right)^2 \quad (18)$$

As the sets A in $\tilde{\mathcal{A}}_T$ are not overlapping, $\sum_{A \in \tilde{\mathcal{A}}_T} n(A) = T$. Replacing this to Eq(18), we get

$$\sqrt{T \left(\sum_{A \in \tilde{\mathcal{A}}_T} \frac{(\sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t)^2}{n(A)} \right)} \geq \sum_{A \in \tilde{\mathcal{A}}_T} \sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t \quad (19)$$

Combining Eq(10), Eq(17) and Eq(19), we get that

$$R_T^+ = \sum_{A \in \tilde{\mathcal{A}}_T} \sum_{t=\phi(A)}^{\phi'(A)} \mathbf{1}\{x_t \in A\} r_t \quad (20)$$

$$\leq \sqrt{T \left(\sum_{A \in \tilde{\mathcal{A}}_T} \tilde{\mathcal{O}}(\gamma_{\phi'(A)}^A) \right)} \quad (21)$$

By Lemma 18 which is proven in Section C.1, the number of the sets $A \in \tilde{\mathcal{A}}_T$ is $\mathcal{O}(T^q)$, where $q = \frac{d(d+1)}{d(d+2)+2\nu}$ (See Section 5.1). By Lemma 17 in Section C.1, $\gamma_{\phi'(A)}^A$ is bounded by a logarithmic function in T for all $A \in \tilde{\mathcal{A}}_T$. Thus, we have $R_T \leq \mathcal{O}(T^{\frac{q+1}{2}})$. Using the definition of $q = \frac{d(d+1)}{d(d+2)+2\nu}$, we get that $R_T = \tilde{\mathcal{O}}(T^{\frac{-2\nu-d}{d(2d+4)+4\nu}})$.

C.1 Auxiliary Lemmas

Lemma 16. *For every $1 \leq i \leq T$, we have that $\gamma_i \leq \gamma_T$.*

Proof. By Lemma 12, for a set of point $A_i = x_1, x_2, \dots, x_i$, we have that $I(y_i; f_i) = \frac{1}{2} \sum_{t=1}^i \ln(1 + \lambda^{-1} \sigma_{t-1}(x_t))$. Consider a set $A_T = A_i \cup \{x_{i+1}, \dots, x_T\}$ containing T elements. We also have that $I(y_{A_T}; f_{A_T}) = \frac{1}{2} \sum_{t=1}^T \ln(1 + \lambda^{-1} \sigma_{t-1}(x_t))$. Thus, for each A_i which contains i elements, we always can construct a set A_T containing T elements such that $I(y_i; f_i) \leq I(y_T; f_T)$. Thus,

$$\gamma_i = \max_{A \in \mathcal{X}: |A|=i} I(y_A : f_A) \leq \max_{A \in \mathcal{X}: |A|=T} I(y_A : f_A) = \gamma_T. \quad \blacksquare$$

Lemma 17 (Lemma 1 of (Janz et al., 2020)). *Let A be a subset of \mathcal{X} and assume that there exists a $1 \leq \phi \leq T$ such that $A \in \mathcal{A}_\tau$. Let $\phi'(A) = \max\{1 \leq t \leq T : A \in \mathcal{A}_t\}$. Then for some $C > 0$, $\gamma_{\phi'(A)}^A \leq C \ln(T) \ln \ln(T)$.*

Lemma 18 (Lemma 2 of (Janz et al., 2020)). *Let A be the covering set at time t . Assume that $\mathcal{A}_2 = \mathcal{O}(T^q)$. Then for T sufficiently large, $|\cup_{t \leq T} \mathcal{A}_t| \leq C_{d,\nu} T^q$, where $C_{d,\nu} > 0$ depends on d and ν only.*

Lemma 19 (Lemma 5 of (Janz et al., 2020)). *Given $\delta \in (0, 1)$. for all $t \leq T$, for all $A \in \cup_{t \leq T} \mathcal{A}_t$, we have*

$$\mathbb{P}(\forall t, \forall x \in A, |f(x) - \mu_{t-1}^A(x)| \leq \hat{\beta}_t^A \sigma_{t-1}(x)) \geq 1 - \delta,$$

where $\hat{\beta}_t^A = B + R \sqrt{2(\gamma_{t-1}^A + 1 + \ln(1/\delta))}$, and $\gamma_{t-1}^A = \frac{1}{2} \ln |I + \lambda^{-1} K_{t-1}^A|$.

Before proving the important Lemma, we need the following lemma to upper bound $\frac{\tau(\frac{\beta_{t_i}^A}{\omega_T})}{\tau(-\frac{\beta_{t_i}^A}{\omega_T})}$ for $1 \leq i \leq n(A)$.

Lemma 20. *Given $\omega_T = \sqrt{\ln(T) \ln \ln(T)}$. Let A be a subset of \mathcal{X} and assume that there exists a $1 \leq \phi \leq T$ such that $A \in \mathcal{A}_\tau$. Let $\phi'(A) = \max\{1 \leq t \leq T : A \in \mathcal{A}_t\}$. There exists a constant $C > 0$ such that for every $T > 0$ and for every $1 \leq i \leq n(A)$ we have*

$$\frac{\tau\left(\frac{\beta_{t_i}^A}{\omega_T}\right)}{\tau\left(-\frac{\beta_{t_i}^A}{\omega_T}\right)} \leq C.$$

Proof. By definition, $\beta_{t_i}^A = B + \sqrt{\gamma_{t_i-1}^A + 1 + \ln(1/\delta)}$. Using the increasing monotonicity of function γ^A , we have that $\gamma_{t_i-1}^A \leq \gamma_{\phi'(A)}^A$ (because $t_i - 1 \leq \phi'(A)$). By Lemma 17, $\gamma_{\phi'(A)}^A \leq C \ln(T) \ln \ln(T)$. Therefore, $\beta_{t_i}^A = B + \sqrt{\gamma_{t_i-1}^A + 1 + \ln(1/\delta)} \leq B + \sqrt{\gamma_{\phi'(A)}^A + 1 + \ln(1/\delta)} \leq B + \sqrt{C \ln(T) \ln \ln(T) + 1 + \ln(1/\delta)}$.

On the other hand, there exists a C_1 large enough such that $B + \sqrt{C \ln(T) \ln \ln(T) + 1 + \ln(1/\delta)} \leq C_1 \sqrt{\ln(T) \ln \ln(T)}$. Combining this with the above result, we imply that $\beta_{t_i}^A \leq C_1 \sqrt{\ln(T) \ln \ln(T)} = C_1 \omega_T$. It is equivalent that $\frac{\beta_{t_i}^A}{\omega_T} \leq C_1$.

Since the function $\tau(z)$ is non-decreasing, we have that $\tau(\frac{\beta_{t_i}^A}{\omega_T}) \leq \tau(C_1)$, and $\tau(-\frac{\beta_{t_i}^A}{\omega_T}) \geq \tau(-C_1)$. Thus,

$$\frac{\tau(\frac{\beta_{t_i}^A}{\omega_T})}{\tau(-\frac{\beta_{t_i}^A}{\omega_T})} \leq \frac{\tau(C_1)}{\tau(-C_1)}.$$

Set $C = \frac{\tau(C_1)}{\tau(-C_1)}$. We now can bound the ratio $\frac{\tau(\frac{\beta_{t_i}^A}{\omega_T})}{\tau(-\frac{\beta_{t_i}^A}{\omega_T})}$ by a constant.

$$\frac{\tau(\frac{\beta_{t_i}^A}{\omega_T})}{\tau(-\frac{\beta_{t_i}^A}{\omega_T})} \leq C.$$

□

Finally, we provide an upper bound for the sum $\sum_{i=1}^{n(A)} r_i^A$.

Lemma 21. *Let $\omega_T = \sqrt{\ln(T) \ln \ln(T)}$. Given a set $A \in \mathcal{A}_t$ and assume that there exists a $1 \leq \phi \leq T$ such that $A \in \mathcal{A}_\tau$. Let $\phi'(A) = \max\{1 \leq t \leq T : A \in \mathcal{A}_t\}$. Then with probability at least $1 - \delta$, we have*

$$\sum_{i=1}^{n(A)} r_i^A \leq \tilde{\mathcal{O}}(\sqrt{n(A) \gamma_{\phi'(A)}^A}),$$

where $n(A) = |\mathcal{D}_T \cap A|$ which is defined as above. The notation $\tilde{\mathcal{O}}$ is a variant of \mathcal{O} , where log factors are suppressed.

Proof. Following the Improved-GP-EI algorithm, each $A \in \mathcal{A}_t$ is fitted by an independent Gaussian process with $N(A)$ observations and the sampled points in $\mathcal{D}_t^A: (x_1^A, y_1^A), \dots, (x_{N(A)}^A, y_{N(A)}^A)$. Therefore, we can apply the results in Section B to the set A .

Similar to the proofs of Lemma 6, Lemma 7, and Lemma 8, we obtain an upper bound on $r_i^A = f(x^*) - f(x_i^{A+})$ as follows.

$$r_i^A \leq \left(\frac{\tau(\frac{\beta_{t_i}^A}{\omega_T})}{\tau(-\frac{\beta_{t_i}^A}{\omega_T})} + \sqrt{2\pi} \right) \max\{0, f(x_{i+1}^A) - \mu_{t_i}^{A+}\} + \left(\frac{\tau(\frac{\beta_{t_i}^A}{\omega_T})}{\tau(-\frac{\beta_{t_i}^A}{\omega_T})} (\beta_{t_i}^A + \omega_T) + \sqrt{2\pi} (3\beta_{t_i}^A + \omega_T) \right) \sigma_{t_i}^A(x_{i+1}^A),$$

where $\mu_{t_i}^{A+} = \max_{x_j^A \in \mathcal{D}_T \cap A, 1 \leq j \leq t_i} \mu_{t_i}(x_j^A)$. Recall that $x_i^{A+} = \operatorname{argmax}_{x_j^A \in \mathcal{D}_T \cap A, 1 \leq j \leq i} f(x_j^A)$.

Using Lemma 21, we can bound the ratio $\frac{\tau(\frac{\beta_{t_i}^A}{\omega_T})}{\tau(-\frac{\beta_{t_i}^A}{\omega_T})}$ by a constant which is independent of $1 \leq t \leq T$ and T . Thus,

by the proofs similar to those of Lemma 14 and Lemma 15, we obtain an upper bound on the sum $\sum_{i=1}^{n(A)} r_i^A$ as

$$\sum_{i=1}^{n(A)} r_i^A \leq \mathcal{O}(\beta_{t_{N(A)}}^A \sqrt{n(A) \gamma_{\phi'(A)}^A}).$$

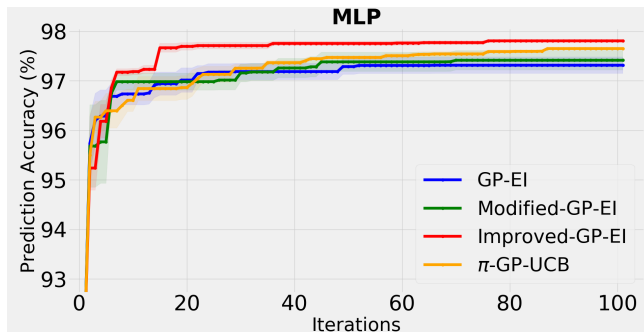


Figure 3: Comparison of methods on MLP.

Using the proof similar to Lemma 21 as above, there exists constant C_1 such that $\beta_{t_{N(A)}}^A \leq C_1 \omega_T = \Theta(\sqrt{\ln(T)\ln\ln(T)})$. We can remove this log component from the upper bound. Finally, we have that

$$\sum_{i=1}^{n(A)} r_i^A \leq \tilde{O}(\sqrt{n(A)\gamma_{\phi'(A)}^A}).$$

□

D Additional Experiments for Real-World Benchmarks

We now take a 2-layer perceptron network (MLP) with 512 neurons/layer and optimize three hyperparameters: the learning rate l and the norm regularization hyperparameters l_{r1} and l_{r2} of the two layers. We train the algorithms using the MNIST train dataset (55000 patterns) and then test the model on the MNIST test dataset (10000 patterns). We plot the optimization results using prediction accuracy in Figure 3.