Regret-Based (ϵ, δ)-optimal Stopping Criteria for Bayesian Optimization

Haowei Wang*
National University of Singapore
Singapore, Singapore
haowei_wang@u.nus.edu

Jingyi Wang*
Lawrence Livermore National Laboratory
CA, USA
wang125@llnl.gov

Qiyu Wei*
The University of Manchester
Manchester, UK
qiyu.wei@postgrad.manchester.ac.uk

Abstract

Bayesian optimization (BO) is a widely used iterative black-box optimization method that utilizes Gaussian process (GP) surrogate models. In practice, BO is typically terminated after a fixed evaluation budget is exhausted, which can incur unnecessary cost and provides no optimality guarantee on solution quality. Recent research in developing a practical stopping criterion has made empirical progress, yet a theoretically sound stopping criterion remains a work in progress. In this work, we present provably tighter instantaneous regret bounds for GP upper confidence bound (GP-UCB) at any given iteration. Then, we propose stopping criteria for GP-UCB based on this tighter bound that ensure an ϵ-optimal solution with high probability $1 - \delta$ upon termination. Numerical experiments are performed to validate and demonstrate the effectiveness and efficiency of our stopping criteria.

1 Introduction

Bayesian optimization (BO) has emerged as a powerful framework for the global optimization of expensive black-box functions. In the classic form, it aims to solve the optimization problem

$$\underset{\boldsymbol{x} \in D}{\text{maximize}} \quad f(\boldsymbol{x}), \tag{1}$$

where $D \subset \mathbb{R}^d$. By combining a probabilistic surrogate model, most commonly a Gaussian process (GP), with an acquisition function that trades off exploration and exploitation to select observation points, BO selects evaluation points iteratively in order to locate the global optimum of the unknown objective function. This iterative paradigm has enabled BO to achieve satisfactory performance in various applications where function evaluation is expensive or time-consuming, including engineering design [Lam, 2018, Wang et al., 2024], simulation-based optimization [Tay and Osorio, 2022, Wang and Papadopoulos, 2025], machine learning hyperparameter tuning [Snoek et al., 2012], and materials discovery [Frazier and Wang, 2015, Jin and Kumar, 2023]. At iteration t, BO updates the GP posterior with past observations, maximizes the acquisition function to choose the next query, observes new data, and repeats. This sequential optimization loop structure immediately raises a fundamental algorithmic question: when should this iterative procedure stop?

Despite its centrality, the question of establishing principled termination criteria for BO has received limited theoretical attention. In practice, the prevailing approach is to prescribe a fixed budget

*Equal contribution; authors listed alphabetically.

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

in advance. This practice is ubiquitous across software packages and empirical studies [Balandat et al., 2020, Lindauer et al., 2022], and is often chosen for convenience rather than algorithmic necessity. However, such total budget-based stopping rules are inherently ad hoc: they may lead to substantial inefficiencies, and more importantly they do not reflect the proximity to optimality, *i.e.*, a sufficiently accurate global optimal solution is not guaranteed. This is especially problematic in settings where the underlying function evaluations arise from time-consuming simulations or costly physical experiments, for which exhausting a preallocated evaluation budget may be wasteful or impractical. This has motivated growing interest in designing BO stopping rules with optimality guarantees at termination [Wilson, 2024, Xie et al., 2025].

We seek stopping rules that are both theoretically grounded and practically effective. Concretely, a desirable stopping rule should i) guarantee convergence to the global optimizer under some standard regularity assumptions, and ii) yield substantial savings in evaluation cost relative to budget-exhaustion strategies. However, existing methods fall short of satisfying these goals. Existing stopping rules in BO often are either heuristic or restrictive. Heuristic approaches, such as using Monte Carlo estimates of simple regret [Wilson, 2024], and its robustness improvement [He et al., 2025], confidence-bound gaps [Makarova et al., 2022], or criteria based on Euclidean distance, probability of improvement, or EI values [Lorenz et al., 2016, Nguyen et al., 2017], offer practical guidance but lack theoretical guarantees. Approaches with theoretical properties address a narrower stopping problem, add assumptions, or certify a weaker target: Bayesian optimal stopping [Dai et al., 2019] terminates the training epochs of each evaluated configuration rather than the BO loop and assumes that performance improves monotonically over epochs; the gap of expected minimum simple regrets [Ishibashi et al., 2023] relies on additional stability assumptions and may stop while the simple regret remains large; and convexity-based termination [McLeod et al., 2018, Li et al., 2023] identifies local convex structure and either switches to local optimization or certifies only local regret. Our criterion does not require these assumptions. In addition, some of these methods can incur considerable additional computational overhead. Wilson [2024] proposed the Probabilistic Regret Bound (PRB) criterion, which estimates via Monte Carlo sampling from the GP posterior. It relies on numerical estimation subject to sampling and optimization errors. Building on PRB, He et al. [2025] proposed MSP, which improves robustness to mean-model misspecification by sampling mean parameters from their posterior.

This gap motivates the present work. To establish guaranteed optimality upon termination, the key idea is to better utilize the theoretical regret upper bounds. Since the seminal work of Srinivas et al. [2009], there have been abundant studies on the regret bounds of GP-UCB, proving their sublinear cumulative regrets. However, the vast majority of the works have focused on the order of the regret upper bound in terms of iteration/sample number $t \in \mathbb{N}$ as $t \to \infty$. In contrast, this paper focuses on developing effective and efficient stopping criteria based on a tight instantaneous regret bound at a finite time T that ensures an ϵ optimal solution with probability greater than $1 - \delta$.

We evaluate the proposed criteria on both synthetic benchmarks and real-world optimization tasks. We compare the proposed stopping rules against widely used heuristics, uncertainty-based rules, and recent state-of-the-art (SOTA) baselines such as PRB [Wilson, 2024]. Empirically, our criterion consistently reduces evaluation iterations while maintaining or improving the quality of the returned solutions. The observed stopping behavior closely aligns with the established theoretical findings and incurs incremental computational cost. Our contributions can be summarized as follows.

- We propose theoretically grounded stopping criteria applicable to BO algorithms that admit subsequences of converging regret bounds including sublinear cumulative regret guarantees, e.g., GP-UCB.
- We establish provably tighter instantaneous regret bounds for GP-UCB at given iteration t. Based on this new bound, we present explicit stopping criteria for GP-UCB.
- We perform extensive experiments on both synthetic and real-world tasks. Results show that our criteria can save evaluations while achieving superior or comparable optimization performance to strong baselines.

2 Background

We consider the maximization problem (1). At iteration t, we query $\boldsymbol{x}_t \in D$ and observe y_t. BO algorithms consist of two major components: a GP surrogate model for the objective function f, and

an acquisition function that measures the benefit of the next sample and guides the sample selection. We introduce both in the subsequent subsections.

2.1 Gaussian Processes

A GP prior assumes that any finite collection of function values follows a multivariate Gaussian distribution. Consider a zero mean GP with the kernel $k(\boldsymbol{x}, \boldsymbol{x}') : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R}$, denoted as $GP(0, k(\boldsymbol{x}, \boldsymbol{x}'))$. At each sample $\boldsymbol{x}_i \in D$, $i = 1, 2, \ldots$, we consider observation noise ϵ_i for f so that the observed function value is $y_i = f(\boldsymbol{x}_i) + \epsilon_i$. The noise is assumed to follow an independent and identically distributed (i.i.d.) Gaussian distribution $\mathcal{N}(0, \sigma^2)$, where $\sigma > 0$ is the noise standard deviation (see Section 2.3 for formal assumptions).

Denote the objective vector and the observation vector as $f_{1:t} = [f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_t)]^T$ and $\boldsymbol{y}_{1:t} = [y_1, \ldots, y_t]^T$, respectively. The posterior distribution of $f(\boldsymbol{x})|\boldsymbol{x}_{1:t}, \boldsymbol{y}_{1:t} \sim \mathcal{N}(\mu_t(\boldsymbol{x}), \sigma_t^2(\boldsymbol{x}))$, where $\mathcal{N}$ denotes the normal distribution, can be inferred using Bayes' rule with $\mu_t(\boldsymbol{x}) = \boldsymbol{k}_t(\boldsymbol{x})^\top (\boldsymbol{K}_t + \sigma^2 \boldsymbol{I})^{-1} \boldsymbol{y}_{1:t}$ and $\sigma_t^2(\boldsymbol{x}) = k(\boldsymbol{x}, \boldsymbol{x}) - \boldsymbol{k}_t(\boldsymbol{x})^\top (\boldsymbol{K}_t + \sigma^2 \boldsymbol{I})^{-1} \boldsymbol{k}_t(\boldsymbol{x})$, where $\boldsymbol{k}_t(\boldsymbol{x}) = [k(\boldsymbol{x}_1, \boldsymbol{x}), \ldots, k(\boldsymbol{x}_t, \boldsymbol{x})]^T$ and $\boldsymbol{K}_t = \left[k(\boldsymbol{x}_i, \boldsymbol{x}_j)\right]_{i,j=1}^t$.

Two of the most popular choices of kernels are the SE and Matérn kernels [Frazier, 2018], given by $k_{SE}(\boldsymbol{x}, \boldsymbol{x}') = \exp(-p^2/(2l^2))$ and $k_{Matérn}(\boldsymbol{x}, \boldsymbol{x}') = \frac{1}{\Gamma(\nu)\, 2^{\nu-1}} (\sqrt{2\nu}\, p/l)^\nu K_\nu (\sqrt{2\nu}\, p/l)$, where $l > 0$ is a length scale hyperparameter, $p = \|\boldsymbol{x} - \boldsymbol{x}'\|_2$, $\nu > 0$ is the smoothness parameter of the Matérn kernel, and K_ν is the modified Bessel function of the second kind.

2.2 Acquisition Function

Denote the acquisition function given $t-1$ samples as $\alpha_{t-1}(\boldsymbol{x})$. The next sample $\boldsymbol{x}_t$ in a typical BO algorithm is chosen by maximizing the acquisition function over set D_t, *i.e.*,

$$\boldsymbol{x}_t = \arg\max_{\boldsymbol{x} \in D_t} \alpha_{t-1}(\boldsymbol{x}).$$

We use D_t for its generality. For GP-UCB, D_t is simply D and thus is independent of t. For GP-TS however, D_t can be a subset of D and viewed as an algorithmic parameter. To solve this maximization, optimization algorithms such as L-BFGS can be used.

The choice of the acquisition function differentiates BO algorithms and has a significant impact on its performance. The definition of UCB for the maximization problem is

$$UCB_t(\boldsymbol{x}) = \mu_{t-1}(\boldsymbol{x}) + \beta_t^{1/2} \sigma_{t-1}(\boldsymbol{x}),$$

where $\beta_t \geq 0$ is an algorithmic parameter crucial to the performance of GP-UCB. It is often chosen based on theoretical cumulative regret results [Srinivas et al., 2009, Chowdhury and Gopalan, 2017]. Naturally, β_t plays an important role in the proposed stopping criteria for GP-UCB.

2.3 Analytic Settings

Let $\boldsymbol{x}^* \in \arg\max_{\boldsymbol{x} \in D} f(\boldsymbol{x})$ and $f^* := f(\boldsymbol{x}^*)$. An output $\hat{\boldsymbol{x}}$ is said to be an (ϵ, δ)-optimal solution of the maximization problem if for given $\epsilon > 0$ and $\delta \in (0, 1)$, $\mathbb{P}(f^* - f(\hat{\boldsymbol{x}}) \leq \epsilon) \geq 1 - \delta$. Our stopping criteria in Section 3 are designed to guarantee this (ϵ, δ)-optimality upon termination.

We consider the optimization history next. Let $\mathcal{H}_{t-1} = \{(\boldsymbol{x}_i, y_i)\}_{i=1}^{t-1}$ be the history and let $\mathcal{F}_{t-1} := \sigma(\mathcal{H}_{t-1})$ denote the filtration generated by the optimization history up to iteration $t-1$. Given $\mathcal{H}_{t-1}$, the posterior mean μ_{t-1} and variance σ_{t-1}^2 are deterministic. We use the standard notations $\mathcal{O}$ and o to denote the orders of quantities. Further, we use $\tilde{\mathcal{O}}$ to suppress the logarithm terms of $\mathcal{O}$.

We present our analytic assumptions below.

Assumption 2.1. The constraint set $D \subseteq [0, r]^d$ is convex and compact for $r > 0$. The objective f is a sample from the Gaussian process $GP(0, k(\boldsymbol{x}, \boldsymbol{x}'))$, where $k(\boldsymbol{x}, \boldsymbol{x}') \leq 1$ and $k(\boldsymbol{x}, \boldsymbol{x}) = 1$. There exist constants $a, b > 0$ such that the objective function f is Lipschitz continuous with Lipschitz constant L (in 1-norm) with probability $\geq 1 - dae^{-L^2/b^2}$. The observation noise is i.i.d., *i.e.*, $\epsilon_t \sim \mathcal{N}(0, \sigma^2), \sigma > 0$ for all $t \in \mathbb{N}$.

3 Stopping Criteria

In this section, we propose our stopping criteria based on regret bounds. Recall that the instantaneous regret is defined as $r_t = f^* - f(\boldsymbol{x}_t)$. From an optimization perspective, it is natural to prefer BO algorithms with at least a converging subsequence of instantaneous regret, *i.e.*, $r_{t_i} \to 0$ for $t_i \to \infty$. Such a property would guarantee that an algorithm converges to the optimal solution f^* at some iterates. In the bandit setting, no-regret is a highly desirable property since it ensures overall convergence performance. A BO algorithm is called no-regret if its cumulative regret is sublinear, *i.e.*, $\lim_{T\to\infty} R_T/T = 0$. Since R_T is the sum of r_t, no regret guarantees a converging subsequence of instantaneous regret.

Consider a BO algorithm that satisfies the following instantaneous regret bound

$$\mathbb{P}\{r_t \leq c_t(\boldsymbol{x}_{1:t-1}, \boldsymbol{y}_{1:t-1}, t, \delta), \text{for all } t \in \mathbb{N}\} \geq 1 - \delta. \tag{2}$$

The function c_t is a t-dependent upper bound on r_t. By our formulation, c_t can be computed using existing information accumulated up to iteration $t-1$, *i.e.*, $\mathcal{F}_{t-1}$. A well-known example of c_t is $c_t(\boldsymbol{x}_{1:t-1}, \boldsymbol{y}_{1:t-1}, t, \delta) = 2\beta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t)$ for GP-UCB [Srinivas et al., 2009].

Given a desired tolerance $\epsilon > 0$, our stopping criterion terminates the algorithm if

$$c_t(\boldsymbol{x}_{1:t-1}, \boldsymbol{y}_{1:t-1}, t, \delta) < \epsilon. \tag{3}$$

The stopping test (3) uses the optimization trajectory and posterior uncertainty quantities available and introduces no additional black-box evaluations beyond those at sample points $\boldsymbol{x}_i, i = 1, \ldots, t-1$. This makes it simple to implement and computationally efficient. We note that an upper bound c_t that relies on the observation y_t can also be feasible.

With the proposed stopping criteria, the generic BO algorithm is given below in Algorithm 1.

Algorithm 1 BO with stopping criteria

Choose the kernel $k(\cdot,\cdot)$ and parameter σ. Choose t_0 initial samples $\boldsymbol{x}_i$ with observations $y_i, i = 0, \ldots, t_0 - 1$.
Train the GP surrogate model for f on the initial samples.
for $t = t_0, t_0 + 1, \ldots$ **do**
    Find $\boldsymbol{x}_t = \arg\max_{\boldsymbol{x}\in D_t} \alpha_{t-1}(\boldsymbol{x})$.
    Observe $y_t = f(\boldsymbol{x}_t) + \epsilon_t$.
    Update the surrogate model with the addition of $\boldsymbol{x}_t$ and y_t.
    Compute the upper bound c_t.
    if (3) is satisfied **then**
        Exit
    end if
end for

Clearly, if $c_t(\boldsymbol{x}_{1:t-1}, \boldsymbol{y}_{1:t-1}, t, \delta)$ in (2) converges, (3) provides the theoretical optimality guarantee, *i.e.*, upon exit at iteration t, $\boldsymbol{x}_t$ is the (ϵ, δ)-optimal solution. We formalize the optimality result below.

Lemma 3.1. *If a BO algorithm generates a converging (to* 0*) subsequence of the instantaneous regret bound c_t as stated in* (2)*, i.e.,* $\liminf_{t\to\infty} c_t = 0$*, then for any $\epsilon > 0$ and $\delta \in (0, 1)$, Algorithm 1 exits in a finite number of iterations $\bar{t}$. Further, $\boldsymbol{x}_{\bar{t}}$ is an (ϵ, δ)-optimal solution of* (1)*.*

Lemma 3.1 asserts the well-posedness of the proposed stopping criteria in that (3) is satisfied for large enough t. Since GP-UCB has been shown to generate sublinear cumulative regert bound in our setting (Assumption 2.1) for commonly used kernels such as the SE and Matérn kernels [Srinivas et al., 2009, Whitehouse et al., 2023], it satisfies Lemma 3.1. Hence, a c_t that follows the actual regret r_t closely can lead to a robust, efficient, and theoretically sound stopping criterion for GP-UCB.

3.1 Tightened Instantaneous Regret Bounds for GP-UCB

In this section, we present tightened instantaneous regret bounds for GP-UCB and the corresponding stopping criteria (3). Importantly, our focus is on improving the finite-time upper bound at a given

iteration t, rather than on the asymptotic order in t as $t \to \infty$, since the tightness of the bound directly impacts the practical performance of the stopping criteria.

To ensure clear presentation and fair comparison, we keep the standard GP-UCB in Section 2.2 and choose the widely adopted choice of β_t in literature [Srinivas et al., 2009]

$$\beta_t = 2\log(4\pi_t/\delta) + 2d\log(dt^2br\sqrt{\log(4da/\delta)}), \tag{4}$$

where $\pi_t = \pi^2 t^2/6$. We note that the discretization size, *i.e.*, number of discrete samples chosen for analytic purposes in (4) is greater than or equal to 1. Thus, implicitly, we have $dt^2br\sqrt{\log(2da/\delta)} \geq 1$. The instantaneous regret bound in Srinivas et al. [2009] is stated next.

Lemma 3.2. *Pick $\delta \in (0,1)$ and set β_t as in* (4)*. Under Assumption 2.1, the instantaneous regret of GP-UCB satisfies*

$$r_t \leq 2\beta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) + \frac{1}{t^2}, \tag{5}$$

for all $t \in \mathbb{N}$, with probability $\geq 1-\delta$.

To proceed, we define four analytic parameters

$$n_{\eta,t}, n_{\alpha,t} > 1, s_t > 0, n_L > 1. \tag{6}$$

We emphasize that these four parameters are not deployed in the UCB acquisition function. The choices of their values will be given later in this section. The first two parameters $n_{\eta,t}$ and $n_{\alpha,t}$ define their corresponding confidence interval parameters η_t and α_t via

$$\pi_t n_{\eta,t}(1-\Phi(\eta_t^{1/2})) = 1,\ \pi_t n_{\alpha,t}|\mathbb{D}_t|(1-\Phi(\alpha_t^{1/2})) = 1, \tag{7}$$

where $\mathbb{D}_t$ is a time-dependent discretization defined below. The parameters η_t and α_t serve similar purposes as β_t and are used to obtain tighter confidence intervals in our analysis.

The last two parameters s_t and n_L parameterize the construction of the discretization [Srinivas et al., 2009]. Specifically, we can choose discretization $\mathbb{D}_t$ of D with cardinality

$$|\mathbb{D}_t| = \left\lceil (rdb)^d \log(n_L da/\delta)^{d/2}\, t^{s_t d} \right\rceil, \tag{8}$$

where the ceiling function $\lceil\cdot\rceil$ denotes the smallest integer not less than its argument. Denote the closest point at $\boldsymbol{x} \in D$ to $\mathbb{D}_t$ as $[\boldsymbol{x}]_t$. Then, we can bound $|f(\boldsymbol{x}) - f([\boldsymbol{x}]_t)|$ for $\forall \boldsymbol{x} \in D$ (Lemma B.8). We recover the widely used discretization with β_t that ensures $|f(\boldsymbol{x}) - f([\boldsymbol{x}]_t)| \leq \frac{1}{t^2}$ if we set $s_t = 2$ and $n_L = 4$. A similar parameterization technique was used in [Wang et al., 2025].

Next, we present the instantaneous regret bound that includes these four parameters.

Lemma 3.3. *Pick $\delta \in (0,1)$, set β_t as in* (4)*, and let $\alpha_t \leq \beta_t$. Under Assumption 2.1, at any given iteration t, the instantaneous regret bound of GP-UCB satisfies*

$$r_t \leq (\beta_t^{1/2} + \eta_t^{1/2})\sigma_{t-1}(\boldsymbol{x}_t) - (\beta_t^{1/2} - \alpha_t^{1/2})\sigma\sqrt{\frac{1}{t+\sigma^2}} + \frac{1}{t^{s_t}}, \tag{9}$$

with probability $\geq 1-\delta/n_L - \frac{1}{\pi_t}(1/n_{\eta,t} + 1/n_{\alpha,t})$, where $|\mathbb{D}_t|$ satisfies (8)*, $\pi_t n_{\eta,t}(1-\Phi(\eta_t^{1/2})) = 1$, and $\pi_t n_{\alpha,t}|\mathbb{D}_t|(1-\Phi(\alpha_t^{1/2})) = 1$.*

The full proof of Lemma 3.3 is given in Appendix B. To establish (9) at a given t and δ, we make multiple necessary modifications to the derivation of instantaneous regret upper bound. For a clear presentation, we list the key proof techniques below to highlight the changes.

Summary of Proof Techniques 1. Given $\mathcal{F}_{t-1}$, we use a different parameter η_t for the confidence interval at $\boldsymbol{x}_t$, whose values can differ from β_t. Similarly, we introduce a new parameter α_t for the confidence interval at $[\boldsymbol{x}^*]_t$, independent of the value of β_t. In addition, we do not take the union bound over $\mathbb{D}_t$ for the confidence interval at $\boldsymbol{x}_t$ with η_t, thereby reducing η_t for the same probability level.
2. We adopt a parameterized discretization $\mathbb{D}_t$ with s_t and n_L. This allows us to design and control the probability of the confidence interval at $[\boldsymbol{x}^*]_t$, since a union bound over all points in $\mathbb{D}_t$ is needed. Moreover, n_L allocates the probability budget of the Lipschitz continuity of f.

3. Minor adjustment: we do not relax the probability of $r > c > 0$ where $r \sim \mathcal{N}(0,1)$ from $1 - \Phi(c)$ to $\frac{1}{2}e^{-c^2/2}$, as in Srinivas et al. [2009]. As c increases, the ratio $\frac{1-\Phi(c)}{0.5e^{-c^2/2}}$ decreases. Moreover, we only use one-sided confidence intervals to obtain an upper bound on r_t, which increases the probability of the same confidence intervals holding. Finally, we incorporate the global lower bound of $\sigma_{t-1}(\boldsymbol{x})$ for any $t \in \mathbb{N}$ that was previously also used in Wang and de Freitas [2014].

Sketch of Proof From the definition of the instantaneous regret bound, we first decompose r_t into $f(\boldsymbol{x}^*) - f([\boldsymbol{x}^*]_t) + f([\boldsymbol{x}^*]_t) - f(\boldsymbol{x}_t)$, where $[\boldsymbol{x}]_t$ denotes the closest point in $\mathbb{D}_t$ to $\boldsymbol{x}$. Then, we bound $f(\boldsymbol{x}^*) - f([\boldsymbol{x}^*]_t)$ through the Lipschitz continuity of f with high probability and the discretization $\mathbb{D}_t$. Next, we use the confidence interval at $[\boldsymbol{x}^*]_t$ and $\boldsymbol{x}_t$ with parameters $\alpha_t^{1/2}$ (from $n_{\alpha,t}$) and $\eta_t^{1/2}$ (from $n_{\eta,t}$), respectively to replace $f([\boldsymbol{x}^*]_t) - f(\boldsymbol{x}_t)$ with terms $\mu_{t-1}([\boldsymbol{x}^*]_t)$, $\mu_{t-1}(\boldsymbol{x}_t)$, $\sigma_{t-1}([\boldsymbol{x}^*]_t)$, and $\sigma_{t-1}(\boldsymbol{x}_t)$. Given our condition $\beta_t \geq \alpha_t$, we use the global lower bound on $\sigma_{t-1}([\boldsymbol{x}^*]_t)$ in Lemma B.4 to arrive at (9).

Optimization Subproblem Formulation The instantaneous regret bound in Lemma 3.3 depends on the parameters. To obtain the smallest upper bound, we solve an optimization subproblem. We note that at iteration t given $\mathcal{F}_{t-1}$, GP-UCB deterministically chooses $\boldsymbol{x}_t$. Thus, $\sigma_{t-1}(\boldsymbol{x}_t)$ and similarly $\sigma\sqrt{\frac{1}{t+\sigma^2}}$ are deterministic given $\mathcal{F}_{t-1}$. Define

$$c_1(t) = \sigma_{t-1}(\boldsymbol{x}_t) \quad \text{and} \quad c_2(t) = \sigma\sqrt{\frac{1}{t+\sigma^2}}.$$

For a given $\delta \in (0,1)$ and t, the optimization subproblem is

$$\begin{aligned}
\min_{n_{\eta,t}, n_{\alpha,t}, s_t \in \mathbb{R}} \quad & \eta_t^{1/2} c_1(t) + \alpha_t^{1/2} c_2(t) + \frac{1}{t^{s_t}} \\
\text{s.t.} \quad & 1/n_{\eta,t} + 1/n_{\alpha,t} \leq \delta(1 - 1/n_L), \\
& \pi_t n_{\eta,t}(1 - \Phi(\eta_t^{1/2})) = 1,\ \pi_t n_{\alpha,t}|\mathbb{D}_t|(1 - \Phi(\alpha_t^{1/2})) = 1, \\
& |\mathbb{D}_t| = (rdb)^d \log(n_L da/\delta)^{d/2}\, t^{s_t d} + 1, \\
& \alpha_t, \eta_t, s_t > 0, \\
& \alpha_t^{1/2} \leq \beta_t^{1/2}, \\
& n_{\eta,t}, n_{\alpha,t} > 1.
\end{aligned} \tag{10}$$

The objective function in (10) is part of the upper bound in (9) after dropping constants independent of the decision variables. The first constraint in (10) allocates the probability budget between the confidence intervals at $\boldsymbol{x}_t$ and $[\boldsymbol{x}^*]_t$. Combined with the weighting sequence π_t, this ensures that the instantaneous regret upper bound holds with probability at least $1 - \delta$ simultaneously for all $t \in \mathbb{N}$. The third constraint is from (8) and the $+1$ term ensures the ceiling function is satisfied. We note that $|\mathbb{D}_t|$ can be relaxed to be a continuous variable in the optimization subproblem. The constraint $\alpha_t \leq \beta_t$ is required to use the lower bound of $\sigma_{t-1}([\boldsymbol{x}^*]_t)$. Finally, the last line of constraints in (10) is explicitly written out for clarity of presentation, but is implicitly enforced by the probability equations.

Problem (10) is a simple three-dimensional ($n_{\eta,t}, n_{\alpha,t}, s_t$) smooth, *i.e.*, continuously differentiable, nonlinear optimization problem. Thus, it can be efficiently solved using standard optimization software such as IPOPT [Wächter and Biegler, 2006] or SciPy's constrained optimization routines [Virtanen et al., 2020], and should only add negligible computational cost. To ensure the robustness of the optimization solver, we relax the lower box constraint on the optimization variables to

$$\alpha_t^{1/2}, \eta_t^{1/2}, s_t \geq \epsilon_b, \qquad n_{\eta,t}, n_{\alpha,t} \geq 1 + \epsilon_b, \tag{11}$$

where $\epsilon_b > 0$ is a small lower-bound tolerance.

Remark 3.4 (Well-posedness of the subproblem). Using the relaxation in (11), which is widely adopted in modern nonlinear optimization algorithms [Wächter and Biegler, 2006], the feasible region of (10) becomes closed. Moreover, the constraints imply that the feasible variables $n_{\alpha,t}$ and s_t are bounded, while $n_{\eta,t}$ is effectively bounded from the objective and (11). Hence, the nonempty feasible level set of (10) is compact. We note that the existing bound (5) is feasible for (10) when

$n_L = 4$ (given that $dt^2br\sqrt{\log(2da/\delta)} \geq 1$ as mentioned above). Therefore, by the continuity of the objective function, a global minimizer of (10) exists [Nocedal and Wright, 2006]. Further, since the conventional β_t (4) already incorporates an equivalent choice of $n_L = 4$, the formula (4) might need to be updated in accordance with n_L to ensure a global optimum.

Remark 3.5 (Choice of n_L). We do not choose to include n_L as an optimization variable for simplicity of the subproblem. Since n_L controls the probability of the Lipschitz continuity of f, it does not require a union bound to hold true for all t. If we let n_L be time-dependent, denoted by $n_{L,t}$, then the optimization subproblem needs to be modified accordingly and could include more constraints. In practice, we recommend a large $n_L \geq 10$ that also ensures $n_L da/\delta > 1$ since its contribution is inside a $\log$ function, whose effect on the bound is mild, for low-dimensional problems.

Denote the solution to (10) as $\bar{n}_{\alpha,t}, \bar{n}_{\eta,t}$, and $\bar{s}_t$. We can now state the tightened instantaneous regret bound. The proof is presented in Appendix B.

Theorem 3.6. *Under Assumption 2.1, let $\delta \in (0,1)$ and define β_t by* (4)*. Then, with probability $\geq 1-\delta$, for all $t \in \mathbb{N}$, the instantaneous regret of GP-UCB satisfies*

$$r_t \leq (\beta_t^{1/2} + \bar{\eta}_t^{1/2})\sigma_{t-1}(\boldsymbol{x}_t) - (\beta_t^{1/2} - \bar{\alpha}_t^{1/2})\sigma\sqrt{\frac{1}{t+\sigma^2}} + \frac{1}{t^{\bar{s}_t}}, \tag{12}$$

where $\pi_t \bar{n}_{\eta,t}(1-\Phi(\bar{\eta}_t^{1/2})) = 1$, $\pi_t \bar{n}_{\alpha,t}|\mathbb{D}_t|(1-\Phi(\bar{\alpha}_t^{1/2})) = 1$, and $\bar{n}_{\eta,t}, \bar{n}_{\alpha,t}, \bar{s}_t$ are the solutions to (10)*.*

For a given t and δ, the new bound is guaranteed to be no worse than (5) since the existing bound (5) is a feasible point of the optimization subproblem and can be recovered by setting $\alpha_t = \eta_t = \beta_t$, $s_t = 2$ and $n_L = 4$. The stopping criterion (3) for GP-UCB is

$$(\beta_t^{1/2} + \bar{\eta}_t^{1/2})\sigma_{t-1}(\boldsymbol{x}_t) - (\beta_t^{1/2} - \bar{\alpha}_t^{1/2})\sigma\sqrt{\frac{1}{t+\sigma^2}} + \frac{1}{t^{\bar{s}_t}} \leq \epsilon. \tag{13}$$

We emphasize that evaluating the stopping criteria does not alter the GP-UCB optimization path prior to termination. Practitioners may choose both the frequency and the iteration indices at which the stopping test (13) is performed (and hence when the optimization subproblem is solved). The computational cost induced by the stopping test is thus negligible.

Remark 3.7 (Predetermined tests). If the stopping test is performed only at a finite set of predetermined iterations, the instantaneous regret bound can be further tightened, since a union bound over all $t \in \mathbb{N}$ is no longer required. In particular, when the test is performed at a single fixed iteration, no union bound over t is needed. Moreover, the improved bound can be evaluated post hoc at the end of an optimization run to obtain a tighter estimate of the instantaneous regret.

Remark 3.8 (Implications on the choice of β_t.). Although we do not modify the GP-UCB algorithm or the commonly used choices of β_t, our analysis suggests that smaller values of β_t may be admissible. In particular, proof techniques 2 and 3 can be incorporated into the design of β_t to obtain tighter confidence intervals for a given finite t (but not in order of t). Any modification to β_t changes the exploration-exploitation trade-off in the GP-UCB acquisition function and therefore alters the algorithmic trajectory. However, since the order of the cumulative regret bound as $t \to \infty$ is not improved, such modifications are unlikely to positively impact asymptotic performance. In contrast, for moderate values of t, the choice of β_t can have a more pronounced effect on performance due to its influence on exploration. We leave a systematic investigation of this question for future work.

4 Experiments

We first numerically illustrate the tightness of our new instantaneous regret bound, then evaluate the proposed regret-bound-based stopping criteria along two desiderata: validity (returning an (ϵ,δ)-optimal solution) and efficiency (saving evaluations). We additionally report the per-iteration computational overhead of each criterion.

4.1 Effect of the New Instantaneous Regret Bound

To quantify the improvement of our new bound (12) over the existing bound (5), we evaluate both at a given iteration t with fixed problem constants $a = b = r = 1$. Table 1 reports a representative

Table 1: Comparison between the existing bound (5) and our new bound (12). Fixed parameters: $\sigma = 10^{-2}$, $n_L = 4$, $a = b = r = 1$, $d = 4$, $\delta = 0.1$, and $\pi_t = \pi^2 t^2/6$.

t	$c_1(t)$	$c_2(t)$	$\sqrt{\beta_t}$	$\sqrt{\eta_t}$	$\sqrt{\alpha_t}$	$\bar{s}_t$	$\bar{n}_{\eta,t}$	$\bar{n}_{\alpha,t}$	(12)	(5)	Ratio
20	2.24×10^{-3}	2.24×10^{-3}	9.267	3.789	9.267	2.264	20.13	39.51	3.03×10^{-2}	4.39×10^{-2}	0.690
20	1.00×10^{-2}	2.24×10^{-3}	9.267	3.720	9.267	2.182	15.25	106.25	1.31×10^{-1}	1.88×10^{-1}	0.699
20	1.00×10^{-1}	2.24×10^{-3}	9.267	3.691	9.267	2.035	13.63	616.21	1.30	1.86	0.699
100	1.00×10^{-3}	1.00×10^{-3}	10.866	4.519	10.035	1.702	19.55	41.94	1.49×10^{-2}	2.18×10^{-2}	0.685
100	1.00×10^{-2}	1.00×10^{-3}	10.866	4.447	10.238	1.706	13.93	309.46	1.53×10^{-1}	2.17×10^{-1}	0.703
100	1.00×10^{-1}	1.00×10^{-3}	10.866	4.438	10.465	1.711	13.39	3044.46	1.53	2.17	0.704
500	4.47×10^{-4}	4.47×10^{-4}	12.258	5.162	10.670	1.400	19.95	40.19	7.25×10^{-3}	1.10×10^{-2}	0.661
500	1.00×10^{-2}	4.47×10^{-4}	12.258	5.091	10.933	1.404	13.62	636.44	1.73×10^{-1}	2.45×10^{-1}	0.706
500	1.00×10^{-1}	4.47×10^{-4}	12.258	5.087	11.147	1.407	13.36	6368.59	1.73	2.45	0.707

configuration ($d = 4$, $\delta = 0.1$, $n_L = 4$, $\sigma = 10^{-2}$); our bound achieves about 30% reduction (Ratio column) across three iteration counts $t \in \{20, 100, 500\}$ and three values of the posterior standard deviation $c_1(t) = \sigma_{t-1}(\boldsymbol{x}_t)$. Comparisons under additional values of δ, d, and n_L are provided in Appendix B.

4.2 Implementation Details

We use a GP surrogate with Matérn-$5/2$ and SE kernels, and use GP-UCB as the acquisition function for baselines with the same parameter β_t. Each configuration is evaluated over 50 independent random seeds. Each run starts from 5 initial points generated by a Sobol sequence, followed by sequential BO iterations up to the budget T of the task. To evaluate general performance, we did not perform post-hoc tuning of the GP parameters or ϵ, and set $\epsilon = 0.1$, $\delta = 0.05$, and $a = b = 1$ uniformly for all synthetic benchmark tasks. Here a, b are the constants of Assumption 2.1 entering β_t; sensitivity to a, b, δ, and ϵ is reported in Appendix C.1. We report the mean stopping iteration and final simple regret over random seeds. We report the success rate as the fraction of runs whose final simple regret at stopping is $\leq \epsilon$. Codes are available in `https://github.com/Qiyu-Wei/Regret_Based_Bo_Stopping.git`.

4.3 Benchmarks

We consider three families of problems. (a) Synthetic functions: We use classical continuous benchmarks with known optima: Branin ($d = 2$, $T = 128$), Rosenbrock ($d = 4$, $T = 96$), Levy ($d = 4$, $T = 96$), and Hartmann ($d = 3$, $T = 64$). (b) GP-sampled objectives: Following Wilson [2024], we additionally evaluate on objectives drawn from GP priors with known optima. We consider two settings, $d = 2$ with $T = 128$ and $d = 6$ with $T = 512$, both at two observation noise levels $\sigma \in \{3 \times 10^{-3}, 5 \times 10^{-3}\}$. For each sampled function, the optimum is computed offline by the benchmark generation procedure and used as ground truth for regret evaluation. We report two variants: GP† and GP, as in Wilson [2024]. In GP†, the BO surrogate is well-specified and uses the true hyperparameters of the GP data-generating process provided by an oracle. In GP, the surrogate hyperparameters are unknown and are estimated online via MAP. (c) Real-world hyperparameter optimization: We tune a random-forest classifier on the digits dataset (RFHPO, $d = 5$, $T = 128$) over the number of trees, the maximum depth, the minimum samples per split and per leaf, and the fraction of features, with a random $80/20$ train/validation split; and an RBF-kernel SVM on Fashion-MNIST (FMNIST-SVM, $d = 2$, $T = 64$) over $\log_{10} C \in [-2, 3]$ and $\log_{10} \gamma \in [-4, 0]$, with $2{,}000$ training and $2{,}000$ validation images. Following Wilson [2024], we set $\epsilon = 0.03$ and 0.165 for these tasks. Consistent with the maximization framework (1), we maximize the negative validation error.

We compare against representative stopping criteria from the literature: PRB [Wilson, 2024], ΔCB [Makarova et al., 2022] and ΔES [Ishibashi et al., 2023], and Acq [Wilson, 2024]. We also include two reference baselines: NOSTOP, which runs to the full budget T without early stopping, and Oracle$_r$, which post-processes each run to determine the earliest iteration at which the regret satisfies the target tolerance; it stops at the first t such that $f^* - \max_{i \leq t} f(\boldsymbol{x}_i) \leq \epsilon$, providing a lower bound on the stopping iteration achievable by any rule. Our method is denoted UCB$_{\mathrm{br}}$, corresponding to the Bayesian stopping rule of Section 3.1. The subscript 'b' indicates the Bayesian analytic setting and 'r' stands for regret.

Table 2: Comparison of stopping criteria under the Matérn-$5/2$ kernel across GP-sampled functions († means known hyperparameters), synthetic benchmarks, and real-world tasks. All methods run GP-UCB with the same β_t, which $\mathrm{UCB}_{\mathrm{br}}$ also uses in its certificate; each cell reports the mean stopping iteration (success rate / final simple regret) over 50 seeds. For each problem, we mark in blue the practical method ($\mathrm{UCB}_{\mathrm{br}}$, PRB, Acq, ΔCB, ΔES) that stops in the fewest iterations among those reaching the target success rate $1-\delta = 95\%$; when no method reaches it, we mark the one with the highest success rate. NOSTOP and Oracle_r are reference baselines and are excluded from the marking, as are methods that never stop before the budget T. For RFHPO and FMNIST-SVM, the final simple regret is the validation error, with $\epsilon = 0.03$ and $\epsilon = 0.165$, respectively.

Problem	D	T	NOSTOP	Oracle_r	$\mathrm{UCB}_{\mathrm{br}}$	PRB	Acq	ΔCB	ΔES
GP† 3×10^{-3}	2	128	128.0 (100% / 0.000347)	6.0 (100% / 0.0162)	76.6 (100% / 0.000347)	6.4 (98% / 0.0164)	67.2 (100% / 0.00794)	9.6 (100% / 0.00798)	6.4 (98% / 0.0183)
GP† 5×10^{-3}	2	128	128.0 (100% / 0.000899)	6.3 (100% / 0.0181)	115.2 (100% / 0.000899)	7.4 (94% / 0.0219)	69.5 (98% / 0.0111)	16.7 (100% / 0.00351)	6.4 (94% / 0.0267)
GP† 3×10^{-3}	6	512	512.0 (100% / 0.00202)	9.6 (100% / 0.0408)	483.3 (100% / 0.00202)	13.9 (96% / 0.0227)	330.1 (66% / 0.0727)	40.2 (100% / 0.00691)	7.4 (34% / 0.148)
GP† 5×10^{-3}	6	512	512.0 (100% / 0.00297)	11.7 (100% / 0.0349)	509.9 (100% / 0.00297)	16.2 (92% / 0.0314)	330.3 (70% / 0.0663)	64.5 (100% / 0.00972)	6.9 (24% / 0.172)
GP 3×10^{-3}	2	128	128.0 (100% / 0.000932)	8.6 (100% / 0.0209)	29.4 (94% / 0.0210)	6.2 (88% / 0.0344)	62.6 (94% / 0.0206)	6.3 (88% / 0.0343)	6.5 (92% / 0.0302)
GP 5×10^{-3}	2	128	128.0 (100% / 0.00276)	6.5 (100% / 0.0212)	38.2 (92% / 0.0232)	6.0 (86% / 0.0388)	62.2 (92% / 0.0229)	6.4 (88% / 0.0368)	7.1 (90% / 0.0341)
GP 3×10^{-3}	6	512	512.0 (100% / 0.00170)	17.2 (100% / 0.0410)	123.9 (72% / 0.0511)	8.0 (20% / 0.174)	300.6 (64% / 0.0831)	9.6 (18% / 0.175)	6.6 (12% / 0.183)
GP 5×10^{-3}	6	512	512.0 (100% / 0.00306)	21.6 (100% / 0.0448)	79.9 (60% / 0.0865)	7.8 (18% / 0.183)	279.6 (58% / 0.0958)	9.5 (18% / 0.181)	6.5 (14% / 0.190)
Branin	2	128	128.0 (100% / 0.00148)	58.2 (100% / 0.0491)	87.0 (88% / 0.313)	62.3 (82% / 1.09)	21.3 (18% / 5.99)	55.8 (66% / 2.04)	82.0 (62% / 4.26)
Rosenbrock	4	96	96.0 (100% / 0.00608)	8.5 (100% / 0.0377)	80.4 (100% / 0.00868)	61.1 (96% / 0.0199)	75.8 (96% / 0.0194)	74.5 (96% / 0.0195)	86.2 (98% / 0.0127)
Levy	4	96	96.0 (98% / 0.0448)	56.0 (98% / 0.0698)	48.8 (42% / 0.273)	45.7 (42% / 0.316)	46.4 (42% / 0.299)	40.3 (36% / 0.340)	55.7 (56% / 0.241)
Hartmann	3	64	64.0 (100% / 0.0167)	46.9 (100% / 0.0464)	53.2 (70% / 0.388)	34.8 (54% / 1.09)	64.0 (100% / 0.0167)	38.5 (52% / 1.18)	49.8 (74% / 0.515)
RFHPO	5	128	128.0 (94% / 0.0174)	25.0 (94% / 0.0250)	57.3 (78% / 0.0233)	15.0 (62% / 0.0294)	6.4 (30% / 0.0376)	16.1 (56% / 0.0301)	6.1 (30% / 0.0383)
FMNIST-SVM	2	64	64.0 (88% / 0.155)	16.9 (88% / 0.160)	32.6 (76% / 0.158)	11.9 (60% / 0.161)	6.0 (48% / 0.165)	10.8 (60% / 0.162)	6.1 (48% / 0.165)

Table 3: GPU runtime breakdown per BO iteration. Each stopping-rule cost is the mean per-iteration cost of evaluating the stopping test, with its share reported as the rule's cost divided by the sum of the rule's cost and the BO baseline.

	BO baseline (ms)				Stopping rule cost time (ms) and share (%)			
Problem	GP fit	Acq opt	Obj eval	Total	**$\mathrm{UCB}_{\mathrm{br}}$**	**PRB**	**ΔCB**	**Acq**
Branin	196	76	0.06	272	**0.95 (0.36)**	81.1 (23.7)	110.7 (29.8)	272.3 (51.0)
Rosenbrock	72	73	0.06	145	**0.92 (0.62)**	80.5 (35.5)	86.8 (37.3)	146.2 (49.8)

4.4 Results

4.4.1 Full Results

Table 2 presents the experimental results under the Matérn-$5/2$ kernel; results for the SE kernel are provided in Appendix C. We report the mean stopping iteration, success rate (fraction of runs achieving final simple regret $\leq \epsilon$), and final simple regret for each method. The reported simple regret evaluates the incumbent $\boldsymbol{x}_t^+ = \arg\max_{i\leq t} f(\boldsymbol{x}_i)$; since $f^* - f(\boldsymbol{x}_t^+) \leq r_t$, satisfying this metric is implied by the instantaneous-regret guarantee underlying our stopping criterion (Section 3.1). $\mathrm{UCB}_{\mathrm{br}}$ maintains the success rate of Oracle_r on the GP† problems and Rosenbrock, achieving an average success rate of 83.7%, which surpasses all competing methods except the oracle baseline. Compared to PRB, $\mathrm{UCB}_{\mathrm{br}}$ generally achieves higher success rates across problems while also attaining lower final regret. Among the criteria that stop before the budget, it attains the lowest final regret on 11 of the 14 problems. In contrast, Acq, ΔCB, and ΔES exhibit inconsistent behavior, at times terminating too aggressively and at others being overly conservative, leading to unreliable performance across different problem settings. This demonstrates a favorable trade-off of our method between evaluation savings and solution quality.

$\mathrm{UCB}_{\mathrm{br}}$ achieves competitive performance across all benchmarks in Table 2. On the harder GP-learned 6D variants where the surrogate hyperparameters are estimated online, $\mathrm{UCB}_{\mathrm{br}}$'s 72% and 60% success contrasts sharply with PRB (18–20%), ΔCB (18%), and ΔES (12–14%), and exceeds that of Acq (58–64%) with less than half of its evaluations. On the synthetic benchmarks Branin and Rosenbrock, $\mathrm{UCB}_{\mathrm{br}}$ remains competitive with success rates of 88% and 100%, respectively, while reducing evaluations relative to NOSTOP. On the two real-world HPO tasks, $\mathrm{UCB}_{\mathrm{br}}$ achieves 78% and 76% success, 16 points above the strongest baseline, while using about half of the evaluation budget.

4.4.2 Test-time Overhead

We quantify the per-iteration wall-clock cost of different stopping rules to measure their practical overhead relative to the BO loop. All experiments are run on a GPU over two standard benchmarks (Branin-2D, Rosenbrock-4D). We consider four stopping methods: $\mathrm{UCB}_{\mathrm{br}}$, PRB, ΔCB, and Acq. For each problem, we run BO for 100 iterations and repeat across 50 random seeds. We decompose each iteration into GP fitting, acquisition optimization, and objective evaluation, and then measure the additional stopping-test time. Table 3 reports averaged per-iteration times. The reported stopping-test time is averaged over the actual stopping iterations. The results show that $\mathrm{UCB}_{\mathrm{br}}$ adds only 0.9 ms per iteration ($< 0.7\%$ of the full iteration time and 0.4% relative to the BO baseline), making its overhead practically negligible. In contrast, PRB and ΔCB incur substantial overhead (roughly 24–37% of the full iteration time), and Acq can consume nearly half of the iteration time.

5 Scope and Limitations

Our (ϵ, δ) guarantee is established under Assumption 2.1, the standard setting of the Bayesian GP-UCB analysis [Srinivas et al., 2009]. Our experiments also cover settings beyond it, with hyperparameters estimated online and objectives that are not GP samples; under a misspecified kernel smoothness, the criterion is ϵ-optimal in all 200 runs (Appendix C.1.3). The tightened bound is developed for GP-UCB, while the stopping principle of Lemma 3.1 applies to any BO algorithm with a computable, converging regret certificate, so tightened certificates for other acquisition functions can follow the same template. Finally, ϵ is task dependent; in practice, it is chosen as the smallest decision-relevant difference on the objective scale.

6 Conclusions

This paper proposes stopping criteria for BO algorithms based on regret bounds. The stopping criteria apply to BO algorithms with a converging subsequence of instantaneous regret bounds and ensure (ϵ, δ)-optimal solutions upon exit. We develop tight instantaneous regret bounds for GP-UCB with novel proof techniques. Using this new bound, we propose detailed forms of stopping criteria for GP-UCB. Numerical experiments are conducted to demonstrate the performance of the proposed stopping criteria.

Acknowledgement

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

Funding on theoretical analysis was provided by the Office of Science, Office of Advanced Scientific Computing Research, Advancements in Artificial Intelligence for Science program. The experimental work was supported by the LLNL LDRD program under project 26-FS-003.

References

[1] Maximilian Balandat, Brian Karrer, Daniel Jiang, Samuel Daulton, Ben Letham, Andrew G Wilson, and Eytan Bakshy. Botorch: A framework for efficient monte-carlo bayesian optimization. *Advances in neural information processing systems*, 33:21524–21538, 2020.

[2] S. R. Chowdhury and A. Gopalan. On kernelized multi-armed bandits. In *International Conference on Machine Learning*, pages 844–853. PMLR, 2017.

[3] T. M. Cover. *Elements of information theory*. John Wiley & Sons, 1999.

[4] Z. Dai, H. Yu, B. K. H. Low, and P. Jaillet. Bayesian optimization meets bayesian optimal stopping. In *International conference on machine learning*, pages 1496–1506. PMLR, 2019.

[5] P. I. Frazier. Bayesian optimization. In *Recent advances in optimization and modeling of contemporary problems*, pages 255–278, October 2018.

[6] Peter I Frazier and Jialei Wang. Bayesian optimization for materials design. In *Information science for materials discovery and design*, pages 45–75. Springer, 2015.

[7] Hans J He, Alec Koppel, Amrit Singh Bedi, Mazen Farhood, and Daniel J Stilwell. Automatically stopping bayesian optimization under uncertain surrogate model means. In *2025 IEEE 64th Conference on Decision and Control (CDC)*, pages 6304–6311. IEEE, 2025.

[8] H. Ishibashi, M. Karasuyama, I. Takeuchi, and H. Hino. A stopping criterion for bayesian optimization by the gap of expected minimum simple regrets. In *International Conference on Artificial Intelligence and Statistics*, pages 6463–6497. PMLR, 2023.

[9] Shogo Iwazaki. Improved regret bounds for gaussian process upper confidence bound in bayesian optimization. *arXiv preprint arXiv:2506.01393*, 2025.

[10] Yimeng Jin and Priyank V Kumar. Bayesian optimisation for efficient material discovery: a mini review. *Nanoscale*, 15(26):10975–10984, 2023.

[11] Remi Roger Alain Paul Lam. *Scaling Bayesian optimization for engineering design: lookahead approaches and multifidelity dimension reduction*. PhD thesis, Massachusetts Institute of Technology, 2018.

[12] Shuang Li, Ke Li, and Wei Li. "why not looking backward?" a robust two-step method to automatically terminate Bayesian optimization. *Advances in Neural Information Processing Systems*, 36:43435–43446, 2023.

[13] Marius Lindauer, Katharina Eggensperger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, Tim Ruhkopf, René Sass, and Frank Hutter. Smac3: A versatile bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research*, 23(54):1–9, 2022.

[14] Romy Lorenz, Ricardo P. Monti, Ines R. Violante, Aldo A. Faisal, Christoforos Anagnostopoulos, Robert Leech, and Giovanni Montana. Stopping criteria for boosting automatic experimental design using real-time fmri with bayesian optimization, 2016. URL `https://arxiv.org/abs/1511.07827`.

[15] A. Makarova, H. Shen, V. Perrone, A. Klein, J. B. Faddoul, A. Krause, M. Seeger, and C. Archambeau. Automatic termination for hyperparameter optimization. In *International Conference on Automated Machine Learning*, pages 7–1. PMLR, 2022.

[16] M. McLeod, S. Roberts, and M. A. Osborne. Optimization, fast and slow: optimally switching between local and bayesian optimization. In *International Conference on Machine Learning*, pages 3443–3452. PMLR, 2018.

[17] V. Nguyen, S. Gupta, S. Rana, C. Li, and S. Venkatesh. Regret for expected improvement over the best-observed value and stopping condition. In *Proceedings of the Ninth Asian Conference on Machine Learning*, volume 77, pages 279–294. PMLR, 2017.

[18] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.

[19] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. *Advances in neural information processing systems*, 25, 2012.

[20] N. Srinivas, A. Krause, S. M. Kakade, and M. Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. *arXiv preprint arXiv:0912.3995*, 2009.

[21] Timothy Tay and Carolina Osorio. Bayesian optimization techniques for high-dimensional simulation-based transportation problems. *Transportation Research Part B: Methodological*, 164:210–243, 2022.

[22] S. Vakili, K. Khezeli, and V. Picheny. On information gain and regret bounds in Gaussian process bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 82–90. PMLR, 2021.

[23] Pauli Virtanen et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

[24] J. Wang and P. Papadopoulos. Finite element analysis-enabled optimization of process parameters in additive manufacturing. *Finite Elements in Analysis and Design*, 244:104282, 2025.

[25] J. Wang, N. Y. Chiang, A. Gillette, and J. L. Peterson. A multifidelity Bayesian optimization method for inertial confinement fusion design. *Phys Plasmas*, 2024.

[26] J. Wang, H. Wang, S. Ng, and C. G. Petra. Bayesian optimization with expected improvement: No regret and the choice of incumbent. *arXiv preprint arXiv:2508.15674*, 2025.

[27] Z. Wang and N. de Freitas. Theoretical analysis of Bayesian optimisation with unknown Gaussian process hyper-parameters, 2014. URL https://arxiv.org/abs/1406.7758.

[28] Justin Whitehouse, Aaditya Ramdas, and Steven Z Wu. On the sublinear regret of gp-ucb. *Advances in Neural Information Processing Systems*, 36:35266–35276, 2023.

[29] J. Wilson. Stopping bayesian optimization with probabilistic regret bounds. *Advances in Neural Information Processing Systems*, 37:98264–98296, 2024.

[30] A. Wächter and L. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.*, 106(1):25–57, 2006.

[31] Q. Xie, L. Cai, A. Terenin, P. I. Frazier, and Z. Scully. Cost-aware stopping for bayesian optimization. *arXiv preprint arXiv:2507.12453*, 2025.

A Proofs of Section 3

The proof of Lemma 3.1 is given below.

Proof. By the condition of the lemma, there exists a subsequence $\{c_{t_i}\}$ such that $c_{t_i} \to 0$ as $i \to \infty$. Hence, for any $\epsilon > 0$, there exists an index $\bar{i}$ such that

$$c_{t_i} < \epsilon, \qquad \forall i \geq \bar{i}.$$

Therefore, the stopping criterion (3) is satisfied at iteration $t_{\bar{i}}$, and Algorithm 1 exits in finite iterations $\bar{t} \leq t_{\bar{i}}$.

From (2),

$$r_{\bar{t}} = f^* - f(\boldsymbol{x}_{\bar{t}}) \leq c_{\bar{t}}$$

with probability $\geq 1 - \delta$. Since $c_{\bar{t}} < \epsilon$, we obtain

$$r_{\bar{t}} \leq \epsilon$$

with probability $\geq 1 - \delta$. Hence, $\boldsymbol{x}_{\bar{t}}$ is an (ϵ, δ)-optimal solution. □

Next, we recall the definition of maximum information gain.

Definition A.1. For any set $A = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T\} \subset D$, let $\boldsymbol{y}_A = [y(\boldsymbol{x}_1), \ldots, y(\boldsymbol{x}_T)]^\top$ and $\boldsymbol{f}_A = [f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_T)]^\top$. The maximum information gain after T samples is $\gamma_T := \max_{A \subset D,\ |A|=T} I(\boldsymbol{y}_A; \boldsymbol{f}_A)$.

Readers are referred to Cover [1999], Srinivas et al. [2009] for a detailed definition of the maximum information gain. γ_t is often used to quantify regret upper bounds [Srinivas et al., 2009], specifically the sum of posterior standard deviation at sample points. For the popular SE and Matérn kernels, the rates of γ_t are well-established [Iwazaki, 2025, Vakili et al., 2021].

The following widely used lemma is useful in establishing cumulative regret bounds for BO algorithms

Lemma A.2 (Srinivas et al. [2009])**.** *The sum of GP posterior variances $\sigma^2_{i-1}(\boldsymbol{x}_i)$ at the next samples $\boldsymbol{x}_i$ for $i = 1, \ldots, t$ satisfies*

$$\sum_{i=1}^{t} \sigma^2_{i-1}(\boldsymbol{x}_i) \leq C_\gamma \gamma_t, \tag{14}$$

where $C_\gamma = \frac{2}{\log(1+\sigma^{-2})}$ and γ_t is the maximum information gain after t iterations or samples.

The following lemma provides the rate on γ_t for SE and Matérn kernels.

Lemma A.3 (Vakili et al. [2021], Iwazaki [2025])**.** *For a GP with t samples, the SE kernel has $\gamma_t = \mathcal{O}(\log^{d+1}(t))$, and the Matérn kernel with $\nu > \frac{1}{2}$ has $\gamma_t = \mathcal{O}(t^{\frac{d}{2\nu+d}} \log^{\frac{2\nu}{2\nu+d}}(t))$.*

B Proofs and Additional Results for Section 3.1

We first recall the definition of the Reproducing kernel Hilbert space (RKHS).

Definition B.1. Let k be a positive definite kernel $k : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ with respect to a finite Borel measure supported on $\mathcal{X}$. A Hilbert space H_k of functions on $\mathcal{X}$ with inner product $\langle \cdot, \cdot \rangle_{H_k}$ is called a RKHS with kernel k if $k(\cdot, \boldsymbol{x}) \in H_k$ for all $\boldsymbol{x} \in \mathcal{X}$, and $\langle f, k(\cdot, \boldsymbol{x}) \rangle_{H_k} = f(\boldsymbol{x})$ for all $\boldsymbol{x} \in \mathcal{X}, f \in H_k$. The induced RKHS norm $\|f\|_{H_k} = \sqrt{\langle f, f \rangle_{H_k}}$ measures the smoothness of f with respect to the kernel k.

Next, the noisy frequentist assumptions are given.

Assumption B.2. The frequentist setting satisfies the following conditions. The function f lies in the RKHS, denoted as $\mathcal{H}_k(D)$ associated with the bounded kernel k, with the norm $\|\cdot\|_{H_k}$. The kernels satisfies $k(\boldsymbol{x}, \boldsymbol{x}') \leq 1$, and $k(\boldsymbol{x}, \boldsymbol{x}) = 1$ for $\forall \boldsymbol{x}, \boldsymbol{x}' \in D$. The RKHS norm is bounded above by a constant $B \geq 0$, *i.e.*, $\|f\|_{H_k} \leq B$. Moreover, the domain D is compact. The noise ϵ_t is conditionally R-sub-Gaussian for a fixed $R > 0$, *i.e.*,

$$\forall t \geq 0, \forall \lambda \in \mathbb{R}, \mathbb{E}[e^{\lambda \epsilon_t} | \mathcal{F}_{t-1}] \leq \exp\left(\frac{\lambda^2 R^2}{2}\right). \tag{15}$$

We first recall important lemmas for confidence intervals under Assumption 2.1.

Lemma B.3. *The PDF and CDF of the standard normal distribution satisfy* $0 < \phi(x) \leq \phi(0), \Phi(x) \in (0,1)$*, for any* $x \in \mathbb{R}$*. Given a random variable* r *that follows the standard normal distribution:* $r \sim \mathcal{N}(0,1)$*, the probability of* $r > c, c > 0$*, satisfies* $\mathbb{P}\{r > c\} = 1 - \Phi(c)$*. Similarly, for* $c < 0$*,* $\mathbb{P}\{r < c\} = \Phi(c)$*.*

The following lemma recalls a global lower bound of $\sigma_t(\boldsymbol{x})$.

Lemma B.4. *For* $\forall \boldsymbol{x} \in D$ *and* $t \in \mathbb{N}$*, the posterior standard deviation has the lower bound*

$$\sigma_t(\boldsymbol{x}) \geq \sigma \sqrt{\frac{1}{t + \sigma^2}}. \tag{16}$$

Lemma B.4 is a widely accepted result and can be found in Wang and de Freitas [2014].

We now recall (see Section 5 in Srinivas et al. [2009]) and introduce changes to the confidence intervals.

Lemma B.5. *For any* $t \in \mathbb{N}$*,* $\boldsymbol{x} \in D$ *deterministic with* $\mathcal{F}_{t-1}$*, and* $\beta > 0$*,*

$$|f(\boldsymbol{x}) - \mu_{t-1}(\boldsymbol{x})| \leq \beta^{1/2} \sigma_{t-1}(\boldsymbol{x}), \tag{17}$$

holds with probability $1 - 2(1 - \Phi(\beta^{1/2}))$*. Moreover, the one-sided inequalities hold with probability* $\Phi(\beta^{1/2})$*, i.e., with probability* $\Phi(\beta^{1/2})$*,* $f(\boldsymbol{x}) - \mu_{t-1}(\boldsymbol{x}) \leq \beta^{1/2} \sigma_{t-1}(\boldsymbol{x})$*, and with probability* $\geq \Phi(\beta^{1/2})$*,* $f(\boldsymbol{x}) - \mu_{t-1}(\boldsymbol{x}) \geq -\beta^{1/2} \sigma_{t-1}(\boldsymbol{x})$*.*

Proof. Under Assumption 2.1, given $\mathcal{F}_{t-1}$, $f(\boldsymbol{x}) \sim \mathcal{N}(\mu_{t-1}(\boldsymbol{x}), \sigma^2_{t-1}(\boldsymbol{x}))$. Then,

$$\mathbb{P}\left\{ f(\boldsymbol{x}) - \mu_{t-1}(\boldsymbol{x}) > \beta^{1/2} \sigma_{t-1}(\boldsymbol{x}) \right\} = 1 - \Phi(\beta^{1/2}). \tag{18}$$

Similarly,

$$\mathbb{P}\left\{ f(\boldsymbol{x}) - \mu_{t-1}(\boldsymbol{x}) < -\beta^{1/2} \sigma_{t-1}(\boldsymbol{x}) \right\} = 1 - \Phi(\beta^{1/2}). \tag{19}$$

Thus, by union bound,

$$\mathbb{P}\left\{ |f(\boldsymbol{x}) - \mu_{t-1}(\boldsymbol{x})| \leq \beta^{1/2} \sigma_{t-1}(\boldsymbol{x}) \right\} = 1 - 2(1 - \Phi(\beta^{1/2})).$$

□

Lemma B.5 is extended to a discrete set $\mathbb{D}_t$ and all $t \in \mathbb{N}$ next.

Lemma B.6. *For any given* $t \in \mathbb{N}$*, a discrete set* $\mathbb{D}_t \subseteq D$*, and* $\alpha_t > 0$*, we have*

$$f(\boldsymbol{x}) - \mu_{t-1}(\boldsymbol{x}) \leq \alpha_t^{1/2} \sigma_{t-1}(\boldsymbol{x}), \tag{20}$$

holds with probability $\geq 1 - |\mathbb{D}_t|(1 - \Phi(\alpha_t^{1/2}))$*, simultaneously for all* $\boldsymbol{x} \in \mathbb{D}_t$*.*

The proof is the same as that for Lemma 5.6 in Srinivas et al. [2009]. Given that $\boldsymbol{x}^*$ is random with respect to $\mathcal{F}_{t-1}$, we take the union bound over $\mathbb{D}_t$, as in Srinivas et al. [2009]. Further, (20) is only one-sided.

We now consider the confidence interval at $\boldsymbol{x}_t$.

Lemma B.7. *For any given* $t \in \mathbb{N}$ *and parameter* $\eta_t > 0$*,*

$$f(\boldsymbol{x}_t) - \mu_{t-1}(\boldsymbol{x}_t) \geq -\eta_t^{1/2} \sigma_{t-1}(\boldsymbol{x}_t), \tag{21}$$

holds with probability $\geq 1 - (1 - \Phi(\eta_t^{1/2}))$*.*

Lemma B.7 is a direct result of Lemma B.5 since $\boldsymbol{x}_t$ is determined with $\mathcal{F}_{t-1}$.

Next, we discuss our parameterized discretization $\mathbb{D}_t$ of D. For any $t \in \mathbb{N}$, $\mathbb{D}_t$ is a predetermined set of finite points in D. The discretization $\mathbb{D}_t$ aims to cover the compact set D with just enough points

so that for any $\boldsymbol{x} \in D$ and $t \in \mathbb{N}$, the distance between $\boldsymbol{x} \in D$ and $\mathbb{D}_t$ is small enough and controlled. As earlier, denote the closest point of $\boldsymbol{x}$ to $\mathbb{D}_t$ as

$$[\boldsymbol{x}]_t := \underset{\boldsymbol{w} \in \mathbb{D}_t}{\operatorname{argmin}} \|\boldsymbol{x} - \boldsymbol{w}\|, \tag{22}$$

and define the parameter h_t as a measure of the distance between D and $\mathbb{D}_t$ by

$$\|\boldsymbol{x} - [\boldsymbol{x}]_t\|_1 \leq h_t, \ \forall \boldsymbol{x} \in D, \tag{23}$$

where $\|\cdot\|_1$ is the one-norm. In this paper, we adopt the standard discretizations with decreasing distances, similar to that in Srinivas et al. [2009]. By (23), $\mathbb{D}_t$ that corresponds to h_t is constructed with $(\frac{rd}{h_t})^d$ points. Given Assumption 2.1 and $D \subseteq [0, r]^d$, this leads to a choice of

$$|\mathbb{D}_t| = \left\lceil \left(\frac{rd}{h_t} \right)^d \right\rceil, \tag{24}$$

where the points in $\mathbb{D}_t$ satisfy (23). We emphasize that the discretization does not play a role in the GP-UCB algorithm. Rather, it is an analytic tool to help analysis. We choose

$$h_t = \frac{1}{b\sqrt{\log(n_L da/\delta)} t^{s_t}},$$

where $s_t > 0$ is a parameter in (6) and $n_L \geq 1$ is the probability associated with Lipschitz continuity. It is easy to verify that the choice (8) satisfies (24). We have the following Lemma on the Lipschitz continuity of f (see also Lemma 5.7 in Srinivas et al. [2009]).

Lemma B.8. *Let the discretization be such that* (8) *is satisfied. Then, under Assumption 2.1, for all* $\boldsymbol{x} \in D$ *and* $t \in \mathbb{N}$*,*

$$|f(\boldsymbol{x}) - f([\boldsymbol{x}]_t)| \leq \frac{1}{t^{s_t}}, \tag{25}$$

with probability $\geq 1 - \delta/n_L$*.*

We emphasize that the discretization $\mathbb{D}_t$ in Srinivas et al. [2009] is recovered if $s_t = 2$ and $n_L = 4$.

Next, we consider a one-sided confidence interval at $[\boldsymbol{x}^*]_t$ using Lemma B.6.

Lemma B.9. *For any given* $t \in \mathbb{N}$*, the discretization* $\mathbb{D}_t \subseteq D$*, and* $\alpha_t > 0$*, we have*

$$f([\boldsymbol{x}^*]_t) - \mu_{t-1}([\boldsymbol{x}^*]_t) \leq \alpha_t^{1/2} \sigma_{t-1}([\boldsymbol{x}^*]_t), \tag{26}$$

holds with probability $\geq 1 - |\mathbb{D}_t|(1 - \Phi(\alpha_t^{1/2}))$.

From the definitions of the parameters (6) and (7), we have

$$\pi_t n_{\eta,t} (1 - \Phi(\eta_t^{1/2})) = 1 \tag{27}$$

and

$$\pi_t n_{\alpha,t} |\mathbb{D}_t| (1 - \Phi(\alpha_t^{1/2})) = 1,$$

The proof of Lemma 3.3 is given below.

Proof. At a given $t \in \mathbb{N}$, we start from the definition of the instantaneous regret bound

$$r_t = f(\boldsymbol{x}^*) - f(\boldsymbol{x}_t).$$

From Lemma B.7 and (27), we have the confidence interval

$$f(\boldsymbol{x}_t) - \mu_{t-1}(\boldsymbol{x}_t) \geq -\eta_t^{1/2} \sigma_{t-1}(\boldsymbol{x}_t), \tag{28}$$

with probability $\geq 1 - \frac{1}{\pi_t n_{\eta,t}}$, where $\eta_t^{1/2} > 0$ is a parameter determined by $n_{\eta,t}$ and can be different from $\beta_t^{1/2}$.

From the discretization and Lemma B.8, we have (25) for $\forall \boldsymbol{x} \in D$ with probability $\geq 1 - \delta/n_L$. Applying (28) and (25) at $\boldsymbol{x}^*$, we can write the instantaneous regret bound as

$$r_t \leq f([\boldsymbol{x}^*]_t) - \mu_{t-1}(\boldsymbol{x}_t) + \eta_t^{1/2} \sigma_{t-1}(\boldsymbol{x}_t) + \frac{1}{t^{s_t}}, \tag{29}$$

with probability $\geq 1 - \frac{1}{\pi_t n_{\eta,t}} - \delta/n_L$. Next, we replace $f([\boldsymbol{x}^*]_t)$ in (29) via Lemma B.6 at $[\boldsymbol{x}^*]_t$ to obtain

$$r_t \leq \mu_{t-1}([\boldsymbol{x}^*]_t) - \mu_{t-1}(\boldsymbol{x}_t) + \alpha_t^{1/2}\sigma_{t-1}([\boldsymbol{x}^*]_t) + \eta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) + \frac{1}{t^{s_t}}, \tag{30}$$

with probability $\geq 1 - \frac{1}{\pi_t n_{\eta,t}} - \frac{1}{\pi_t n_{\alpha,t}} - \delta/n_L$. From the definition of UCB, we have

$$\mu_{t-1}([\boldsymbol{x}^*]_t) + \beta_t^{1/2}\sigma_{t-1}([\boldsymbol{x}^*]_t) \leq \mu_{t-1}(\boldsymbol{x}_t) + \beta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t). \tag{31}$$

Using (31) in (30), we have

$$\begin{aligned} r_t \leq & \beta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) - \beta_t^{1/2}\sigma_{t-1}([\boldsymbol{x}^*]_t) + \alpha_t^{1/2}\sigma_{t-1}([\boldsymbol{x}^*]_t) + \eta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) + \frac{1}{t^{s_t}} \\ \leq & \beta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) - (\beta_t^{1/2} - \alpha_t^{1/2})\sigma_{t-1}([\boldsymbol{x}^*]_t) + \eta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) + \frac{1}{t^{s_t}}, \end{aligned} \tag{32}$$

with probability $\geq 1 - \frac{1}{\pi_t n_{\eta,t}} - \frac{1}{\pi_t n_{\alpha,t}} - \delta/n_L$. Since $\alpha_t \leq \beta_t$, by the global lower bound of $\sigma_{t-1}(\boldsymbol{x})$ in Lemma B.4, (32) implies

$$r_t \leq \beta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) - (\beta_t^{1/2} - \alpha_t^{1/2})\sigma\sqrt{\frac{1}{t+\sigma^2}} + \eta_t^{1/2}\sigma_{t-1}(\boldsymbol{x}_t) + \frac{1}{t^{s_t}}, \tag{33}$$

with probability $\geq 1 - \frac{1}{\pi_t n_{\eta,t}} - \frac{1}{\pi_t n_{\alpha,t}} - \delta/n_L$, which leads to (9).

□

Remark B.10. Without the assumption $\alpha_t \leq \beta_t$, the coefficient of $\sigma_{t-1}([\boldsymbol{x}^*]_t)$ in (32) may become positive, in which case the global lower bound on $\sigma_{t-1}([\boldsymbol{x}^*]_t)$ is no longer useful. From the first inequality in (32), decreasing $\alpha_t^{1/2}$ tightens the upper bound on r_t. Thus, we impose $\alpha_t^{1/2} \leq \beta_t^{1/2}$, which makes the coefficient nonpositive and allows us to use the global lower bound for $\sigma_{t-1}([\boldsymbol{x}^*]_t)$.

Proof of Theorem 3.6 is given next.

Proof. For each t, Lemma 3.3 applies with $\bar{n}_{\eta,t}$, $\bar{n}_{\alpha,t}$, and $\bar{s}_t$ obtained from (10). By the first constraint in (10),

$$\frac{1}{\bar{n}_{\eta,t}} + \frac{1}{\bar{n}_{\alpha,t}} \leq \delta\left(1 - \frac{1}{n_L}\right). \tag{34}$$

Therefore, at a given t, we have

$$\begin{aligned} \mathbb{P}\left\{(9) \text{ holds at } t \in \mathbb{N}\right\} &\geq 1 - \frac{1}{\pi_t}\left(\frac{1}{\bar{n}_{\eta,t}} + \frac{1}{\bar{n}_{\alpha,t}}\right) - \frac{\delta}{n_L} \\ &\geq 1 - \frac{1}{\pi_t}\delta\left(1 - \frac{1}{n_L}\right) - \frac{\delta}{n_L}. \end{aligned} \tag{35}$$

The Lipschitz/discretization inequality in Lemma B.8 holds simultaneously for all $t \in \mathbb{N}$ with probability $\geq 1 - \delta/n_L$. Applying union bound to the t-dependent confidence intervals leads to

$$\begin{aligned} \mathbb{P}\left\{(9) \text{ holds for all } t \in \mathbb{N}\right\} &\geq 1 - \sum_{t=1}^{\infty}\frac{1}{\pi_t}\delta\left(1 - \frac{1}{n_L}\right) - \frac{\delta}{n_L} \\ &= 1 - \delta, \end{aligned} \tag{36}$$

where the equality uses $\sum_{t=1}^{\infty} 1/\pi_t = 1$. □

Remark B.11. We note that the probability (35) can be derived first by conditioning on the history $\mathcal{H}_{t-1}$ and filtration $\mathcal{F}_{t-1}$. Then, the conditional probability lead to probability by taking the expectation of both sides. Such technique is standard in BO literature [Srinivas et al., 2009].

More comparison between the existing bound (5) and our new bound (12) with different constants are given below, where the latter consistently reduces the upper bound.

Table 4: Comparison between the existing bound (5) and our new bound (12). Fixed parameters: $\sigma = 10^{-2}$, $n_L = 10$, $a = b = r = 1$, $d = 2$, $\delta = 0.1$, and $\pi_t = \pi^2 t^2/6$.

t	$c_1(t)$	$c_2(t)$	$\sqrt{\beta_t}$	$\sqrt{\eta_t}$	$\sqrt{\alpha_t}$	$\bar{s}_t$	$\bar{n}_{\eta,t}$	$\bar{n}_{\alpha,t}$	(12)	(5)	Ratio
20	2.24×10^{-3}	2.24×10^{-3}	7.075	3.751	7.050	2.464	17.28	31.13	2.48×10^{-2}	3.41×10^{-2}	0.726
20	1.00×10^{-2}	2.24×10^{-3}	7.075	3.680	7.075	2.346	13.04	75.08	1.08×10^{-1}	1.44×10^{-1}	0.753
20	1.00×10^{-1}	2.24×10^{-3}	7.075	3.648	7.075	2.103	11.51	322.26	1.07	1.42	0.758
100	1.00×10^{-3}	1.00×10^{-3}	8.328	4.499	7.710	1.797	17.78	29.62	1.25×10^{-2}	1.68×10^{-2}	0.744
100	1.00×10^{-2}	1.00×10^{-3}	8.328	4.410	7.961	1.803	11.75	205.60	1.27×10^{-1}	1.67×10^{-1}	0.764
100	1.00×10^{-1}	1.00×10^{-3}	8.328	4.399	8.248	1.811	11.17	2028.62	1.27	1.67	0.764
500	4.47×10^{-4}	4.47×10^{-4}	9.417	5.144	8.317	1.473	18.13	28.71	6.13×10^{-3}	8.43×10^{-3}	0.727
500	1.00×10^{-2}	4.47×10^{-4}	9.417	5.057	8.640	1.479	11.41	426.00	1.44×10^{-1}	1.88×10^{-1}	0.767
500	1.00×10^{-1}	4.47×10^{-4}	9.417	5.052	8.907	1.483	11.14	4288.44	1.45	1.88	0.768

Table 5: Comparison between the existing bound (5) and our new bound (12). Fixed parameters: $\sigma = 10^{-2}$, $n_L = 4$, $a = b = r = 1$, $d = 2$, $\delta = 0.01$, and $\pi_t = \pi^2 t^2/6$.

t	$c_1(t)$	$c_2(t)$	$\sqrt{\beta_t}$	$\sqrt{\eta_t}$	$\sqrt{\alpha_t}$	$\bar{s}_t$	$\bar{n}_{\eta,t}$	$\bar{n}_{\alpha,t}$	(12)	(5)	Ratio
20	2.24×10^{-3}	2.24×10^{-3}	7.450	4.340	7.426	2.481	213.64	354.70	2.69×10^{-2}	3.58×10^{-2}	0.751
20	1.00×10^{-2}	2.24×10^{-3}	7.450	4.274	7.450	2.366	158.28	845.84	1.18×10^{-1}	1.51×10^{-1}	0.779
20	1.00×10^{-1}	2.24×10^{-3}	7.450	4.244	7.450	2.124	138.45	3606.20	1.17	1.49	0.785
100	1.00×10^{-3}	1.00×10^{-3}	8.649	5.005	8.055	1.806	218.00	343.31	1.33×10^{-2}	1.74×10^{-2}	0.765
100	1.00×10^{-2}	1.00×10^{-3}	8.649	4.921	8.293	1.812	141.44	2327.61	1.36×10^{-1}	1.73×10^{-1}	0.783
100	1.00×10^{-1}	1.00×10^{-3}	8.649	4.911	8.568	1.819	134.12	22827.85	1.36	1.73	0.784
500	4.47×10^{-4}	4.47×10^{-4}	9.702	5.596	8.639	1.479	221.15	335.78	6.47×10^{-3}	8.68×10^{-3}	0.745
500	1.00×10^{-2}	4.47×10^{-4}	9.702	5.512	8.948	1.484	137.07	4884.65	1.52×10^{-1}	1.94×10^{-1}	0.783
500	1.00×10^{-1}	4.47×10^{-4}	9.702	5.508	9.205	1.489	133.70	49016.90	1.52	1.94	0.784

C Additional Results

As mentioned in Section 4, Table 7 presents experimental results under the SE kernel. As in Table 2, the marking rule highlights the practical method that stops in the fewest iterations among those reaching success rate $\geq 1 - \delta$, and otherwise the one with the highest success rate; see the table caption for the full specification. The overall trends are consistent with those observed under the Matérn kernel. UCB_{br} continues to achieve a favorable balance between evaluation savings and success rate across most problem settings. On the synthetic benchmarks, UCB_{br} achieves 96% success rate on Rosenbrock, above the target $1 - \delta = 95\%$. For the HPO tasks, UCB_{br} achieves 72% and 76% success while reducing evaluations by 53% and 57% compared to NOSTOP. On the GP-learned 6D benchmarks, UCB_{br} achieves 74% and 76% success—well above PRB (16–22%), Acq (64%), and ΔCB (18–26%)—demonstrating robustness in the harder, hyperparameter-learning regime. Overall, UCB_{br} generally outperforms or matches the baseline stopping methods across the majority of test cases, confirming the robustness of the proposed criterion across different kernel choices.

C.1 Ablation Study

C.1.1 Bayesian Variant Ablation

Table 8 reports the sensitivity to the constants a, b of Assumption 2.1 at the default $\delta = 0.05$. Table 9 reports δ-sensitivity at the default $a = b = 1$, and Table 10 reports ϵ-sensitivity at both defaults. The default configuration $(a = b = 1, \delta = 0.05, \epsilon = 0.1)$ is the same column in all three tables and is highlighted in bold.

The certificate has no calibration constant beyond δ and the constants a, b of Assumption 2.1, which enter β_t and the discretization size $|D_t|$ only through $\log\log(\cdot)$ and $\log b$, respectively. Perturbing either by a decade keeps the success rate within 2 points of the default on the GP d=2 and Rosenbrock rows, and raises it on the learned GP d=6 rows under both b perturbations ($72\% \to 78\%$ and 76%; $60\% \to 72\%$) and on Hartmann at b=0.1 ($70\% \to 86\%$) (Table 8).

Table 6: Comparison between the existing bound (5) and our new bound (12). Fixed parameters: $\sigma = 10^{-2}$, $n_L = 10$, $a = b = r = 1$, $d = 6$, $\delta = 0.01$, and $\pi_t = \pi^2 t^2/6$.

t	$c_1(t)$	$c_2(t)$	$\sqrt{\beta_t}$	$\sqrt{\eta_t}$	$\sqrt{\alpha_t}$	$\bar{s}_t$	$\bar{n}_{\eta,t}$	$\bar{n}_{\alpha,t}$	(12)	(5)	Ratio
20	2.24×10^{-3}	2.24×10^{-3}	11.431	4.284	11.431	2.178	165.74	337.12	3.66×10^{-2}	5.36×10^{-2}	0.683
20	1.00×10^{-2}	2.24×10^{-3}	11.431	4.222	11.431	2.119	125.48	970.01	1.58×10^{-1}	2.31×10^{-1}	0.685
20	1.00×10^{-1}	2.24×10^{-3}	11.431	4.198	11.431	2.014	113.07	6412.58	1.57	2.29	0.684
100	1.00×10^{-3}	1.00×10^{-3}	13.257	4.942	12.266	1.657	157.28	378.50	1.77×10^{-2}	2.66×10^{-2}	0.665
100	1.00×10^{-2}	1.00×10^{-3}	13.257	4.882	12.435	1.660	115.61	2852.55	1.81×10^{-1}	2.65×10^{-1}	0.683
100	1.00×10^{-1}	1.00×10^{-3}	13.257	4.875	12.623	1.663	111.55	27972.48	1.81	2.65	0.684
500	4.47×10^{-4}	4.47×10^{-4}	14.859	5.539	12.922	1.365	159.91	364.08	8.46×10^{-3}	1.33×10^{-2}	0.637
500	1.00×10^{-2}	4.47×10^{-4}	14.859	5.478	13.143	1.368	113.24	5923.30	2.03×10^{-1}	2.97×10^{-1}	0.682
500	1.00×10^{-1}	4.47×10^{-4}	14.859	5.475	13.322	1.370	111.32	59045.61	2.03	2.97	0.684

Table 7: Comparison of stopping criteria under the SE kernel across GP-sampled functions († means known hyperparameters), synthetic benchmarks, and real-world tasks. Each cell reports the mean stopping iteration (success rate / final simple regret); settings and marking follow Table 2, with the SE kernel as the surrogate model.

Problem	D	T	NOSTOP	Oracle$_r$	UCB$_{\text{br}}$	PRB	Acq	ΔCB	ΔES
GP† 3×10^{-3}	2	128	128.0 (100% / 0.000395)	6.0 (100% / 0.0164)	70.3 (100% / 0.000395)	6.3 (98% / 0.0172)	67.2 (100% / 0.00820)	8.6 (100% / 0.0101)	6.4 (98% / 0.0183)
GP† 5×10^{-3}	2	128	128.0 (100% / 0.000470)	6.3 (100% / 0.0173)	112.7 (100% / 0.000470)	7.2 (96% / 0.0185)	67.2 (98% / 0.0114)	15.8 (100% / 0.00345)	6.4 (94% / 0.0267)
GP† 3×10^{-3}	6	512	512.0 (100% / 0.00204)	9.2 (100% / 0.0368)	285.0 (100% / 0.00215)	11.5 (98% / 0.0233)	330.0 (68% / 0.0695)	25.3 (98% / 0.00722)	7.1 (32% / 0.151)
GP† 5×10^{-3}	6	512	512.0 (100% / 0.00229)	10.5 (100% / 0.0334)	440.7 (100% / 0.00229)	14.3 (94% / 0.0303)	320.4 (66% / 0.0762)	41.8 (100% / 0.00693)	6.7 (14% / 0.195)
GP 3×10^{-3}	2	128	128.0 (100% / 0.000506)	6.5 (100% / 0.0189)	35.7 (96% / 0.0104)	6.0 (88% / 0.0343)	64.8 (96% / 0.0173)	6.3 (90% / 0.0319)	6.6 (90% / 0.0295)
GP 5×10^{-3}	2	128	128.0 (100% / 0.00114)	7.1 (100% / 0.0209)	34.0 (94% / 0.0163)	6.1 (84% / 0.0440)	64.7 (90% / 0.0266)	6.3 (86% / 0.0420)	7.2 (90% / 0.0345)
GP 3×10^{-3}	6	512	512.0 (100% / 0.00204)	17.4 (100% / 0.0447)	106.6 (74% / 0.0518)	7.1 (22% / 0.175)	299.6 (64% / 0.0823)	8.6 (26% / 0.165)	6.8 (20% / 0.170)
GP 5×10^{-3}	6	512	512.0 (100% / 0.00425)	21.8 (100% / 0.0458)	160.4 (76% / 0.0641)	7.2 (16% / 0.190)	297.3 (64% / 0.0894)	8.1 (18% / 0.188)	6.5 (12% / 0.193)
Branin	2	128	128.0 (100% / 0.000193)	43.2 (100% / 0.0357)	33.3 (42% / 3.29)	28.0 (40% / 4.33)	11.6 (2% / 7.92)	29.8 (44% / 3.97)	88.7 (68% / 3.89)
Rosenbrock	4	96	96.0 (100% / 0.00560)	8.4 (100% / 0.0379)	78.9 (96% / 0.0156)	57.1 (96% / 0.0223)	75.7 (96% / 0.0206)	59.9 (96% / 0.0221)	70.8 (98% / 0.0117)
Levy	4	96	96.0 (98% / 0.0140)	46.9 (98% / 0.0478)	45.4 (42% / 0.245)	28.3 (26% / 0.350)	39.8 (40% / 0.275)	33.9 (26% / 0.346)	49.2 (50% / 0.235)
Hartmann	3	64	64.0 (96% / 0.0297)	45.8 (96% / 0.0586)	50.7 (64% / 0.471)	32.4 (42% / 1.27)	64.0 (96% / 0.0297)	37.3 (48% / 1.17)	50.3 (68% / 0.365)
RFHPO	5	128	128.0 (94% / 0.0181)	22.8 (94% / 0.0251)	60.4 (72% / 0.0247)	12.1 (50% / 0.0319)	6.5 (32% / 0.0367)	13.6 (48% / 0.0326)	6.2 (32% / 0.0374)
FMNIST-SVM	2	64	64.0 (90% / 0.156)	17.0 (90% / 0.160)	27.3 (76% / 0.158)	9.3 (56% / 0.163)	6.1 (48% / 0.165)	9.8 (58% / 0.163)	6.1 (48% / 0.165)

The effect of δ on the success rate is small: δ enters β_t only logarithmically, and varying it by a factor of 20 keeps the success rate of UCB$_{\text{br}}$ within 2 points of the default on the GP d=2 and Rosenbrock rows and raises it on the GP d=6 rows (Table 9). We restrict the sweep to Matérn-5/2 on the GP and analytic families: the well-specified GP† rows of Table 2 reach 100% success and admit no informative sweep at ϵ=0.1, and SE/RBF results in Table 7 follow the same qualitative pattern.

Table 10 varies the target accuracy ϵ. On every row, a larger ϵ stops no later and succeeds at least as often: at $\epsilon = 0.2$, UCB$_{\text{br}}$ reaches 100% success on both GP d=2 rows using 9–11% of the budget, and 96% and 88% on the GP d=6 rows using 10% of the budget.

C.1.2 Tightened versus Existing Bound

Table 11 isolates the effect of our new bound (12). The criterion is evaluated with it and with the existing bound (5) using the same β_t along the same BO trajectories, so the two differ only in the bound. Across all 14 problems of Table 2, the new bound is never wider than the existing one at any of the 151,700 recorded iterations, and its median width is 36% smaller. On the well-specified GP† problems, the existing bound triggers on 18 of the 50 seeds at d=2, σ=3×10^{-3} and on none of the seeds at d=2, σ=5×10^{-3} and d=6, σ=3×10^{-3}, so the criterion would run the full budget there; the new bound triggers on 48, 24, and 12 seeds with the same final regret.

C.1.3 Kernel-Smoothness Misspecification

Table 12 misspecifies the kernel smoothness of the surrogate. UCB$_{\text{br}}$ is ϵ-optimal in all 200 runs, matching NOSTOP while saving 22–56% of the evaluations, and its final regret is 3–10 times smaller than that of PRB, which fails on 2–3 of the 50 seeds in every setting.

Table 8: Sensitivity of $\mathrm{UCB}_{\mathrm{br}}$ to the constants a and b of Assumption 2.1, which enter β_t, at fixed $\delta = 0.05$, under the Matérn-$5/2$ kernel. Each arm changes one constant by a decade and keeps the other at 1. Cells follow the format of Table 2: mean stopping iteration (success rate / final simple regret). The default (a=b=1) column is in bold and repeats the $\mathrm{UCB}_{\mathrm{br}}$ column of Table 2. Each cell averages over 50 independent random seeds.

Problem	D	T	$\boldsymbol{a = b = 1}$ (default)	$a = 10$	$b = 0.1$	$b = 10$
Branin	2	128	**87.0 (88% / 0.313)**	86.5 (84% / 1.01)	77.1 (84% / 1.57)	91.0 (86% / 0.905)
Rosenbrock	4	96	**80.4 (100% / 0.00868)**	81.2 (98% / 0.0119)	84.7 (98% / 0.0101)	78.7 (98% / 0.0144)
Hartmann	3	64	**53.2 (70% / 0.388)**	49.9 (62% / 0.417)	56.5 (86% / 0.210)	47.6 (62% / 0.523)
GP 3×10^{-3}	2	128	**29.4 (94% / 0.0210)**	35.5 (92% / 0.0203)	26.4 (92% / 0.0185)	35.5 (92% / 0.0197)
GP 5×10^{-3}	2	128	**38.2 (92% / 0.0232)**	31.9 (92% / 0.0204)	29.4 (90% / 0.0212)	45.4 (92% / 0.0190)
GP 3×10^{-3}	6	512	**123.9 (72% / 0.0511)**	158.8 (76% / 0.0456)	150.7 (78% / 0.0409)	168.8 (76% / 0.0568)
GP 5×10^{-3}	6	512	**79.9 (60% / 0.0865)**	105.9 (58% / 0.0952)	137.5 (72% / 0.0627)	191.5 (72% / 0.0581)

Table 9: Sensitivity of $\mathrm{UCB}_{\mathrm{br}}$ to the failure probability δ at fixed $a = b = 1$, under the Matérn-$5/2$ kernel. Cells follow the format of Table 2. The default (δ=0.05) column is in bold and repeats the $\mathrm{UCB}_{\mathrm{br}}$ column of Table 2. Each cell averages over 50 independent random seeds.

Problem	D	T	$\delta = 0.01$	$\boldsymbol{\delta = 0.05}$ (default)	$\delta = 0.2$
Branin	2	128	84.1 (78% / 1.67)	**87.0 (88% / 0.313)**	80.7 (80% / 2.10)
Rosenbrock	4	96	81.1 (100% / 0.0105)	**80.4 (100% / 0.00868)**	79.9 (98% / 0.0138)
Hartmann	3	64	50.0 (62% / 0.433)	**53.2 (70% / 0.388)**	53.4 (72% / 0.301)
GP 3×10^{-3}	2	128	35.7 (92% / 0.0191)	**29.4 (94% / 0.0210)**	31.4 (92% / 0.0191)
GP 5×10^{-3}	2	128	43.3 (90% / 0.0225)	**38.2 (92% / 0.0232)**	30.6 (90% / 0.0230)
GP 3×10^{-3}	6	512	150.1 (78% / 0.0451)	**123.9 (72% / 0.0511)**	160.4 (74% / 0.0519)
GP 5×10^{-3}	6	512	160.8 (68% / 0.0649)	**79.9 (60% / 0.0865)**	167.4 (66% / 0.0692)

Table 10: Sensitivity of $\mathrm{UCB}_{\mathrm{br}}$ to the target accuracy ϵ at fixed $\delta = 0.05$ and $a = b = 1$, under the Matérn-$5/2$ kernel. Cells follow the format of Table 2, with the success rate measured against the column's ϵ. The default (ϵ=0.1) column is in bold and repeats the $\mathrm{UCB}_{\mathrm{br}}$ column of Table 2. Each cell averages over 50 independent random seeds.

Problem	D	T	$\epsilon = 0.05$	$\epsilon = 0.075$	$\boldsymbol{\epsilon = 0.1}$ (default)	$\epsilon = 0.15$	$\epsilon = 0.2$
Branin	2	128	105.6 (84% / 0.311)	93.2 (84% / 0.312)	**87.0 (88% / 0.313)**	80.9 (88% / 0.314)	79.3 (88% / 0.315)
Rosenbrock	4	96	80.4 (98% / 0.00868)	80.4 (100% / 0.00868)	**80.4 (100% / 0.00868)**	80.4 (100% / 0.00868)	80.4 (100% / 0.00868)
Hartmann	3	64	53.2 (70% / 0.388)	53.2 (70% / 0.388)	**53.2 (70% / 0.388)**	53.2 (72% / 0.388)	53.0 (76% / 0.388)
GP 3×10^{-3}	2	128	49.8 (90% / 0.0210)	42.0 (92% / 0.0210)	**29.4 (94% / 0.0210)**	15.4 (98% / 0.0219)	12.1 (100% / 0.0233)
GP 5×10^{-3}	2	128	47.0 (88% / 0.0232)	44.6 (90% / 0.0232)	**38.2 (92% / 0.0232)**	22.6 (98% / 0.0245)	14.6 (100% / 0.0261)
GP 3×10^{-3}	6	512	149.1 (66% / 0.0511)	140.9 (72% / 0.0511)	**123.9 (72% / 0.0511)**	76.6 (90% / 0.0511)	51.6 (96% / 0.0511)
GP 5×10^{-3}	6	512	84.6 (44% / 0.0865)	84.1 (52% / 0.0865)	**79.9 (60% / 0.0865)**	67.4 (74% / 0.0865)	51.0 (88% / 0.0865)

Table 11: $\mathrm{UCB}_{\mathrm{br}}$ with our new bound (12) and with the existing bound (5), both evaluated with the same β_t along the same BO trajectories, on the well-specified $\mathrm{GP}^{\dagger}$ problems under the Matérn-$5/2$ kernel. Cells follow the format of Table 2; the number of the 50 seeds on which the certificate triggered is given in brackets.

Problem	D	T	New bound (12)	Existing bound (5)
$\mathbf{GP}^{\dagger}$ 3×10^{-3}	2	128	76.6 (100% / 0.000347) [48]	122.1 (100% / 0.000347) [18]
$\mathbf{GP}^{\dagger}$ 5×10^{-3}	2	128	115.2 (100% / 0.000899) [24]	128.0 (100% / 0.000899) [0]
$\mathbf{GP}^{\dagger}$ 3×10^{-3}	6	512	483.3 (100% / 0.00202) [12]	512.0 (100% / 0.00202) [0]
$\mathbf{GP}^{\dagger}$ 5×10^{-3}	6	512	509.9 (100% / 0.00297) [1]	512.0 (100% / 0.00297) [0]

Table 12: Robustness of $\mathrm{UCB}_{\mathrm{br}}$ to kernel-smoothness misspecification. The objective is a Matérn-$5/4$ draw while the surrogate assumes $\nu = 1.5$ or $\nu = 2.5$; in GP† the true lengthscale, outputscale, and noise are supplied, so the smoothness is the only misspecified quantity. Cells follow the format of Table 2; 50 seeds each.

Surrogate ν	Problem	D	T	NOSTOP	$\mathrm{UCB}_{\mathrm{br}}$	PRB
1.5	**GP**† 3×10^{-3}	2	128	128.0 (100% / 0.00255)	99.6 (100% / 0.00255)	6.6 (96% / 0.0269)
1.5	**GP** 3×10^{-3}	2	128	128.0 (100% / 0.00403)	61.2 (100% / 0.0115)	6.4 (94% / 0.0398)
2.5	**GP**† 3×10^{-3}	2	128	128.0 (100% / 0.00559)	66.2 (100% / 0.00645)	6.3 (96% / 0.0280)
2.5	**GP** 3×10^{-3}	2	128	128.0 (100% / 0.00460)	56.0 (100% / 0.0100)	6.2 (96% / 0.0335)

NeurIPS Paper Checklist

1. **Claims**

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

   Answer: [Yes]

   Justification: The main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

   Guidelines:

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

2. **Limitations**

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

   Answer: [Yes]

   Justification: The assumptions and algorithms that our work apply to are clearly stated.

   Guidelines:

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

3. **Theory assumptions and proofs**

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

   Answer: [Yes]

Justification: We provide assumptions and proofs.

Guidelines:

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

4. **Experimental result reproducibility**

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

   Answer: [Yes]

   Justification: We provide all the details and supplemental codes.

   Guidelines:

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

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

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

Answer: [Yes]

Justification: NA

Guidelines:

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

6. **Experimental setting/details**

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

   Answer: [Yes]

   Justification: NA

   Guidelines:

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

7. **Experiment statistical significance**

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

   Answer: [Yes]

   Justification: NA

   Guidelines:

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

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

8. **Experiments compute resources**

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

   Answer: [Yes]

   Justification: NA

   Guidelines:

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

9. **Code of ethics**

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

   Answer: [Yes]

   Justification: NA

   Guidelines:

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

10. **Broader impacts**

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

    Answer: [N/A]

    Justification: This paper focuses on the theoretical analysis and algorithm design for Bayesian optimization.

    Guidelines:

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

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

11. **Safeguards**

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

    Answer: [N/A]

    Justification: We do not have such data.

    Guidelines:

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

12. **Licenses for existing assets**

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

    Answer: [Yes]

    Justification: NA

    Guidelines:

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

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

13. **New assets**

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

    Answer: [Yes]

    Justification: NA

    Guidelines:

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

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

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

    Answer: [N/A]

    Justification: NA

    Guidelines:

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

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

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

    Answer: [N/A]

    Justification: NA

    Guidelines:

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

16. **Declaration of LLM usage**

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

Answer: [N/A]

Justification: NA

Guidelines:

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