Bayesian Optimization with Noise-Free Observations: Improved Regret Bounds via Random Exploration

Duke University

Hwanwoo Kim & Daniel Sanz-Alonso Contact: hwanwoo.kim@duke.edu

The University of Chicago 1890

Overview

- We introduce two novel Bayesian Optimization (BO) algorithms almost **attaining the optimal simple regret bounds in [Bul11].**
- Our algorithms (GP-UCB+ and EXPLOIT+) share the **simplicity and ease of implementation** of the standard BO algorithm. In addition, EXPLOIT+ achieves **competitive empirical performance** to existing BO algorithms **without any hyperparameter tuning**.

Background

Goal Maximize a function $f : \mathcal{X} \to \mathbb{R}$, where

- Mathematical expression of f is not necessarily known.
- f is not necessarily convex nor differentiable.
- Only source of information about f is through its evaluations, which are typically expensive.

Gaussian Process Denote generic query locations by $X_t = \{x_1, \ldots, x_t\}$ and the corresponding noise-free observations by $F_t = [f(x_1), \ldots, f(x_t)]^\top$. Given a positive-definite kernel function k, Gaussian process interpolation with a prior $\mathcal{GP}(0, k)$ yields the following posterior predictive mean and variance:

$$\mu_t(x) = k_t(x)^\top K_{tt}^{-1} F_t,$$
$$\sigma_t^2(x) = k(x, x) - k_t(x)^\top K_{tt}^{-1} k_t(x),$$

where $k_t(x) = [k(x, x_1), \ldots, k(x, x_t)]^\top$ and K_{tt} is a $t \times t$ matrix with entries $(K_{tt})_{i,j} = k(x_i, x_j)$.

Gaussian Process Upper Confidence Bound (GP-UCB)

Input: Kernel k; Total number of iterations T; Initial query locations X_0; Initial noise-free observations F_0

For $t = 1, \cdots, T$:

1. Compute posterior mean/variance using all query locations and function evaluations (X_{t-1}, F_{t-1}).
2. Obtain $x_t = \arg\max \mu_{t-1}(x) + \beta_t \sigma_{t-1}(x)$, for $\beta_t \in \mathbb{R}_+$.
3. Set $X_t = X_{t-1} \cup \{x_t\}$, $F_t = F_{t-1} \cup \{f(x_t)\}$.

Output: $\arg\max_{x \in X_T} f(x)$.

Performance Metric: The simple regret is defined as

$$s_t = \max_{\mathcal{X}} f(x) - \max_{t=1,\cdots,T} f(x_t).$$

Optimal Regret Bounds

Under RKHS assumption: The best possible BO strategy yields

$$s_T = \Theta\left(T^{-\frac{\nu}{d}}\right)$$

for a Matérn kernel with smoothness parameter $\nu > 0$ [Bul11]. The popular GP-UCB algorithm, with $\beta_t = \|f\|_{\mathcal{H}_k}$, yields

$$s_T = \begin{cases} \mathcal{O}\left(T^{-\frac{\nu}{2\nu+d}} \log^{\frac{\nu}{2\nu+d}} T\right), \\ \mathcal{O}\left(T^{-\frac{1}{2}} \log^{\frac{d+1}{2}} T\right), \end{cases}$$

for Matérn and squared exponential kernels [LYT19].

GP-UCB+ & EXPLOIT+

Improved Exploration via Random Sampling

Input: Kernel k; Total number of iterations T; Initial query locations X_0; Initial noise-free observations F_0; Prob distribution P on $\mathcal{X}$.

For $t = 1, \cdots, T$:

1. Compute posterior mean/variance using all query locations and function evaluations (X_{t-1}, F_{t-1}).
2. Obtain
   - GP-UCB+: $x_t = \arg\max \mu_{t-1}(x) + \|f\|_{\mathcal{H}_k} \sigma_{t-1}(x)$.
   - EXPLOIT+: $x_t = \arg\max \mu_{t-1}(x)$.
3. Sample $\tilde{x}_t \sim P$.
4. Set $X_t = X_{t-1} \cup \{x_t, \tilde{x}_t\}$, $F_t = F_{t-1} \cup \{f(x_t), f(\tilde{x}_t)\}$.

Output: $\arg\max_{x \in X_T} f(x)$.

Under RKHS assumption: For a probability measure P with a strictly positive density, both algorithm yields

$$\mathbb{E}_P[s_T] = \begin{cases} \mathcal{O}\left(T^{-\frac{\nu}{d}+\varepsilon}\right), \\ \mathcal{O}\left(\exp\left(-CT^{\frac{1}{d}-\varepsilon}\right)\right), \end{cases}$$

for Matérn and squared exponential kernels [KSA24]. Here, $\varepsilon > 0$ can be arbitrarily small.

Remark The exact rate by replacing the random sampling step in GP-UCB+ and EXPLOIT+ with a more computationally expensive quasi-uniform sampling scheme.

Numerical Experiments

- We compare with other BO algorithms: Expected Improvement (EI), Probability of Improvement (PI), Pure Exploitation (EXPLOIT), and GP-UCB.
- The new algorithms need two noise-free observations per iteration, but the methods we compare with only need one; we run the new algorithms for half as many iterations to ensure a fair comparison.

Benchmark Ackley function (10-dim):

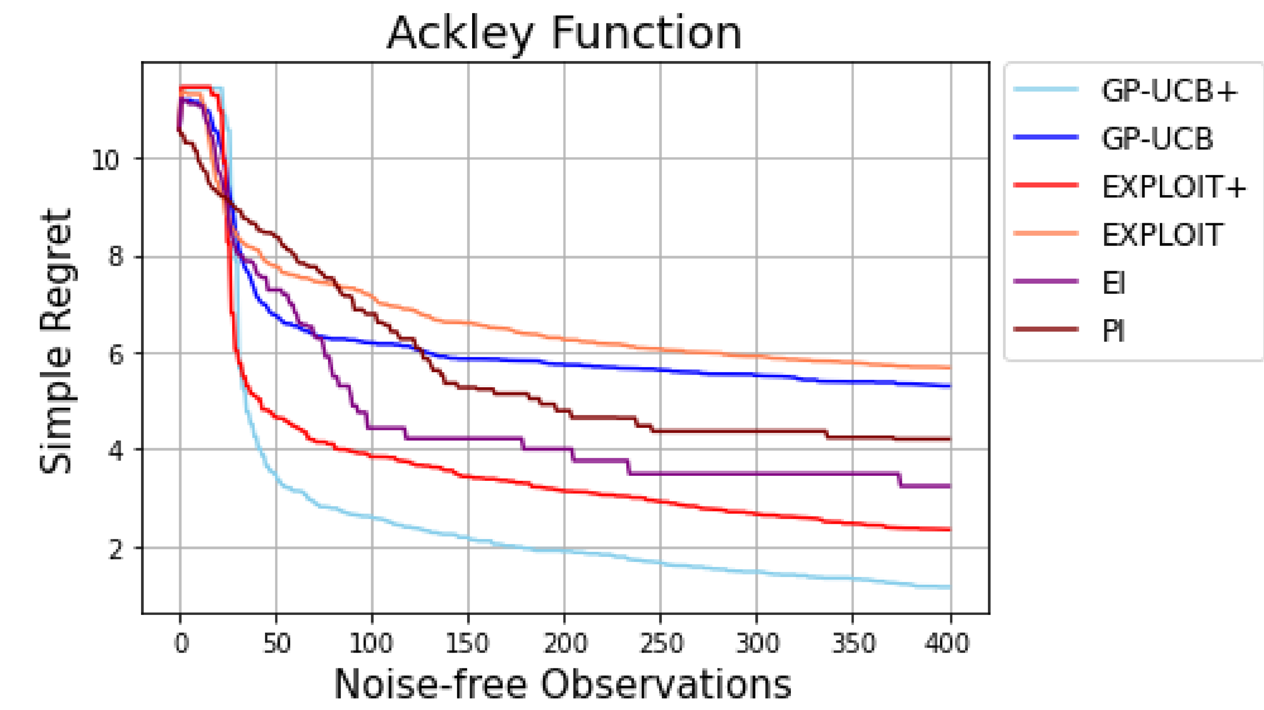

Maximizing the range of garden sprinkler (8-dim)

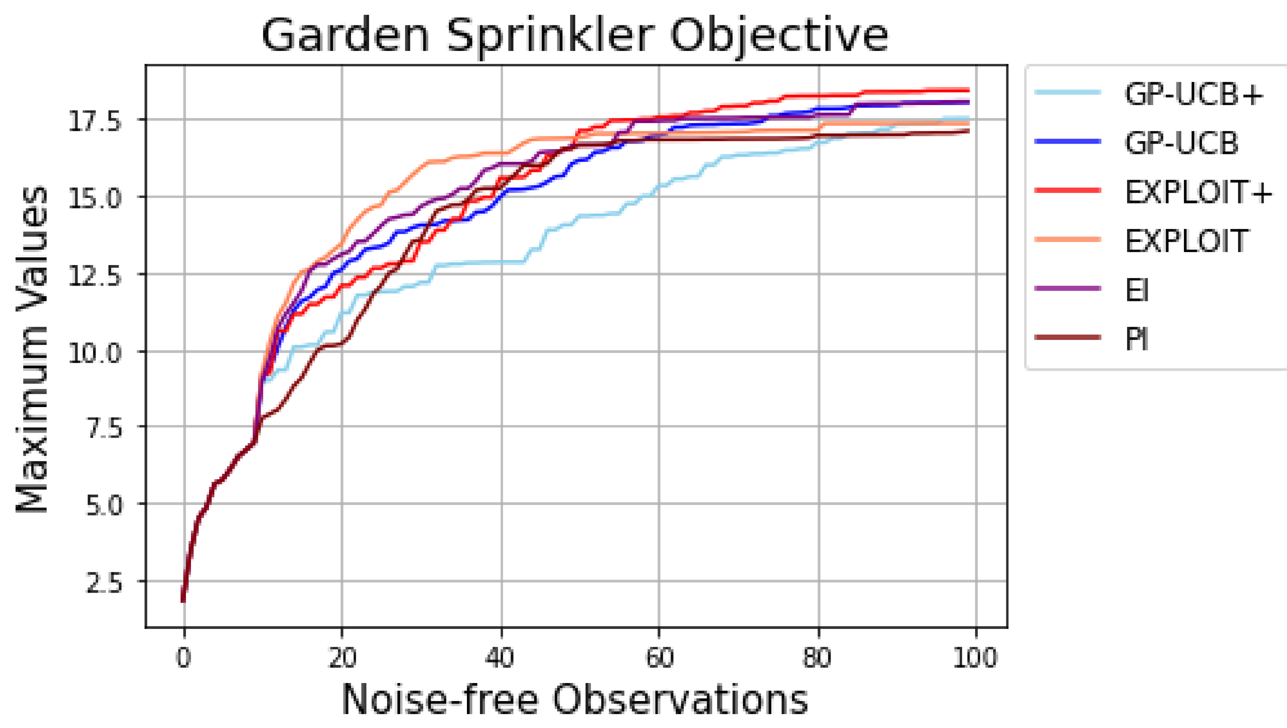

References

[Bul11] Adam D Bull. Convergence rates of efficient global optimization algorithms. *Journal of Machine Learning Research*, 12(10), 2011.

[KSA24] Hwanwoo Kim and Daniel Sanz-Alonso. Enhancing gaussian process surrogates for optimization and posterior approximation via random exploration. *arXiv preprint arXiv:2401.17037*, 2024.

[LYT19] Yueming Lyu, Yuan Yuan, and Ivor W Tsang. Efficient batch black-box optimization with deterministic regret bounds. *arXiv preprint arXiv:1905.10041*, 2019.