Cost-aware Stopping for Bayesian Optimization

Qian Xie[1,*] **Linda Cai**[2,*] **Alexander Terenin**[1] **Peter I. Frazier**[1] **Ziv Scully**[1]

[1]Cornell University
[2]University of California, Berkeley

Abstract In automated machine learning and other applications of Bayesian optimization, deciding when to stop evaluating expensive black-box functions is practically important. Existing adaptive stopping rules lack guarantees ensuring they stop before incurring excessive function evaluation costs. We propose a *cost-aware stopping rule* for Bayesian optimization that is free of heuristic tuning and grounded in a theoretical connection to state-of-the-art cost-aware acquisition functions, namely the Pandora's Box Gittins Index (PBGI) and log expected improvement per cost. We prove a theoretical guarantee bounding the expected cumulative evaluation cost incurred by our stopping rule when paired with these two acquisition functions. Empirical results on hyperparameter optimization and neural architecture size search show that combining our stopping rule with the PBGI acquisition function consistently matches or outperforms other acquisition-function–stopping-rule pairs in terms of *cost-adjusted simple regret*, a metric capturing trade-offs between solution quality and cumulative evaluation cost.

1 Introduction

Bayesian optimization is a framework designed to efficiently solve optimization problems involving expensive-to-evaluate black-box functions, particularly in hyperparameter tuning [13]. It iteratively updates a probabilistic model of the objective based on current data and selects new evaluation points via an acquisition function balancing exploration and exploitation.

This work considers the *cost-aware* setting, where each evaluation incurs an input-dependent costs. For example, hyperparameter tuning in cloud environments incurs CPU/GPU hour costs depending on chosen hyperparameters, creating a tradeoff between the compute cost of training and the business value of deploying an accurate model. We therefore aim to design adaptive stopping rules that optimizes expectation of *cost-adjusted simple regret*, which defined as the sum of simple regret and the cumulative cost of data collection.

Several stopping rules have been proposed for Bayesian optimization. Simple heuristics—such as fixing a maximum number of iterations or stopping when improvement falls below a threshold—are widely used in practice [9, 10], but can either stop too early or lead to unnecessary evaluations. Other more advanced approaches have been proposed [4, 7, 9–11], but none of them allow for varying function evaluation costs.

In this work, we study how to design cost-aware stopping rules, motivated by two primary factors. First, state-of-the-art cost-aware acquisition functions such as the Pandora's Box Gittins Index (PBGI) [16] and log expected improvement per cost (LogEIPC) [1] have not yet been studied in the adaptive stopping setting. This is important because—as our experiments in Section 4 will show—for best performance, one should pair different acquisition functions with different stopping rules. Second, while certain stopping rules, such as UCB-LCB [10], are guaranteed to achieve a low simple regret, they are not necessarily guaranteed to do so with low evaluation costs. This is important because—as our experiments will show—UCB-LCB will often incur high evaluation costs, resulting in a high cost-adjusted simple regret.

Our main methodological contribution is a stopping rule that is provably Bayesian-optimal in the independent evaluation case, and empirically strong in the general case. This stopping rule is constructed based on ideas from Pandora's Box theory, which forms the theoretical foundation for the recently-proposed PBGI acquisition function. It also aligns closely with an existing stopping rule in the uniform-cost setting, namely one based on expected improvement [11]; our work thus extends this stopping rule to the varying-cost setting. Our specific contributions are as follows:

1. *A Novel Cost-Aware Stopping Criterion.* We propose an adaptive stopping rule derived naturally from Pandora's box theory, establishing a unified and principled framework applicable to cost-aware Bayesian optimization.
2. *Theoretical Guarantee.* We prove in Theorem 1 that our stopping rule, when paired with the PBGI or LogEIPC acquisition functions, satisfies a theoretical upper bound on expected cumulative cost up to stopping, which constitutes the first theoretical guarantee of this type for any adaptive stopping rule for Bayesian optimization.
3. *Empirical Validation.* Our experiments show that combining the PBGI acquisition function with our proposed stopping rule consistently matches or outperforms other combinations of acquisition functions and stopping rules on two AutoML benchmarks [3, 17].

2 Bayesian Optimization and Adaptive Stopping

In classic finite-horizon Bayesian optimization the goal is to adaptively evaluate an unknown function $f : X \to \mathbb{R}$, which is assumed to be sampled from a Gaussian process (GP) [12], at T points $x_1, \ldots, x_T$, with the general aim of *minimizing* the lowest objective value observed. A specific metric that measures this is *expected simple regret* [5, Sec. 10.1].

We study a *cost-aware* variant of Bayesian optimization that differs in three ways from the finite-horizon version above. First, *different points have different evaluation costs*, with a cost function $c : X \to \mathbb{R}^+$ mapping each point x to its evaluation cost $c(x)$. In this abstract, we assume the cost function is known to the algorithm, but it is possible to extend our method to unknown costs. Second, instead of a fixed number of time steps, *the algorithm decides when to stop.* That is, in addition to adaptively choosing the points x_t, it also adaptively chooses a *stopping time* $\tau \geq 0$. Specifically, after inspecting x_t and observing $y_t = f(x_t)$, the algorithm must decide whether to stop now—in which case $\tau = t$—or to continue. Third, instead of the evaluation costs being part of a constraint, *the costs are incorporated into the objective.* Specifically, we aim to optimize a metric we call *expected cost-adjusted simple regret*, defined by

$$\mathcal{R}_c = \mathbb{E}\Bigg[\underbrace{\min_{1 \leq t \leq T} y_t - \inf_{x \in X} f(x)}_{\text{simple regret}} + \underbrace{\sum_{t=1}^{T} c(x_t)}_{\text{cumulative cost}} \Bigg]. \tag{1}$$

2.1 Acquisition Functions

Most Bayesian optimization algorithms select their next points using an *acquisition function*, denoted $\alpha_t(x)$, which gives a "rating" to each point x at each time t. When using an acquisition function, the point x_{t+1} is chosen to maximize or minimize of the acquisition function $\alpha_t(x)$.

In this work, we consider a number of acquisition functions—chiefly, *log expected improvement per cost (LogEIPC)* [1] and the *Pandora's Box Gittins Index (PBGI)* [16]. These, respectively, are

$$\alpha_t^{\text{LogEIPC}}(x) = \log \frac{\mathrm{EI}_{f \mid x_{1:t}, y_{1:t}}\left(x; y^*_{1:t}\right)}{c(x)} \qquad \text{(choose maximizer)} \tag{2}$$

$$\alpha_t^{\text{PBGI}}(x) = g \quad \text{where } g \text{ solves} \quad \mathrm{EI}_{f \mid x_{1:t}, y_{1:t}}(x; g) = c(x) \qquad \text{(choose minimizer)} \tag{3}$$

where $\mathrm{EI}_\psi(x; y) = \mathbb{E}\left[\max(0, y - \psi(x))\right]$ is the expected improvement at point x with respect to some random function $\psi : X \to \mathbb{R}$ and a comparison value y; $y^*_{1:t} = \min_{1 \leq s \leq t} y_s$ is the best observed value up to time t; and $c(x)$ is the evaluation cost at point x. In the classical uniform-cost setting, we also consider the classical *lower confidence bound (LCB)* and *Thompson sampling (TS)* acquisition functions: see Garnett [5] for additional details.

2.2 Adaptive Stopping Rules

A number of stopping rules have been proposed for Bayesian optimization. These include methods that stop when they estimate that convergence has occurred [2], when acquisition function values fail to cross a threshold [4, 7–9, 11], and when they estimate that regret is low [7, 10, 15]. We use several of these methods as baselines in our experiments in Section 4. However, to the best of our knowledge, all of these prior methods do not explicitly account for varying evaluation costs.

3 A Stopping Rule Based on the Gittins Index

In this work, we propose a new stopping rule tailored for two state-of-the-art acquisition functions used in cost-aware Bayesian optimization: LogEIPC and PBGI, introduced in Section 2.

The inspiration for our stopping rule comes from the *Pandora's Box* problem [14], which Xie et al. [16] observed can be viewed as a (generalization of a) special case of cost-aware Bayesian optimization with two significant simplifications: (1) there are only finitely many points $X = \{1, \dots, N\}$; and (2) the objective function has no correlations, i.e. $f(x)$ and $f(y)$ are independent if $x \neq y$. In this special case, using the PBGI acquisition function together with the following stopping rule is *Bayes-optimal* for cost-adjusted regret:

$$\text{stop at } \tau = t \text{ if} \quad \min_{x \in X} \alpha_t^{\mathrm{PBGI}}(x) \geq y_{1:t}^*. \tag{4}$$

Our proposal is to use this stopping rule in general Bayesian optimization. The intuition behind it is that the PBGI acquisition value $\alpha_t^{\mathrm{PBGI}}(x)$, called the *Gittins index* of point x [6], is in some sense the "fair value" of that point. If $x = x_t$ has been evaluated, its Gittins index is simply its revealed value $y_t = f(x_t)$. If x has not been evaluated, its Gittins index summarizes the utility of being able to reveal a new random value $f(x)$ by paying cost $c(x)$. From this view, (4) says to *stop when the best Gittins index is at an evaluated point.*

Connection to LogEIPC. Using the fact that the PBGI acquisition function is defined in terms of a root-finding problem related to the expected improvement function, one can show that (4) is equivalent to a stopping rule based on LogEIPC:

$$\text{stop at } \tau = t \text{ if} \quad \max_{x \in X} \alpha_t^{\mathrm{LogEIPC}}(x; y_{1:t}^*) \leq 0. \tag{5}$$

Unwrapping the logarithm, this is in turn equivalent to

$$\text{stop at } \tau = t \text{ if} \quad \mathrm{EI}\big(x; y_{1:t}^*\big) \leq c(x) \quad \text{for all } x \in X \setminus \{x_1, \dots, x_t\}. \tag{6}$$

This gives an intuitive interpretation of the stopping rule: *stop when no unevaluated point's expected improvement outweighs its evaluation cost.* We call the stopping rule given by the equivalent conditions (4), (5), and (6) the *PBGI/LogEIPC stopping rule.*

Bounding the Cost up to Stopping. The following result bounds the expected cumulative cost incurred when using it in conjunction with the PBGI or LogEIPC acquisition functions.

Theorem 1. *Consider optimizing $f \sim \mathcal{GP}(\mu, K)$ on a compact, non-empty domain X, and let c be the cost function. Suppose at time 1, a fixed point x_1 is evaluated, and thereafter the acquisition function is either PBGI or LogEIPC. Let τ be the stopping time under the PBGI/LogEIPC stopping rule from (4). Then the expected cumulative cost starting at time 2 is bounded by*

$$\mathbb{E}\left[\sum_{t=2}^{\tau} c(x_t)\right] \leq \mu - \mathbb{E}\Big[\min_{x \in X} f(x)\Big]. \tag{7}$$

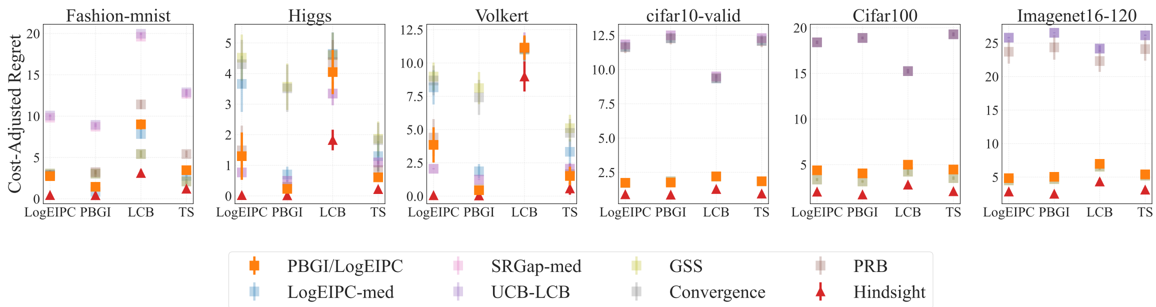

Figure 1: Cost-adjusted simple regret of different acquisition-function–stopping-rule pairs on LCBench (left) and NATS-Bench (right). Squares show means, and bars show two times standard error.

While there are several results bounding the expected simple regret after stopping under various stopping rules [7, 10, 11, 15], Theorem 1 is, to the best of our knowledge, the first theoretical guarantee on the *expected cumulative cost*. We emphasize that our result requires using the PBGI or LogEIPC acquisition functions (or, potentially, other related acquisition functions), reflecting the importance of studying stopping-rule–acquisition-function pairs together.

4 Experiments on AutoML Benchmarks

Acquisition functions. We consider four common acquisition functions that were discussed in Section 2: *log expected improvement per cost (LogEIPC)*, *Pandora's Box Gittins Index (PBGI)*, *lower confidence bound (LCB)*, and *Thompson sampling (TS)*—chosen for their competitive performance, computational efficiency, and close connections to existing stopping rules.

Baselines. We compare our PBGI/LogEIPC stopping rule against several stopping rules from prior work: *UCB-LCB* [10], *LogEIPC-med* [7], *SRGap-med* [7], and *PRB* [15]. Additionally, we include two simple heuristics used in practice: *convergence* (stop when the best value remains unchanged for a fixed number of iterations) and *GSS* (stop when the improvement is no longer significant relative to the inter-quartile range) used in Meta's Ax platform [2]. Each experiment is repeated with 50 random seeds to assess variability. Each trial (run with a distinct random seed) is capped at 200 iterations; if a stopping rule is not triggered within this limit, the stopping time is set to the cap.

Benchmarks. We test our algorithms on empirical objective functions from two AutoML benchmarks: *LCBench* [17], a hyperparameter tuning data set; and *NATS-Bench* [3], a neural architecture search benchmark. Each of these benchmarks feature data about thousands of models trained on several tasks; we show results for three tasks from each benchmark. While we use a GP model, specifically one with a Matérn 5/2 kernel, to capture correlations between similar data points, we restrict our search space to points that correspond to configurations in the benchmark data, so our acquisition function optimization is over a discrete (but large) space.

During acquisition, we use validation error for the objective value, but when computing regret, we use test error for the objective value. For the evaluation cost, we use an estimated training cost based on the number of model parameters for LCBench or number of FLOPs (floating point operations) for NATS-Bench.

Experiment results. Figure 1 compares cost-adjusted regret for acquisition function and stopping rule pairings. Overall, our PBGI/LogEIPC stopping rule consistently performs competitively in terms of cost-adjusted simple regret—often close to the hindsight optimum—especially when paired with the PBGI acquisition function. When combined with either the LogEIPC or PBGI acquisition function, it also avoids excessive spending, in line with Theorem 1, whereas other acquisition-based and regret-based methods do not, particularly on NATS-Bench.

References

[1] Sebastian Ament, Samuel Daulton, David Eriksson, Maximilian Balandat, and Eytan Bakshy. Unexpected improvements to expected improvement for bayesian optimization. *Advances in Neural Information Processing Systems*, 36:20577–20612, 2023.

[2] Eytan Bakshy, Lili Dworkin, Brian Karrer, Konstantin Kashin, Benjamin Letham, Ashwin Murthy, and Shaun Singh. Ae: A domain-agnostic platform for adaptive experimentation. In *Conference on neural information processing systems*, pages 1–8, 2018.

[3] Xuanyi Dong, Lu Liu, Katarzyna Musial, and Bogdan Gabrys. Nats-bench: Benchmarking nas algorithms for architecture topology and size. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3634–3646, 2021.

[4] Peter Frazier and Warren B Powell. The knowledge-gradient stopping rule for ranking and selection. In *2008 Winter Simulation Conference*, pages 305–312. IEEE, 2008.

[5] Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.

[6] John C. Gittins, Kevin D. Glazebrook, and Richard R. Weber. *Multi-Armed Bandit Allocation Indices*. Wiley, 2011.

[7] Hideaki Ishibashi, Masayuki Karasuyama, Ichiro Takeuchi, and Hideitsu 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.

[8] Marco Locatelli. Bayesian algorithms for one-dimensional global optimization. *Journal of Global Optimization*, 10:57–76, 1997.

[9] 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. *arXiv preprint arXiv:1511.07827*, 2015.

[10] Anastasia Makarova, Huibin Shen, Valerio Perrone, Aaron Klein, Jean Baptiste Faddoul, Andreas Krause, Matthias Seeger, and Cedric Archambeau. Automatic termination for hyperparameter optimization. In *International Conference on Automated Machine Learning*, pages 7–1. PMLR, 2022.

[11] Vu Nguyen, Sunil Gupta, Santu Rana, Cheng Li, and Svetha Venkatesh. Regret for expected improvement over the best-observed value and stopping condition. In *Asian conference on machine learning*, pages 279–294. PMLR, 2017.

[12] Carl Edward Rasmussen and Christopher KI Williams. *Gaussian Processes for Machine Learning*. Cambridge University Press, 2006.

[13] Jasper Snoek, Hugo Larochelle, and Ryan P Adams. Practical bayesian optimization of machine learning algorithms. *Advances in Neural Information Processing Systems*, 2012.

[14] Martin L Weitzman. Optimal search for the best alternative. *Econometrica*, 1979.

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

[16] Qian Xie, Raul Astudillo, Peter Frazier, Ziv Scully, and Alexander Terenin. Cost-aware bayesian optimization via the pandora's box gittins index. *Advances in Neural Information Processing Systems*, 37:115523–115562, 2024.

[17] Lucas Zimmer, Marius Lindauer, and Frank Hutter. Auto-pytorch tabular: Multi-fidelity metalearning for efficient and robust autodl. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):3079 – 3090, 2021.