Cost-aware Stopping for Bayesian Optimization

Qian Xie*[1], Linda Cai*[2], Alexander Terenin[1], Peter Frazier[1], and Ziv Scully[1]
[1]Cornell University [2]University of California, Berkeley

Cost-aware Bayesian Optimization with Adaptive Stopping

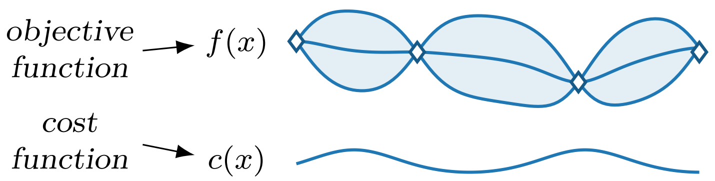

Cost-adjusted simple regret: $\underbrace{\min_{1\le t\le\tau} f(x_t) - \inf_{x\in X} f(x)}_{\text{simple regret}} + \underbrace{\sum_{t=1}^{\tau} c(x_t)}_{\text{cumulative cost}}$

Goal: *Adaptively* evaluate $x_1, x_2, \cdots$ and stop at time τ to minimize expected cost-adjusted simple regret.

Existing Acquisition Rules

Expected Improvement (EI): Evaluate the point with maximum expected improvement over the current best observed value $y^*_{1:t}$

$$\alpha_t^{\mathrm{EI}}(x) = \mathrm{EI}_{f|x_{1:t},y_{1:t}}(x; y^*_{1:t}) \quad \text{where} \quad \mathrm{EI}_\psi(x;y) = \mathbb{E}\left[(y-\psi(x))^+\right].$$

Expected Improvement per Cost (EIPC): Evaluate the point with maximum expected improvement divided by cost

$$\alpha_t^{\mathrm{EIPC}}(x) = \alpha_t^{\mathrm{EI}}(x)/c(x).$$

Pandora's Box Gittins Index (PBGI): (Xie et al., 2024) Evaluate the point with the minimum index given by

$$\alpha_t^{\mathrm{PBGI}}(x) = g \qquad \text{where } g \text{ solves} \qquad \mathrm{EI}_{f|x_{1:t},y_{1:t}}(x;g) = c(x).$$

Intuition: (Weitzman, 1979) Is opening the box worth the cost?

Decision	Disutility
Open	$\mathbb{E}\min(f,g)+c$
Not Open	g

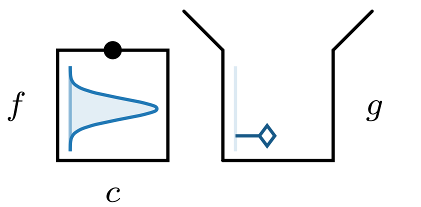

Should one open the closed box? Depends on outside option g!
If *both* opening and not are optimal: g is a *fair value*.
α_t^{PBGI}: pick points according to their fair values.

Other acquisition rules: lower confidence bound (LCB), Thompson sampling (TS), ...

PBGI/EIPC Stopping Rule

Existing EI stopping rule: (Nguyen et al., 2017) Under the *uniform-cost* setting, stop when $\alpha_t^{\mathrm{EI}}(x) \le c$ where c is the unit cost.

Our PBGI/EIPC stopping rule: Stopping when the PBGI index at *every* unevaluated point is at least the current best observed value

$$\min_{x\in X\setminus\{x_1,\dots,x_t\}} \alpha_t^{\mathrm{PBGI}}(x) \ge y^*_{1:t}. \tag{1}$$

Bayesian-optimal when paired with the PBGI acquisition rule under the independent-value setting.

Behavior Illustration

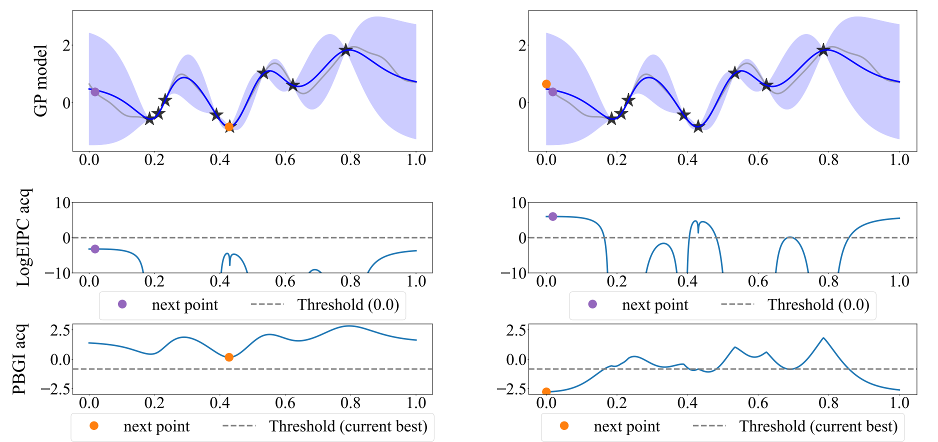

(a) $c(x) \equiv 0.1$ (b) $c(x) \equiv 0.0001$

Theoretical Guarantee

Theorem 1 (No worse than stopping-immediately)
When optimizing a random function f with a constant prior mean, our stopping rule with PBGI or EIPC achieves expected cost-adjusted regret no worse than stopping immediately after initial evaluation.

$$\mathbb{E}\left[y^*_{1:\tau} - \min_{x\in X} f(x) + \sum_{t=1}^{\tau} c(x_t)\right] \le \mathbb{E}\left[y_1 - \min_{x\in X} f(x) + c(x_1)\right].$$

Key proof idea: Using our stopping rule, both PBGI and EIPC are guaranteed to evaluate only points whose one-step expected improvement is worth the evaluation cost before stopping.

Implication: Matches the best we can hope for in the worst case, and avoids over-spending — properties many cost-unaware rules lack.

Performance

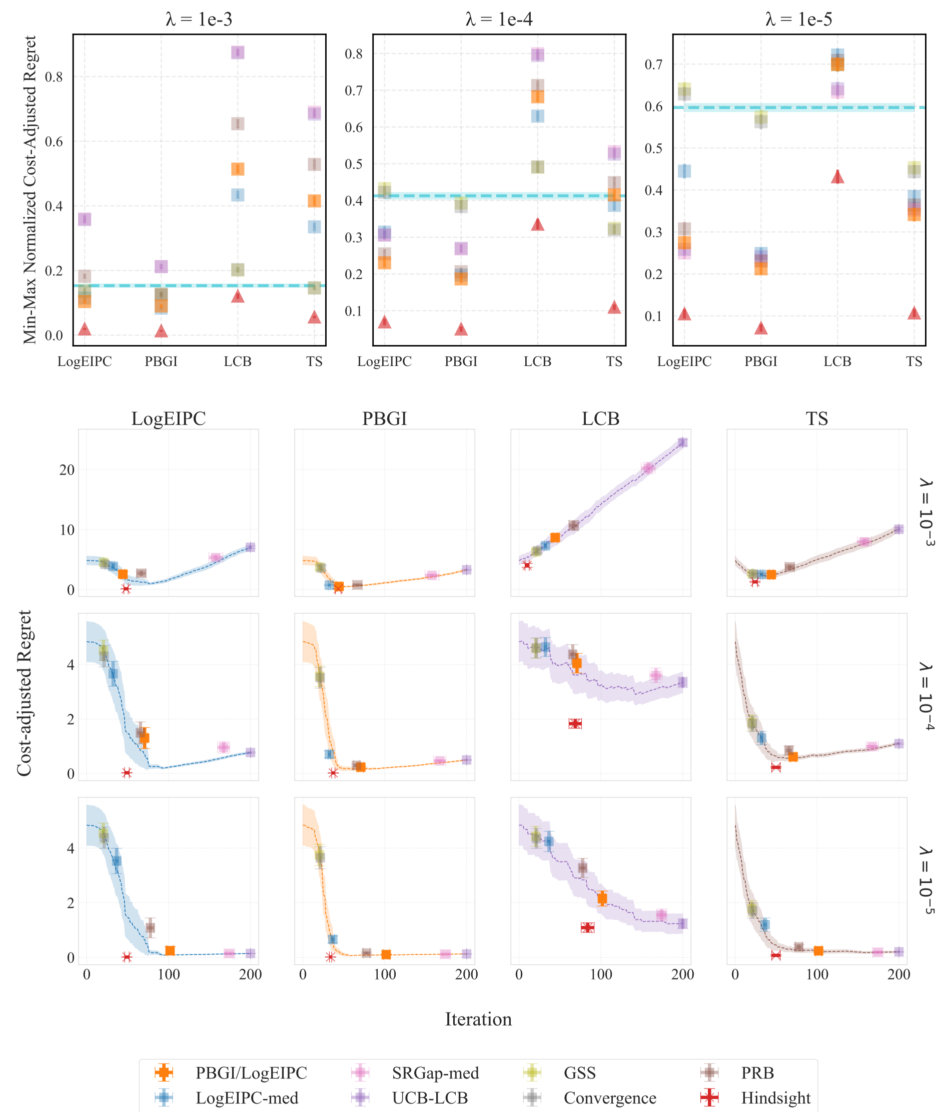

Computation Time