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

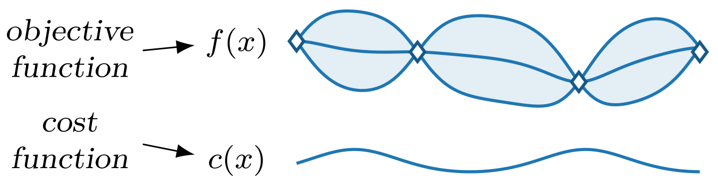

Cost-adjusted simple regret: $\underbrace{\min_{1\le t\le\tau} f(x_t) - \inf_{x\in X} f(x)}_{\text{simple regret}} + \underbrace{\sum_{t=1}^{\tau} c(x_t)}_{\text{cumulative cost}}$

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

Existing Adaptive Stopping Rules

Simple heuristics: stop when the best observed value remains unchanged or improvement is not statistically significant.

Acquisition-based: stop when PI, EI or KG falls below a threshold.

Regret-based: stop when regret bounds drop below a threshold (with some probability), e.g., UCB-LCB (Makarova et al., 2022).

SOTA Cost-aware Acquisition Rules

LogEIPC: (Snoek et al., 2012) evaluate the point with the large expected improvement relative to evaluation cost

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

PBGI: (Xie et al., 2024) evaluate the point with the smallest index given by

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

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

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

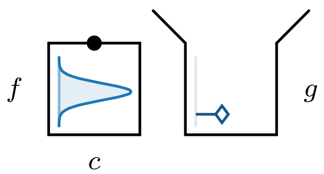

Should one open the closed box? Depends on the outside option g!
If *both* opening and not are optimal: g is a *fair value.*
α_t^{PBGI}: pick points according to their fair values.

PBGI/LogEIPC Stopping Rule

Our PBGI/LogEIPC stopping rule:

$$\min_{x\in X\setminus\{x_1,\ldots,x_t\}} \alpha_t^{\text{PBGI}}(x) \ge y_{1:t}^* \iff \max_{x\in X\setminus\{x_1,\ldots,x_t\}} \alpha_t^{\text{LogEIPC}}(x) \le 0.$$

- Stop when the best fair value is already evaluated (equiv. to *no unevaluated point's expected improvement is worth its cost*).
- Recovers an existing EI stopping rule under the unit-cost setting (Nguyen et al., 2017; Zhou et al., 2024).
- *Bayesian-optimal* when paired with the PBGI acquisition rule under the independent-value setting.

Behavior Illustration

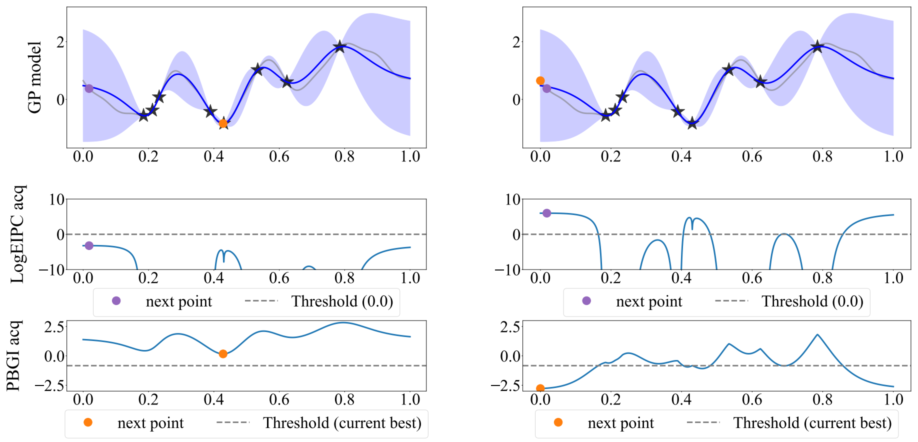

(a) $c(x) \equiv 0.1$ (b) $c(x) \equiv 0.0001$

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.

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

- Matches the best we can hope for in the worst case.
- *Avoids over-spending*, a property many cost-unaware rules lack.

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

Performance

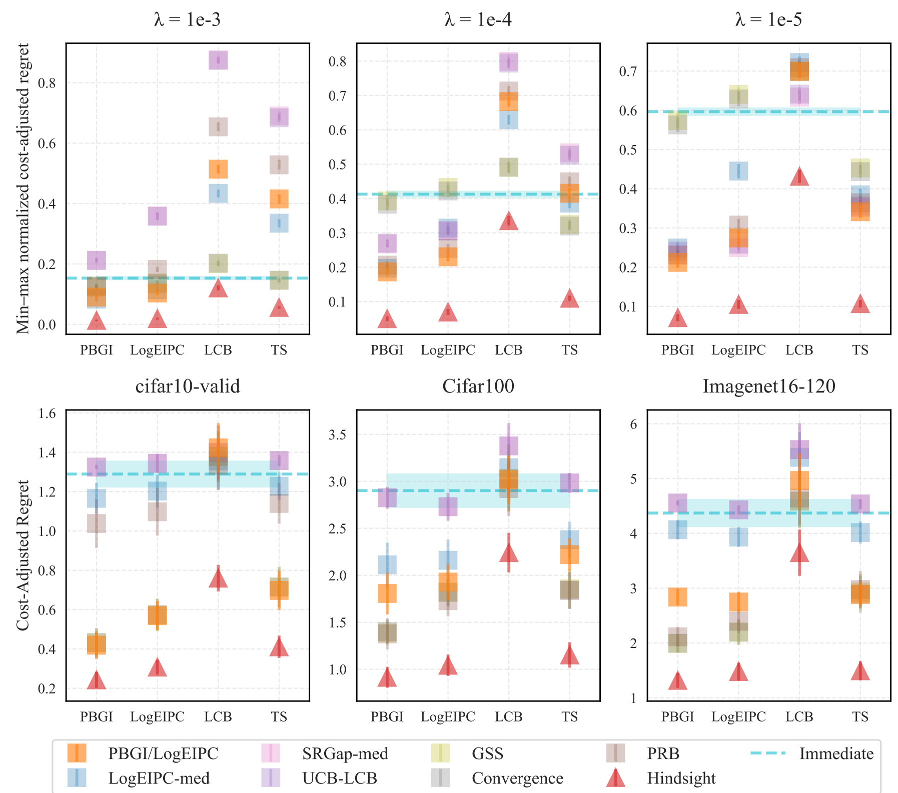

LCBench Rank Frequency of Our Stopping Rule

λ	Acq	Top 1		Top 2		Top 3	
		Large	All	Large	All	Large	All
10^{-3}	PBGI	20.0%	20.0%	70.0%	48.6%	**80.0%**	60.0%
	LogEIPC	40.0%	31.4%	65.0%	45.7%	**70.0%**	62.8%
10^{-4}	PBGI	40.0%	28.6%	70.0%	51.5%	**80.0%**	62.9%
	LogEIPC	30.0%	22.9%	50.0%	45.8%	**85.0%**	68.7%
10^{-5}	PBGI	50.0%	28.6%	70.0%	51.5%	**90.0%**	74.4%
	LogEIPC	20.0%	20.0%	50.0%	48.6%	**85.0%**	71.5%

LCBench has 35 datasets; large datasets have > 10,000 instances

Takeaways: (1) Our stopping rule paired with PBGI or LogEIPC ranks Top-3 in **70–90**% of large LCBench tasks; (2) The matched PBGI pair is the most robust default under model misspecification.

Conclusions

- Connection to Pandora's box gives a shared cost-aware stopping rule for PBGI and LogEIPC.
- Our stopping rule paired with PBGI or LogEIPC has a safeguard theoretical guarantee and strong empirical performance.