"Cost-aware Bayesian Optimization via the Pandora's Box Gittins Index." *NeurIPS'24.*

Gittins Indices for Cost-aware and Freeze-thaw Bayesian Optimization

Ongoing work

Qian Xie (Cornell ORIE)

Admission to Candidacy Exam ("A Exam")

Collaborators

Cost-aware Bayesian Optimization

Raul Astudillo
Caltech Postdoc

Peter Frazier
Committee member

Ziv Scully
Advisor

Alexander Terenin
Cornell RAP

Theo Brown
UCL PhD

Freeze-thaw Bayesian Optimization

World of Parameter Optimization

Hyperparameter tuning:
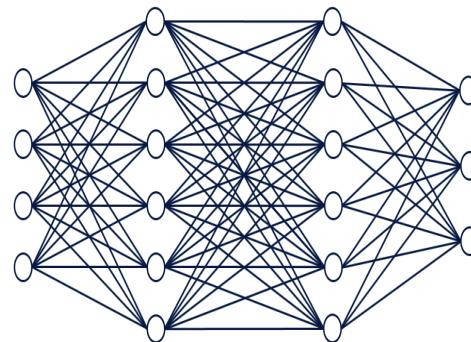
Training parameters $\longrightarrow$ $\longrightarrow$ Accuracy

Control optimization:

Control variables $\longrightarrow$ $\longrightarrow$ Reward
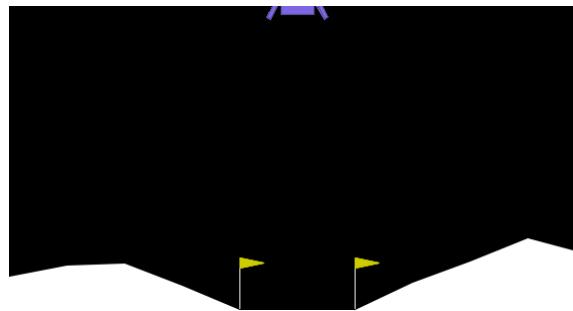
Plasma physics:
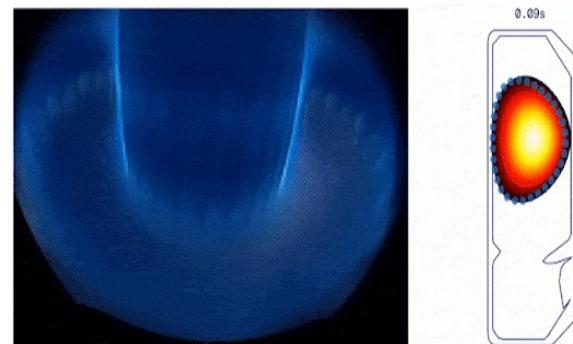
Reactor parameters $\longrightarrow$ $\longrightarrow$ Stability

World of Parameter Optimization

Black-box optimization:

Input parameters $\longrightarrow$ $\longrightarrow$ Performance metric
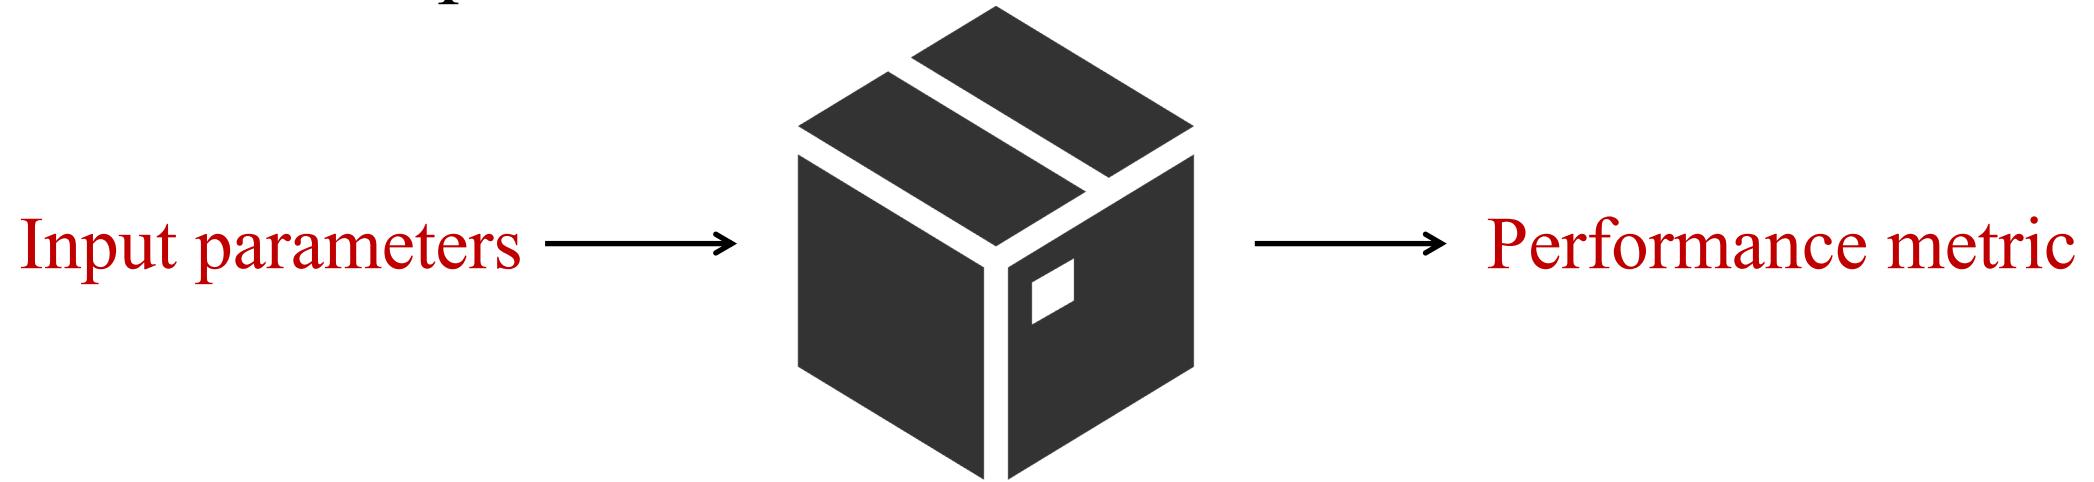
World of Parameter Optimization

Black-box function:
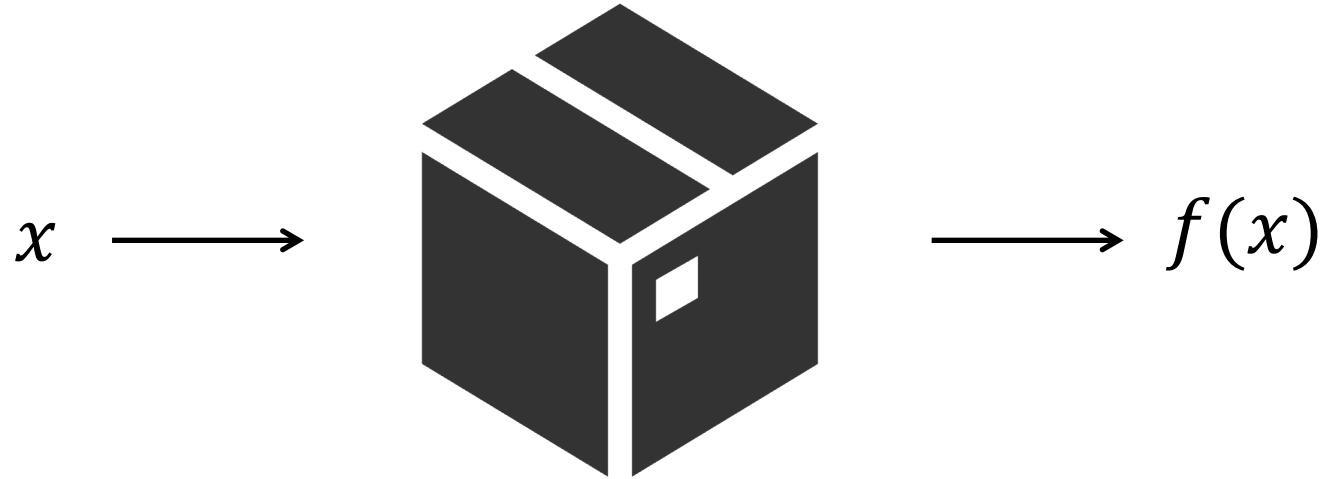
$x \longrightarrow$ [black box] $\longrightarrow f(x)$

$f \sim$ Stochastic Process

Optimizing Black-box Functions

$$x \longrightarrow \boxed{} \longrightarrow f(x)$$

Goal: $\max_{x \in \mathcal{X}} f(x)$

$f \sim$ Stochastic Process

Optimizing Black-box Functions

$$x \longrightarrow \text{[box]} \longrightarrow f(x)$$

Goal: $\max_{x \in \mathcal{X}} f(x)$

$f \sim$ Stochastic Process

Grid Search? Random search?

Optimizing Black-box Functions

$x \longrightarrow$ [black box] $\longrightarrow f(x)$

expensive-to-evaluate

Goal: $\max_{x \in \mathcal{X}} f(x)$

$f \sim$ Stochastic Process

Grid Search? Random search?

Optimizing Black-box Functions

$x \longrightarrow$ $\longrightarrow f(x)$

limited #evaluations

Goal: $\textbf{sup}\ \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$

policy

$f \sim$ Stochastic Process

Grid Search? Random search?

Optimizing Black-box Functions

$$x \longrightarrow \text{[box]} \longrightarrow f(x)$$

limited #evaluations

Goal: $\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$
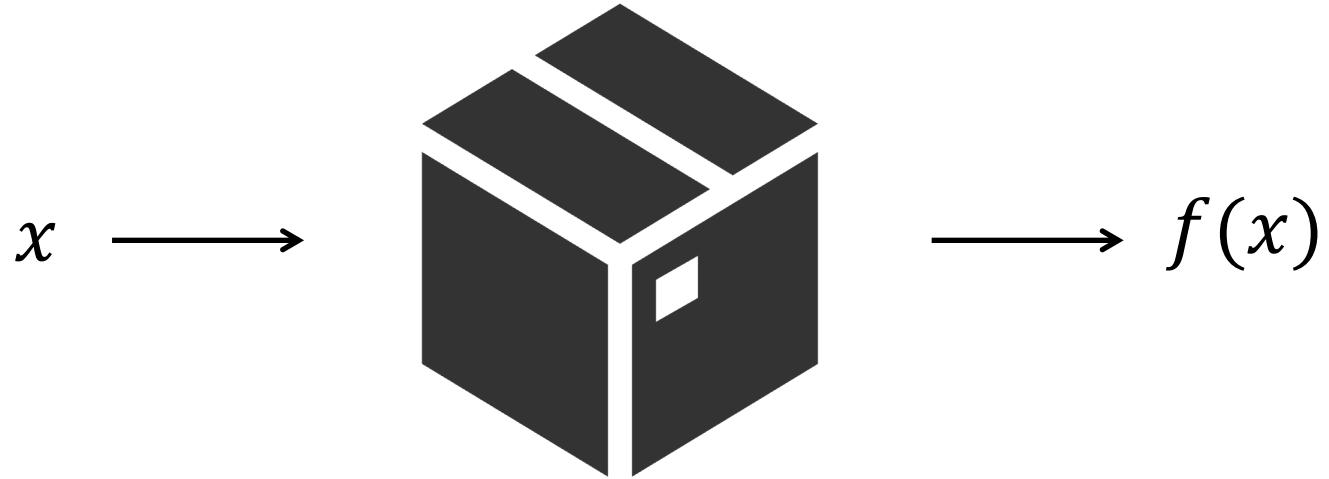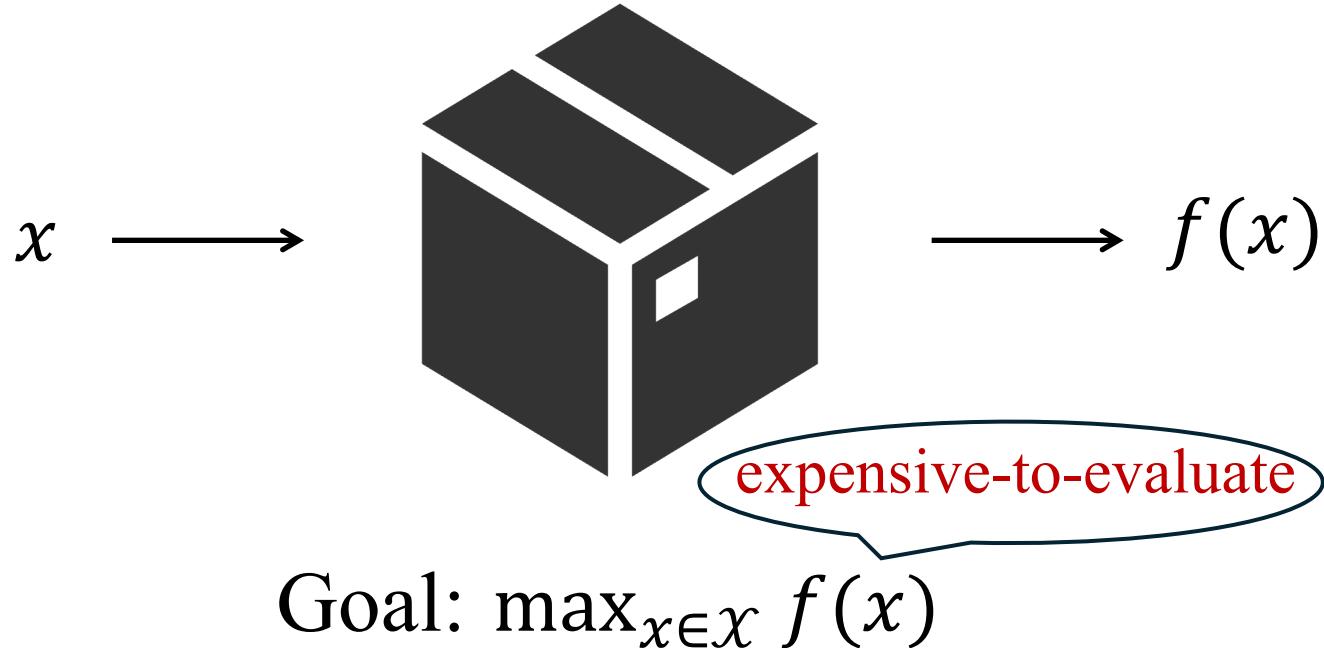
$f \sim$ Stochastic Process
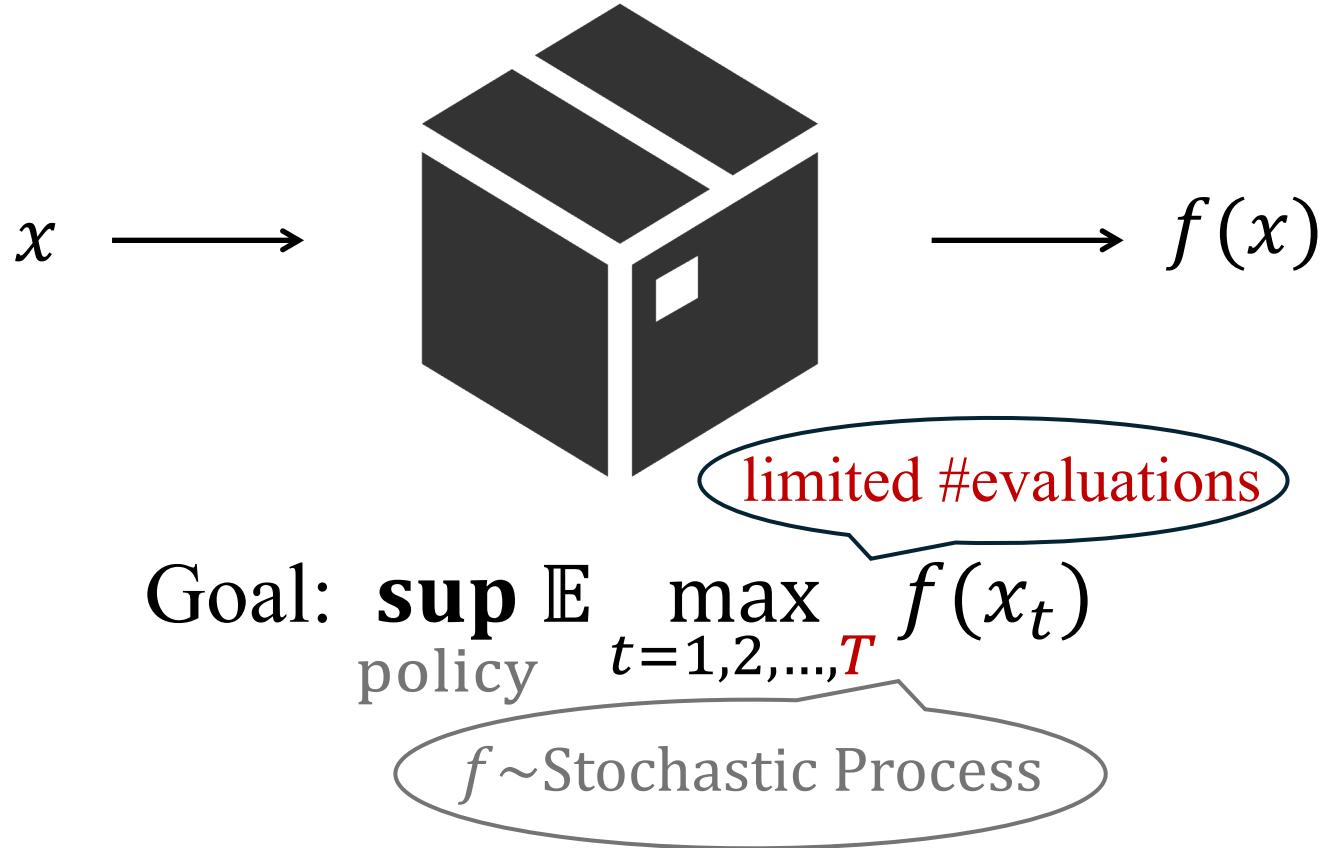
More efficient: Bayesian optimization!

Bayesian Optimization

$$x \longrightarrow \qquad \longrightarrow f(x)$$

limited #evaluations

Goal: $\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$
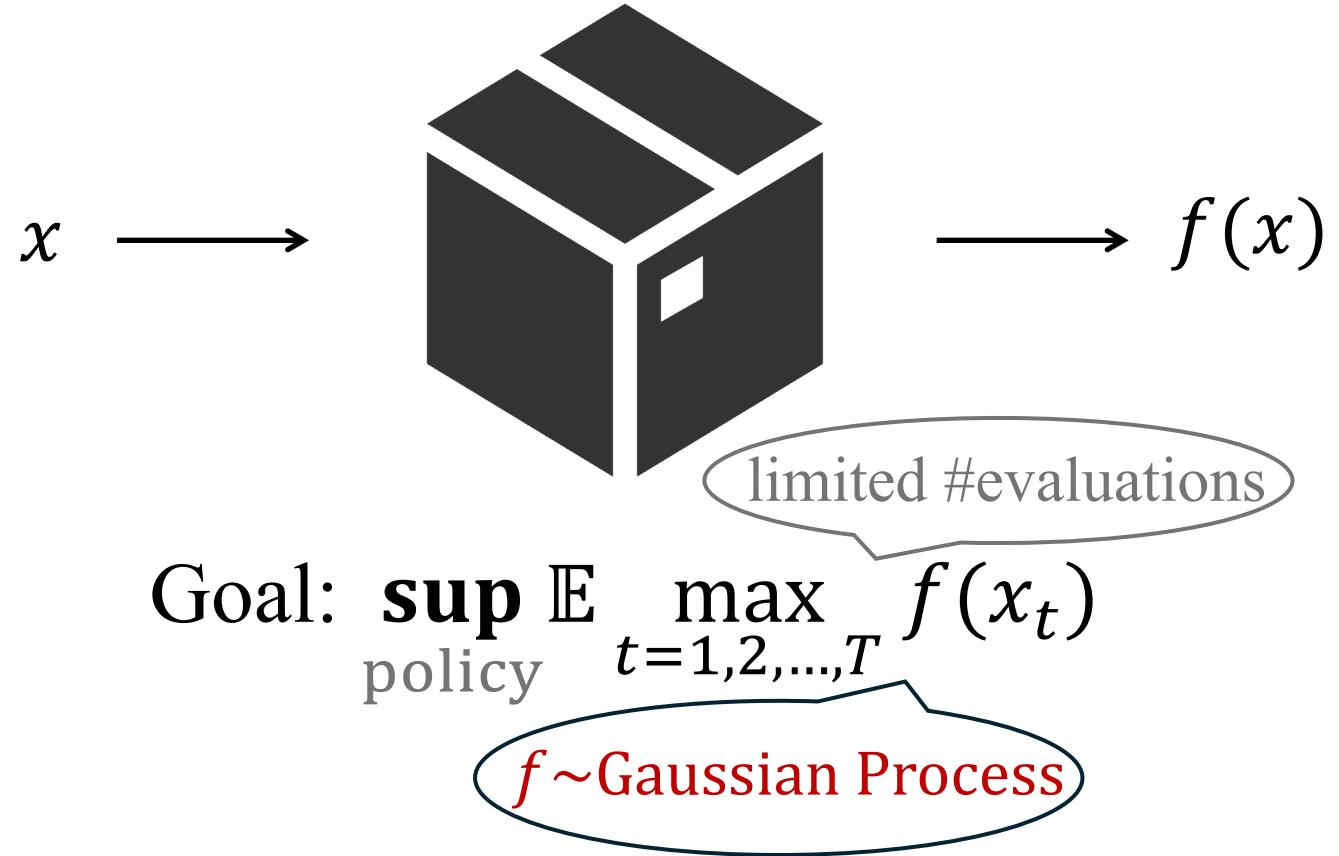
$f \sim$ Gaussian Process

Key idea: maintain posterior belief about f

Bayesian Optimization
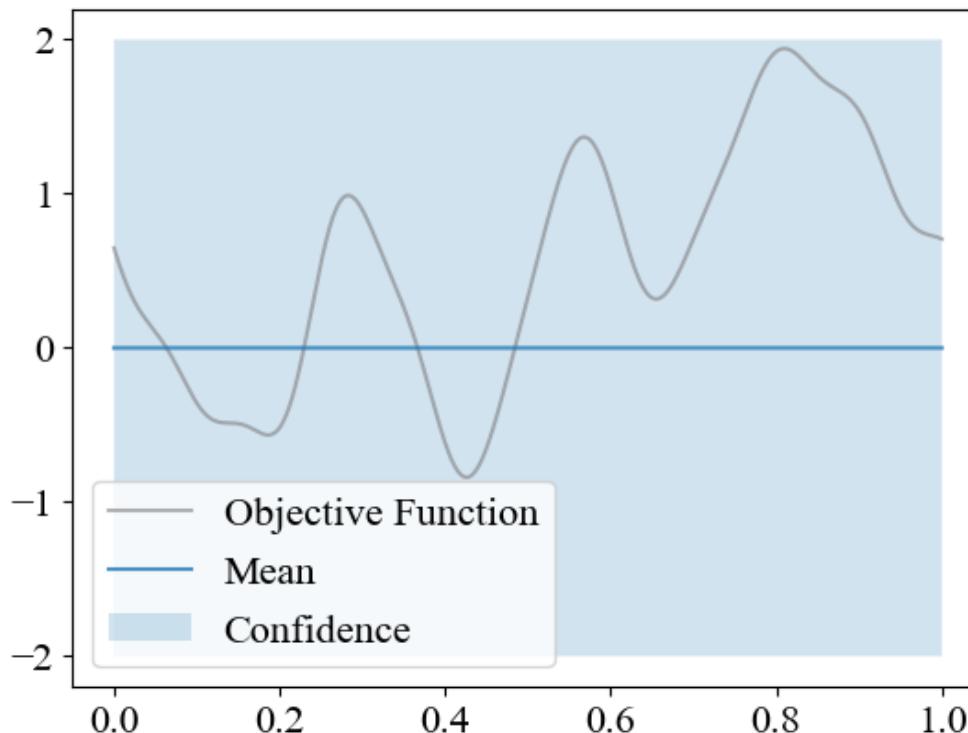
Goal: $\mathbf{sup}\limits_{\text{policy}} \ \mathbb{E} \ \max\limits_{t=1,2,\ldots,T} \ f(x_t)$

$f \sim$ Gaussian Process

Bayesian Optimization

Goal: $\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\dots,T} f(x_t)$
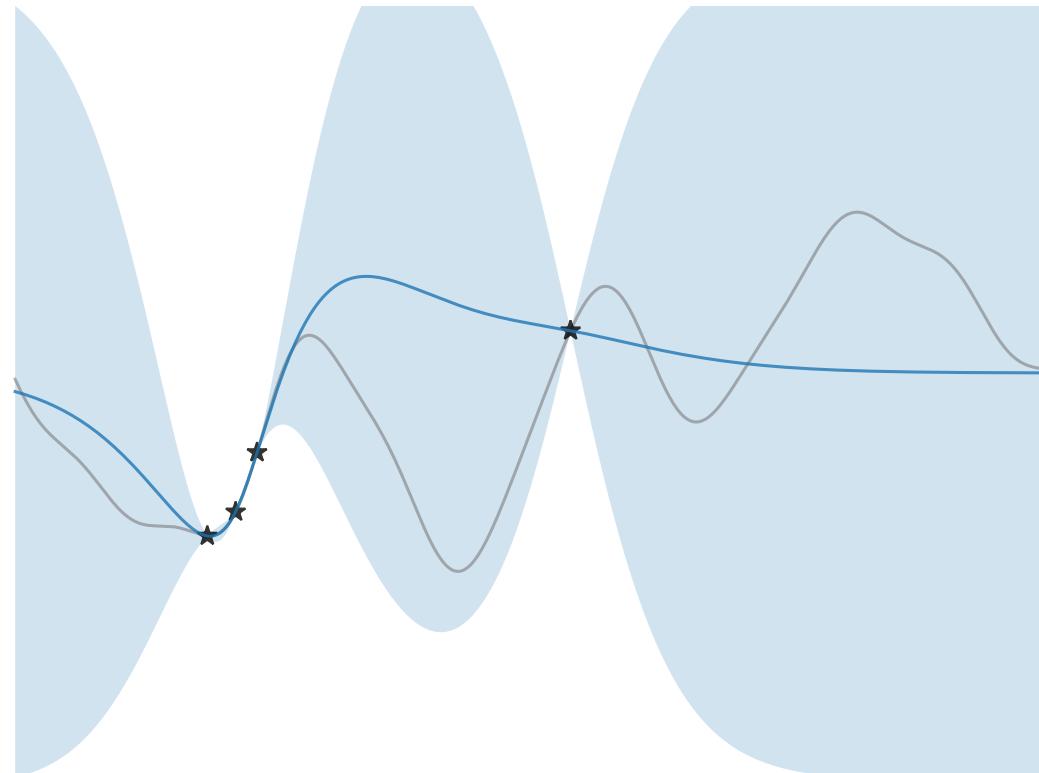
$f \sim \text{Gaussian Process}$

# Bayesian Optimization
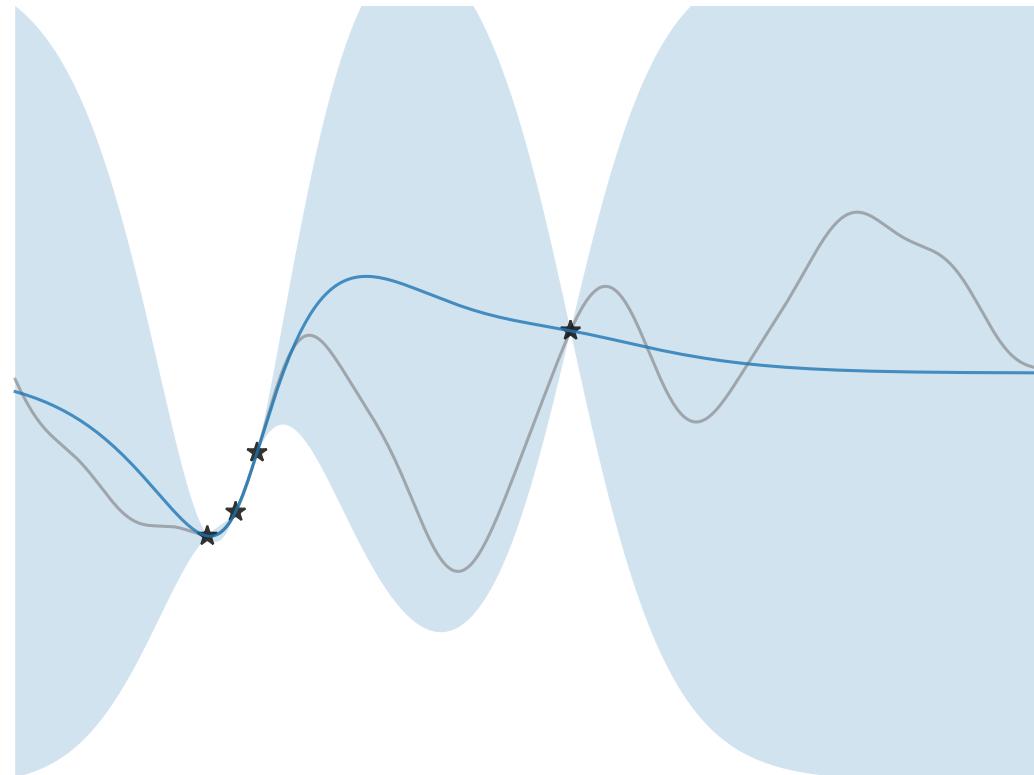
What to evaluate next?

Bayesian Optimization
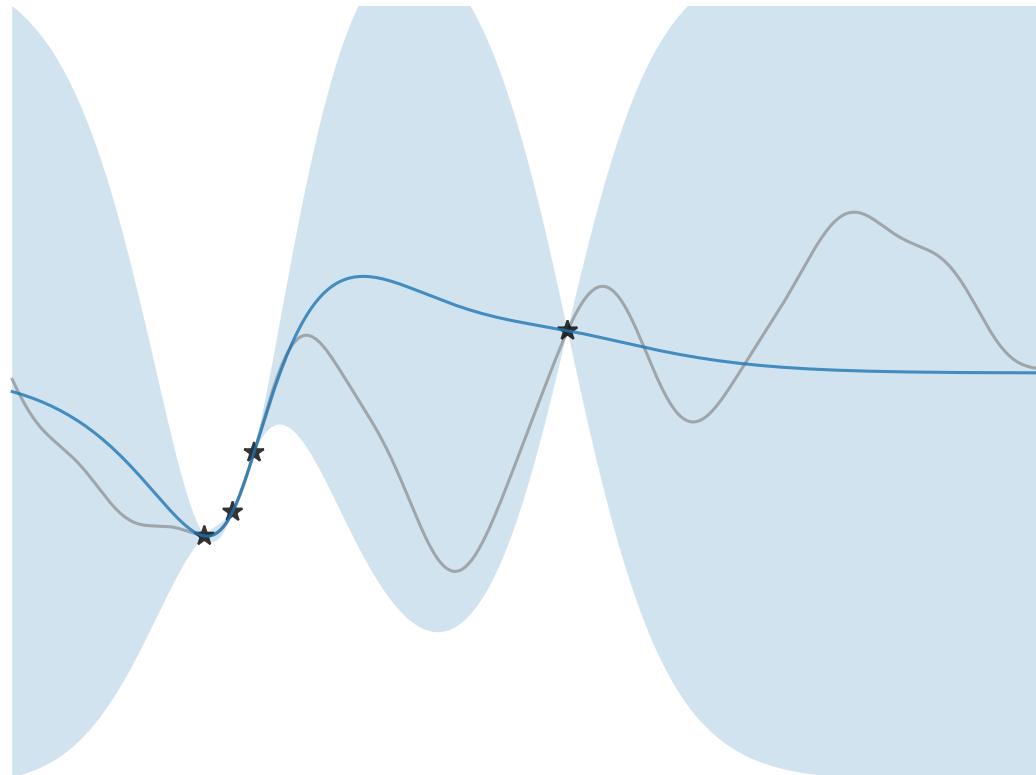

What to evaluate next?

Optimal policy?

Properties of Bayesian Optimization
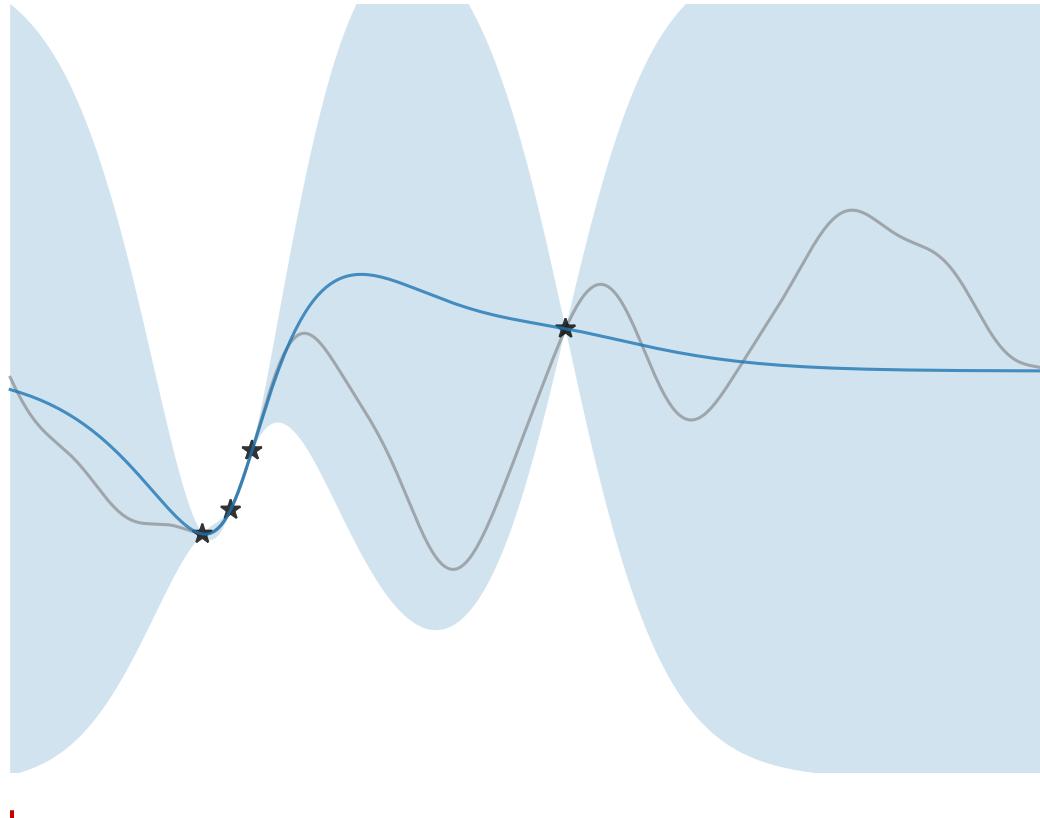


Continuous search domain

Properties of Bayesian Optimization

Correlated values

Continuous search domain
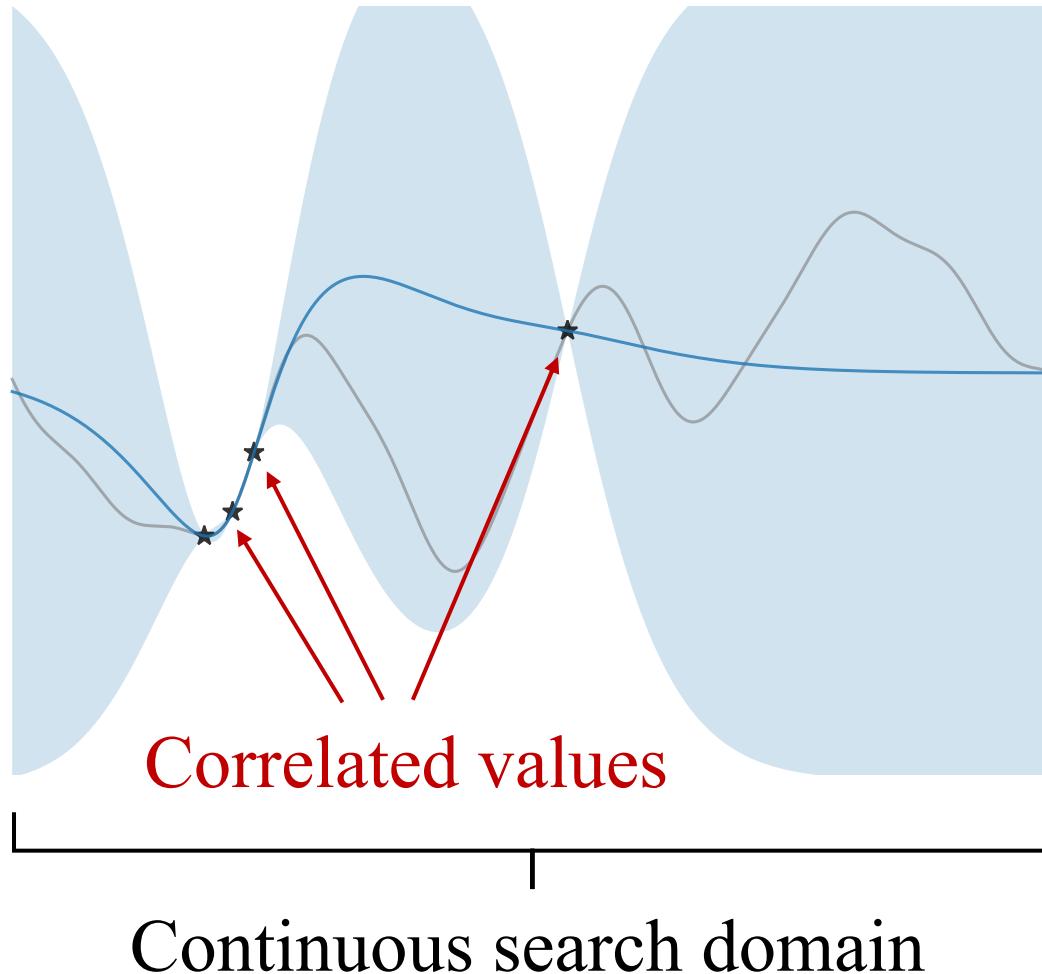
# Properties of Bayesian Optimization
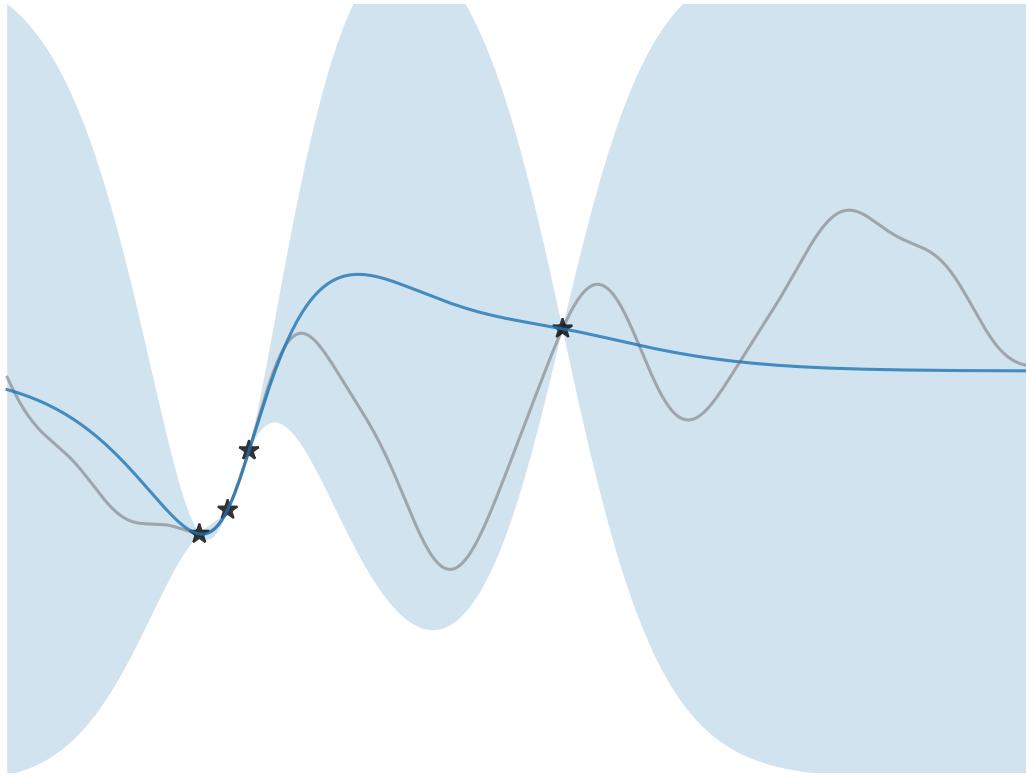
Correlation & continuity $\Longrightarrow$ Intractable MDP

Properties of Bayesian Optimization
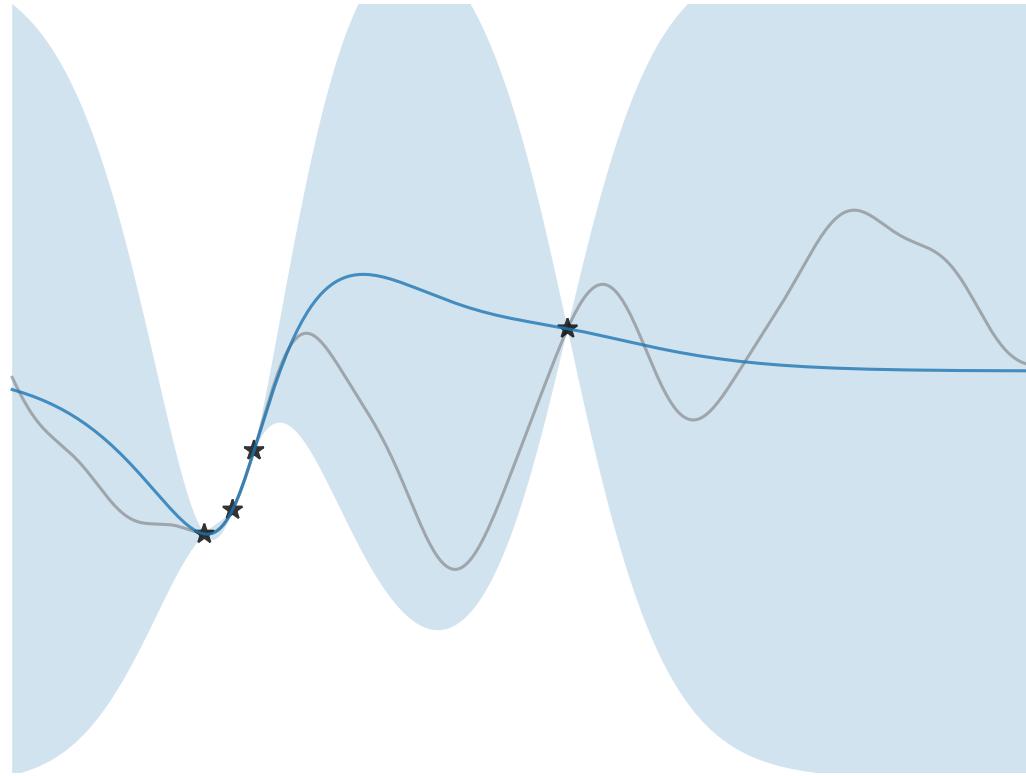
Intractable MDP $\Longrightarrow$ Optimal policy unknown

Popular Policy: Expected Improvement

current best observed data
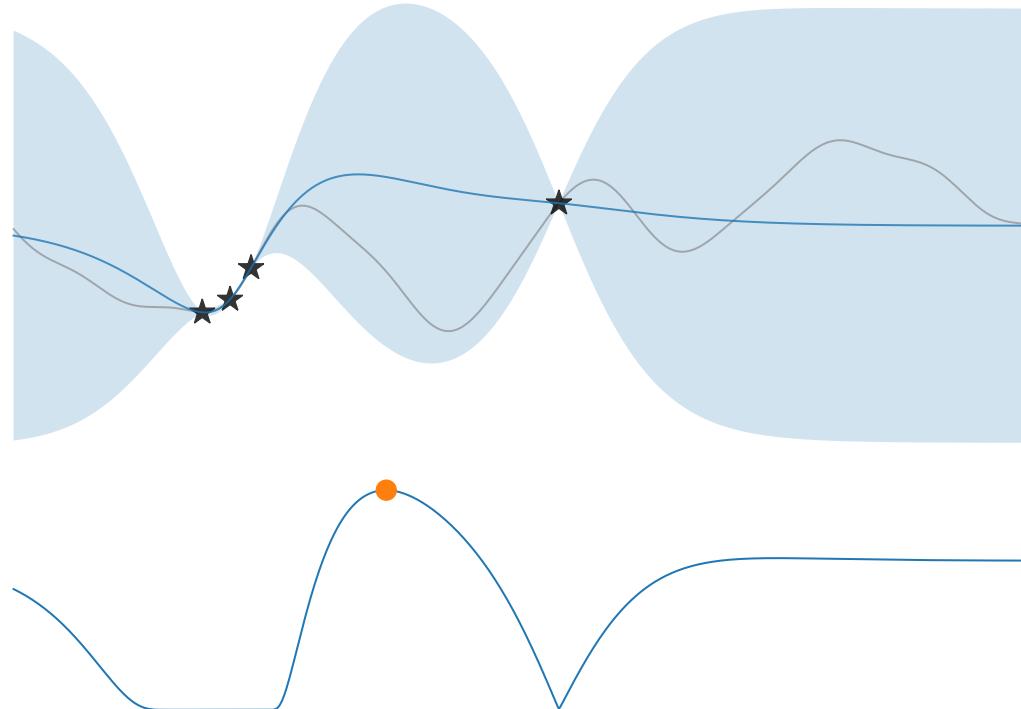
$$\mathrm{EI}(x) = \mathbb{E}[\max(f(x) - y_{\mathrm{best}}, 0) \mid D]$$

"improvement"

$$\max_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$$

posterior distribution

One-step approximation to MDP

Popular Policy: Expected Improvement
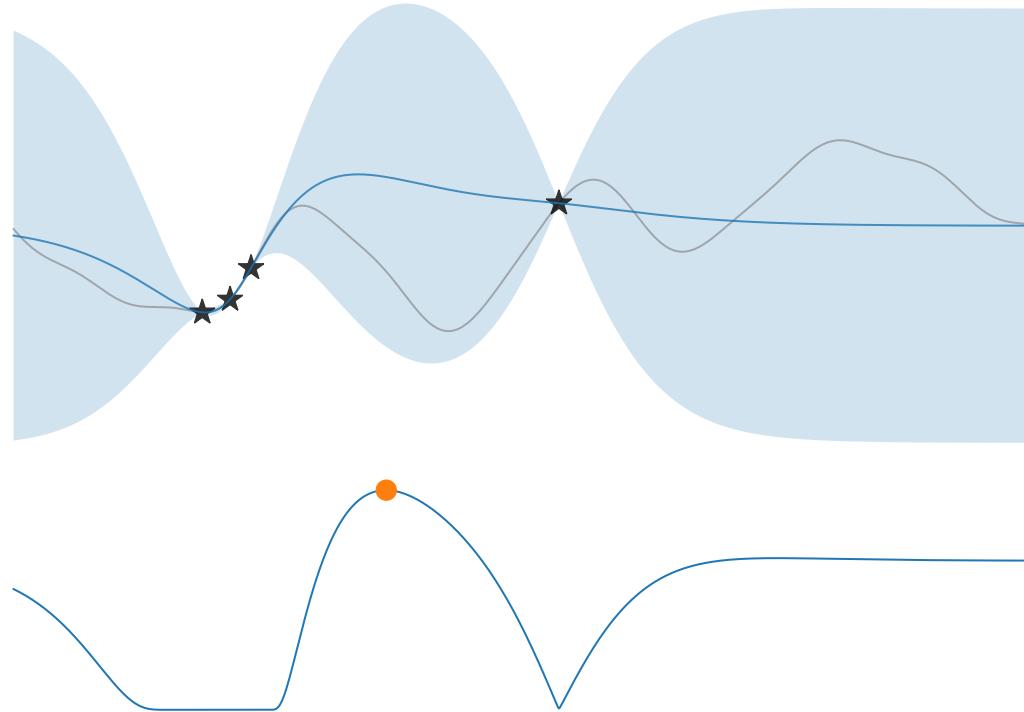
Other improvement-based policies:
- Probability of Improvement
- Knowledge Gradient
- Multi-step Lookahead EI
- ⋮

multi-step approximation to MDP

Approaches to Bayesian Optimization

- Improvement-based:
  - Expected Improvement
  - Probability of Improvement
  - Knowledge Gradient
  - Multi-step Lookahead EI

Approaches to Bayesian Optimization

- Improvement-based:
  - Expected Improvement
  - Probability of Improvement
  - Knowledge Gradient
  - Multi-step Lookahead EI

- Entropy-based:
  - Max-value Entropy Search
  - Predictive Entropy Search

Approaches to Bayesian Optimization

- Improvement-based:
    - Expected Improvement
    - Probability of Improvement
    - Knowledge Gradient
    - Multi-step Lookahead EI

- Entropy-based:
    - Max-value Entropy Search
    - Predictive Entropy Search

- Upper Confidence Bound

Approaches to Bayesian Optimization

- Improvement-based:
  - Expected Improvement
  - Probability of Improvement
  - Knowledge Gradient
  - Multi-step Lookahead EI

- Entropy-based:
  - Max-value Entropy Search
  - Predictive Entropy Search

- Upper Confidence Bound

- Thompson Sampling

Approaches to Bayesian Optimization

- Improvement-based
- Entropy-based
- Upper Confidence Bound
- Thompson Sampling

Approaches to Bayesian Optimization

- Improvement-based

- Entropy-based

- Upper Confidence Bound

- Thompson Sampling

- Our work: Gittins Index

Approaches to Bayesian Optimization

- Improvement-based

- Entropy-based

- Upper Confidence Bound

- Thompson Sampling

- Our work: Gittins Index

Why another approach?

Approaches to Bayesian Optimization

- Improvement-based
- Entropy-based
- Upper Confidence Bound
- Thompson Sampling
- Our work: Gittins Index

Why another approach?

Less-explored Practical Factors!

Outline: Two Less-explored Practical Factors

Part I: Varying Evaluation Costs

      - Cost-aware Bayesian Optimization

Outline: Two Less-explored Practical Factors

Part I: Varying Evaluation Costs
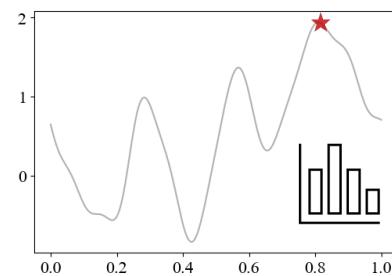
- Cost-aware Bayesian Optimization

Part II: Observable Partial Feedback

- Freeze-thaw Bayesian Optimization
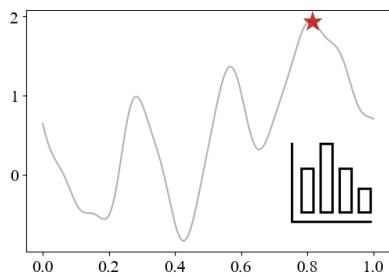
Outline: Two Less-explored Practical Factors

Part I: Varying Evaluation Costs

    - Cost-aware Bayesian Optimization

Part II: Observable Partial Feedback

    - Freeze-thaw Bayesian Optimization

New design principle: Gittins index
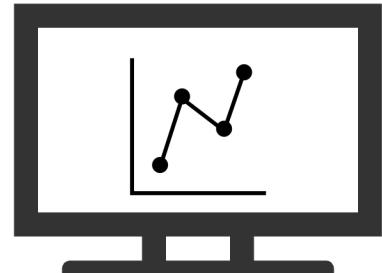
Part I: Varying Evaluation Costs

Hyperparameter tuning:
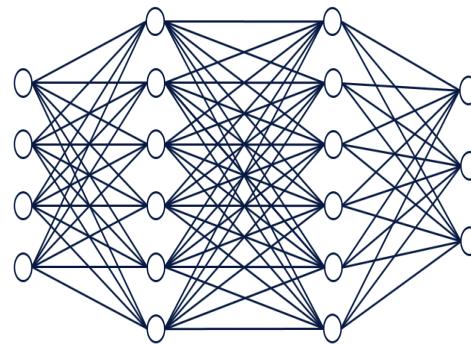
Training parameters $\longrightarrow$ $\longrightarrow$ Accuracy

Control optimization:

Control variables $\longrightarrow$ $\longrightarrow$ Reward
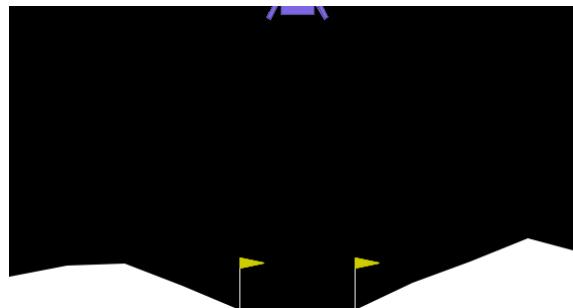
Plasma physics:

Reactor parameters $\longrightarrow$ $\longrightarrow$ Stability
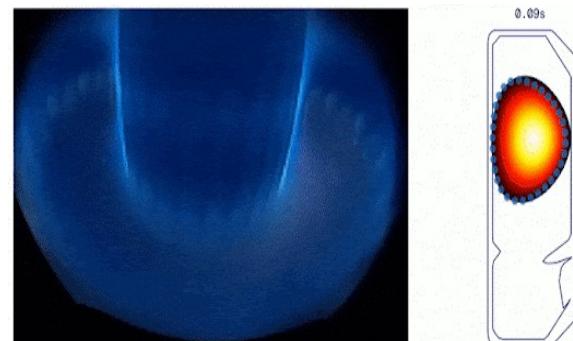
Practical Factor: Varying Evaluation Costs

$$x \longrightarrow \boxed{} \longrightarrow f(x)$$

$c(x)$ — varying

Cost-aware Bayesian Optimization

[Lee, Perrone, Archambeau, Seeger'21]

$$x \longrightarrow \text{[box]} \longrightarrow f(x)$$
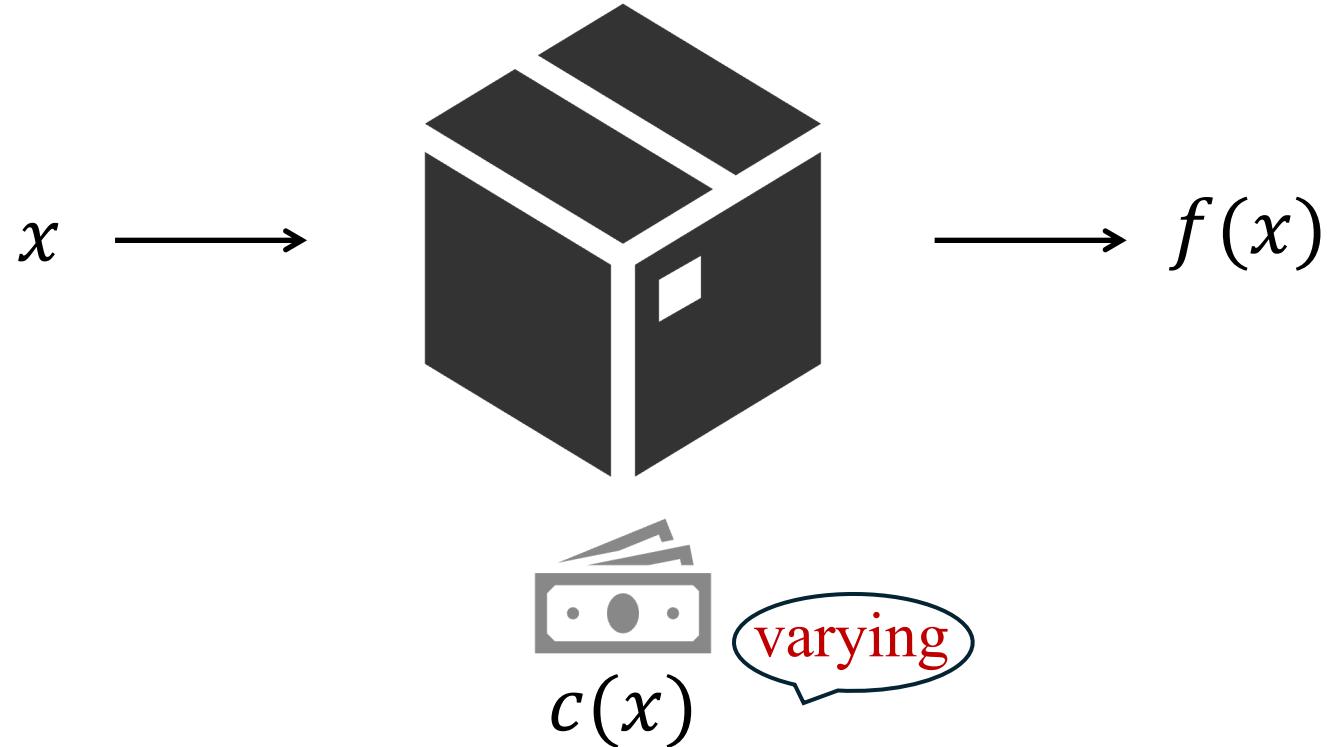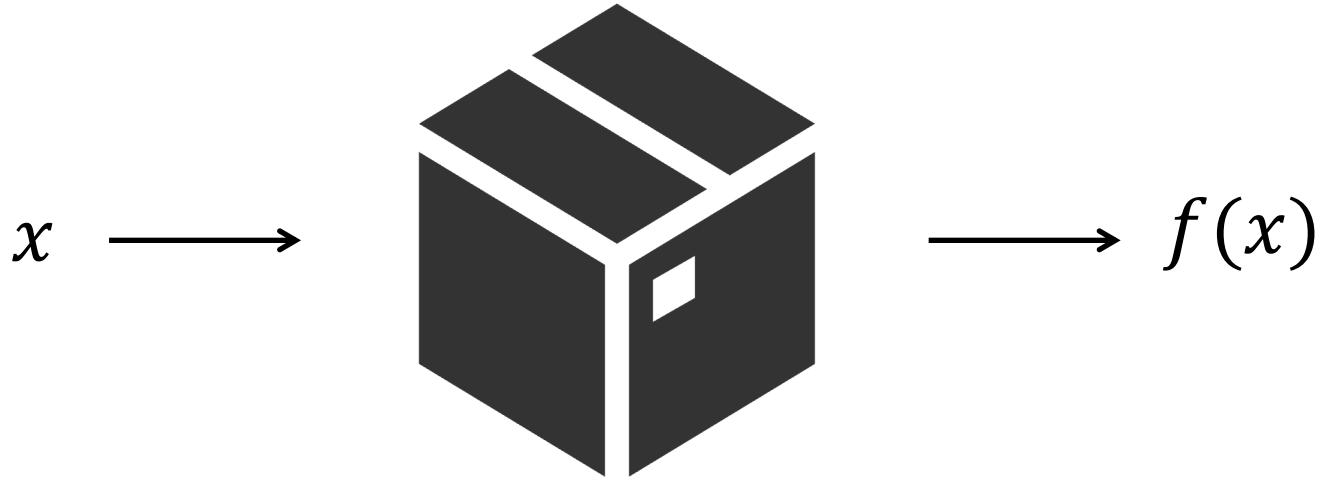
Goal: $\displaystyle \sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\dots,T} f(x_t)$

s.t. $\displaystyle \mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$

Cost-aware Bayesian Optimization

[Lee, Perrone, Archambeau, Seeger'21]

$$x \longrightarrow \boxed{} \longrightarrow f(x)$$
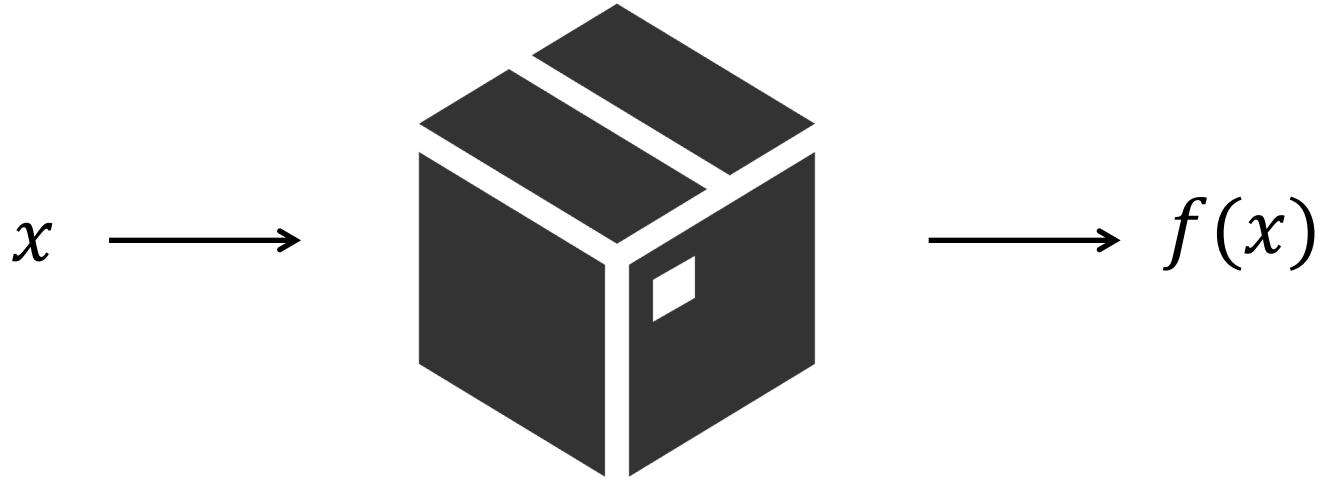
Goal: $\mathbf{sup}_{\text{policy}} \; \mathbb{E} \max_{t=1,2,\dots,T} f(x_t)$

s.t. $\mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$ Expected budget constraint

"Multi-step Budgeted Bayesian Optimization with Unknown Evaluation Costs"
[Astudillo, Jiang, Balandat, Bakshy, Frazier'21]

Cost-aware Bayesian Optimization

Uniform costs   Varying costs

One-step   Expected improvement

$$\max_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$$

Cost-aware Bayesian Optimization

Uniform costs	Varying costs
One-step Expected improvement	Expected improvement per cost
$$\max_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$$	[Snoek, Larochelle, Adams'21]

Cost-aware Bayesian Optimization

	Uniform costs	Varying costs
One-step	Expected improvement	Expected improvement per cost
	$\max_x \text{EI}_{f|D}(x; y_{\text{best}})$	$\max_x \text{EI}_{f|D}(x; y_{\text{best}})/c(x)$

Cost-aware Bayesian Optimization

	Uniform costs	Varying costs

One-step

Expected improvement

$$\max_x \mathrm{EI}_{f|D}(x; y_{\text{best}})$$

Expected improvement per cost

$$\max_x \mathrm{EI}_{f|D}(x; y_{\text{best}})/c(x)$$

Why divide?

Cost-aware Bayesian Optimization

| | Uniform costs | Varying costs |

Uniform costs   Varying costs

One-step   Expected improvement   Expected improvement per cost

$$\max_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$$

$$\max_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})/c(x)$$

arbitrarily bad

[Astudillo, Jiang, Balandat, Bakshy, Frazier'21]

Cost-aware Bayesian Optimization

	Uniform costs	Varying costs
One-step	Expected improvement	Expected improvement per cost
Multi-step	Multi-step Lookahead EI	Budgeted Multi-step Lookahead EI

slow

[Astudillo, Jiang, Balandat, Bakshy, Frazier'21]

Cost-aware Bayesian Optimization

	Uniform costs	Varying costs
One-step	Expected improvement	Expected improvement per cost
Multi-step	Multi-step Lookahead EI	Budgeted Multi-step Lookahead EI
	Upper Confidence Bound	?
	Thompson Sampling	?

Cost-aware Bayesian Optimization

	Uniform costs	Varying costs
One-step	Expected improvement	Expected improvement per cost
Multi-step	Multi-step Lookahead EI	Budgeted Multi-step Lookahead EI
	Upper Confidence Bound	?
	Thompson Sampling	?
	⋮	⋮

Our view: lack of a guidance to incorporate costs

Cost-aware Bayesian Optimization

	Uniform costs	Varying costs
One-step	Expected improvement	Expected improvement per cost
Multi-step	Multi-step Lookahead EI	Budgeted Multi-step Lookahead EI
	Upper Confidence Bound	?
	Thompson Sampling	?
	⋮	⋮

New design principle: Gittins Index

Cost-aware Bayesian Optimization

	Uniform costs	Varying costs
One-step	Expected improvement	Expected improvement per cost
Multi-step	Multi-step Lookahead EI	Budgeted Multi-step Lookahead EI
	Upper Confidence Bound	?
	Thompson Sampling	?
	⋮	⋮

New design principle: Gittins Index

inherently cost-aware

Cost-aware Bayesian Optimization

$f(x)$

$c(x)$
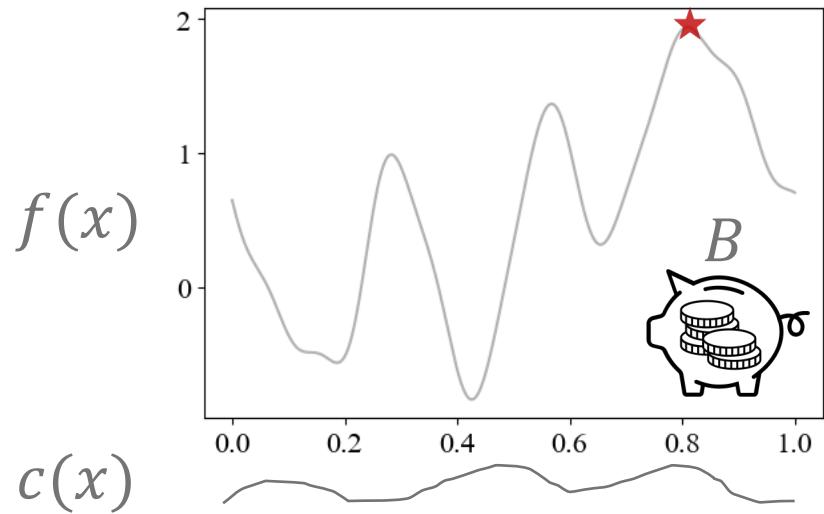
Continuous

Correlated

Cost-aware Bayesian Optimization

Pandora's Box

$f(x)$

$c(x)$

$\implies$

$f(x)$

$c(x)$
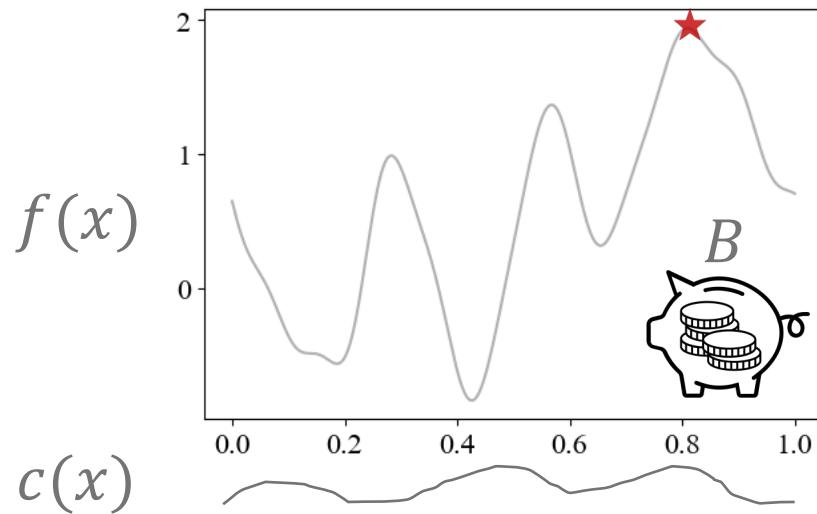
Continuous $\implies$ Discrete

Correlated $\implies$ Independent

Pandora's Box

$t = 0$
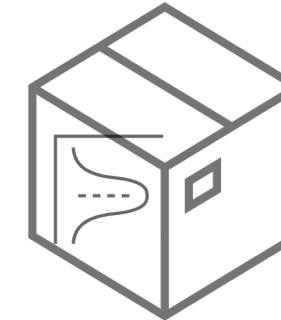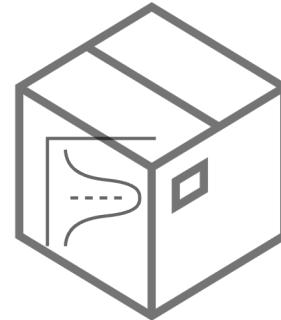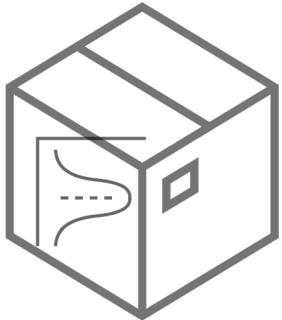
$$\sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\ldots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Pandora's Box

$t = 1$

$f(x_1)$
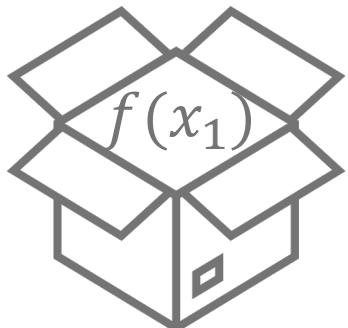
$c(x_1)$
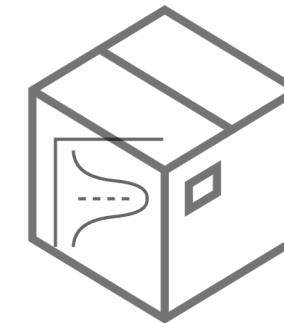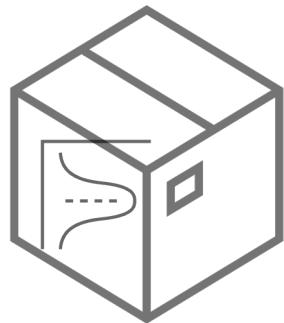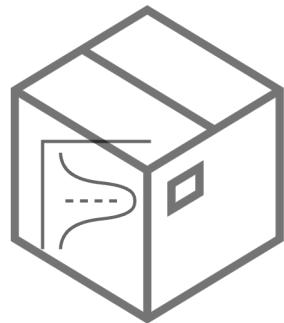
$$\sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\dots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Pandora's Box

$t = 2$

$f(x_1)$

$c(x_1)$

$f(x_2)$

$c(x_2)$

$$\sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\dots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Pandora's Box

$t = 3$
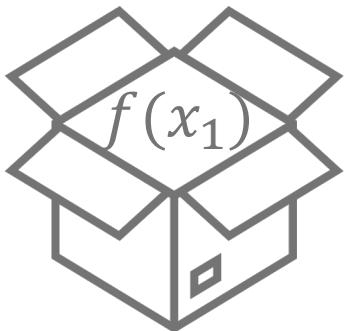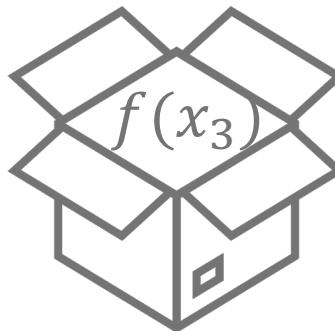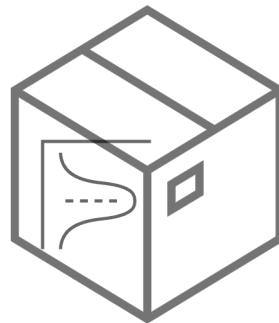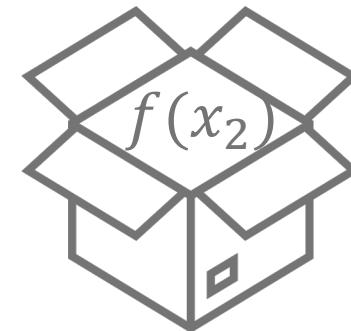
$f(x_1)$

$f(x_3)$

$f(x_2)$

$c(x_1)$

$c(x_3)$

$c(x_2)$

$$\sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\dots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Pandora's Box

$t = T,$ stop
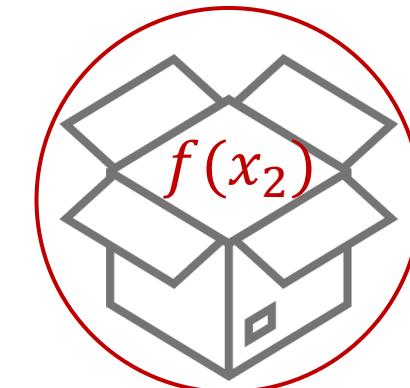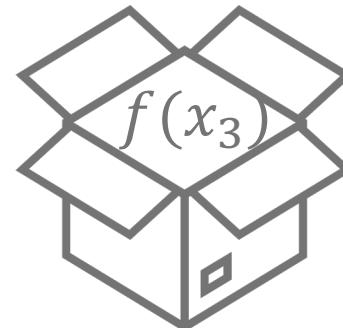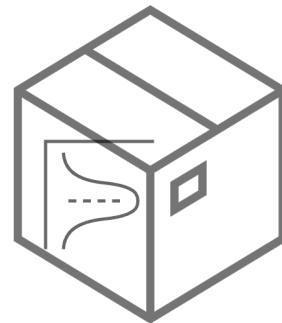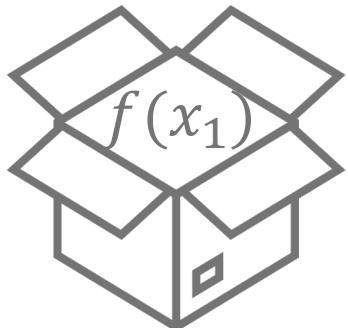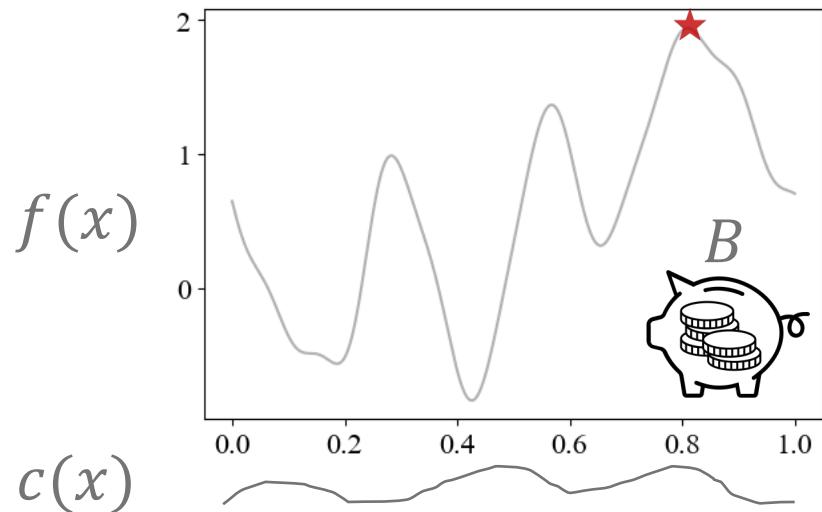
$f(x_1)$

$f(x_3)$

$f(x_2)$

$c(x_1)$

$c(x_3)$

$c(x_2)$

$$\sup_{\text{policy}} \mathbb{E}\left(\max_{t=1,2,\ldots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Cost-aware Bayesian Optimization

$f(x)$

$c(x)$

Pandora's Box

$f(x)$

$c(x)$

Continuous	Discrete
Correlated	Independent
Expected-budget-constrained	**Cost-per-sample**

$$\sup_{\text{policy}} \ \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$
$$\text{s.t. } \mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$$
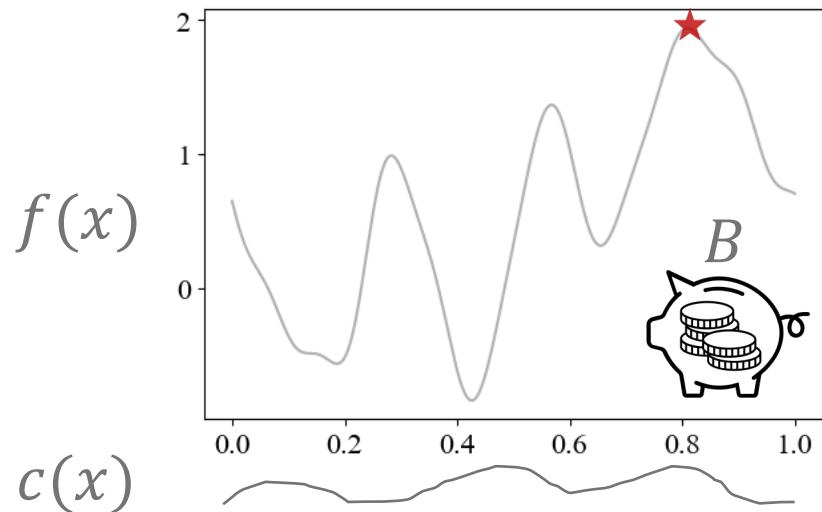
$$\sup_{\text{policy}} \ \mathbb{E} \left(\max_{t=1,2,\ldots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Cost-aware Bayesian Optimization

Pandora's Box

[Weitzman'79]

$f(x)$

B

$c(x)$

$f(x)$

$c(x)$

Continuous

Correlated

Expected-budget-constrained

Discrete

Independent

Cost-per-sample

Optimal policy: Gittins index

Optimal Policy: Gittins Index

Step 1: Assign each box a Gittins index (higher is better)

$f(x)$

$c(x)$

g
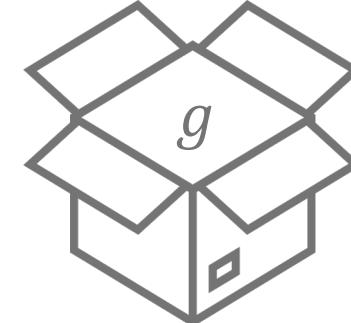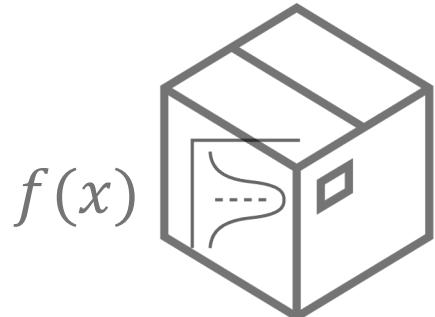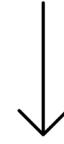
$\text{GI}_f(x; c(x))$

g

Optimal Policy: Gittins Index

Step 2: Open the box with highest index if it is closed
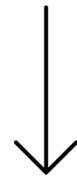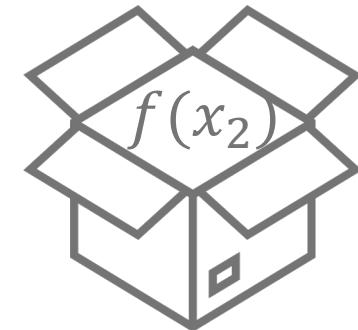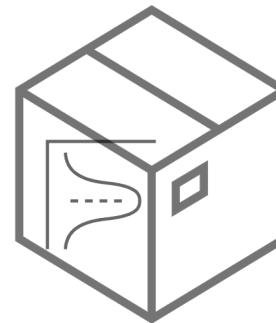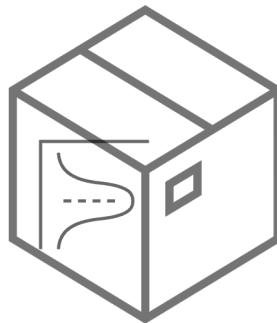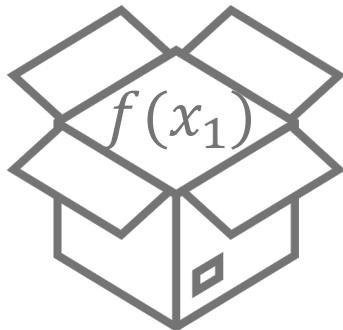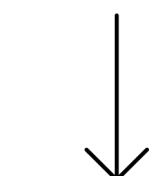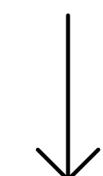
$f(x_1)$ $\qquad$ $\text{GI}_f(x; c(x))$ $\qquad$ $\text{GI}_f(x'; c(x'))$ $\qquad$ $f(x_2)$

Optimal Policy: Gittins Index

Step 2': Select the box with highest index if it is opened and stop
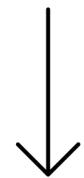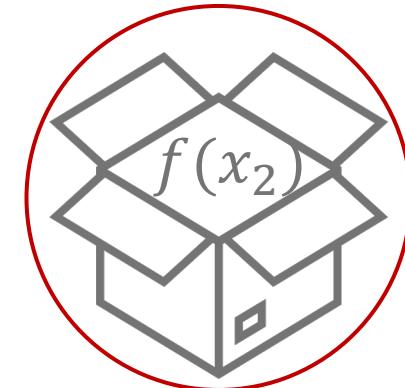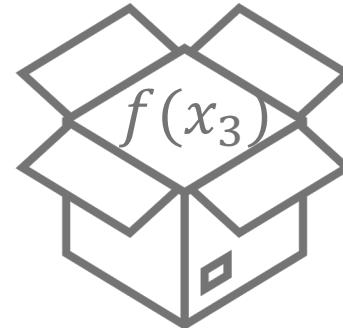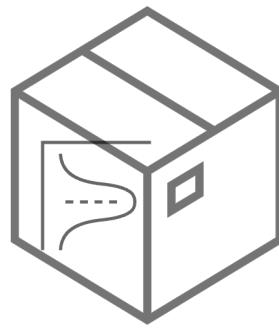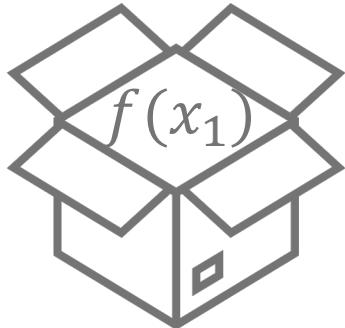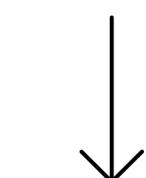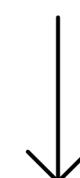

$f(x_1)$ $\qquad$ $\mathrm{GI}_f(x; c(x))$ $\qquad$ $f(x_3)$ $\qquad$ $\textcolor{red}{f(x_2)}$

Cost-aware Bayesian Optimization
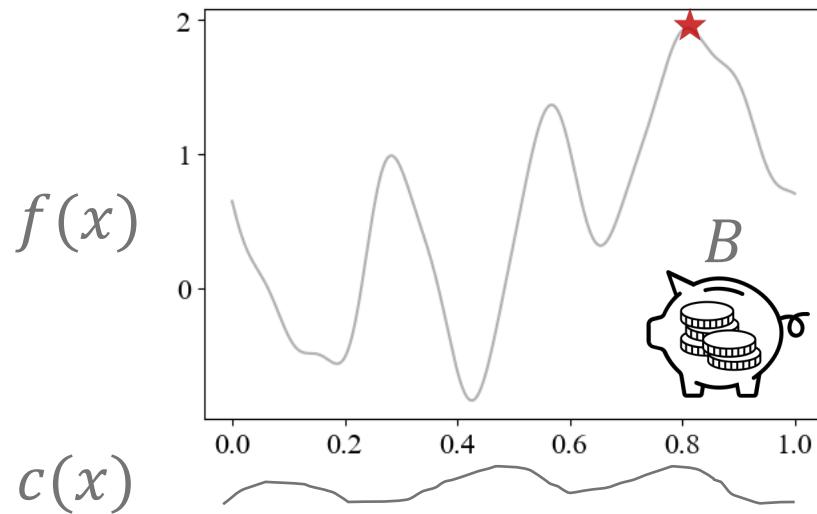
Pandora's Box

$f(x)$

$c(x)$

$f(x)$

$c(x)$

Continuous

Discrete

Correlated

Independent

Expected-budget-constrained

Cost-per-sample

How to translate?

$\Longleftarrow$ Optimal policy: $\text{GI}_f(x; c(x))$

Pandora's Box

[Weitzman'79]

$f(x)$

$c(x)$
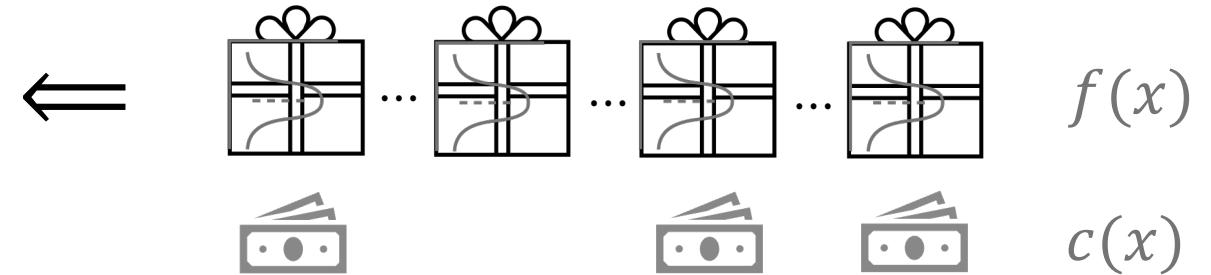
How to convert?

Expected-budget-constrained $\Longleftarrow$ Cost-per-sample

Budgeted Pandora's Box

[Aminian, Manshadi, Niazadeh'24]

Pandora's Box

[Weitzman'79]

$f(x)$ $\Longleftarrow$ $f(x)$

B $c(x)$

Expected-budget-constrained   Cost-per-sample

Budgeted Pandora's Box

[Aminian, Manshadi, Niazadeh'24]

Pandora's Box

[Weitzman'79]

$$\mathbf{sup}_{\text{policy}} \; \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$
$$\mathbf{s.t.} \; \mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$$

Expected-budget-constrained
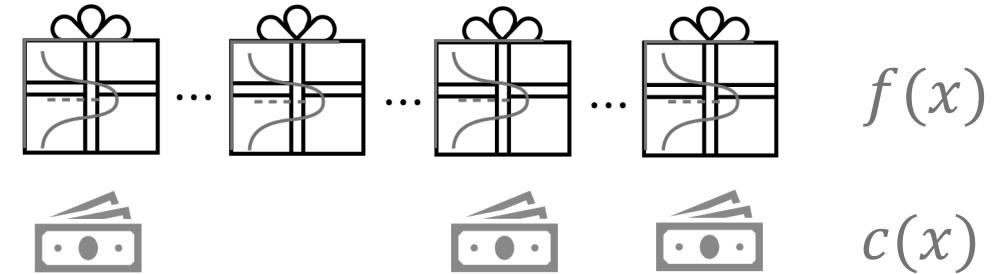
$$\mathbf{sup}_{\text{policy}} \; \mathbb{E} \left(\max_{t=1,2,\ldots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Cost-per-sample

Budgeted Pandora's Box
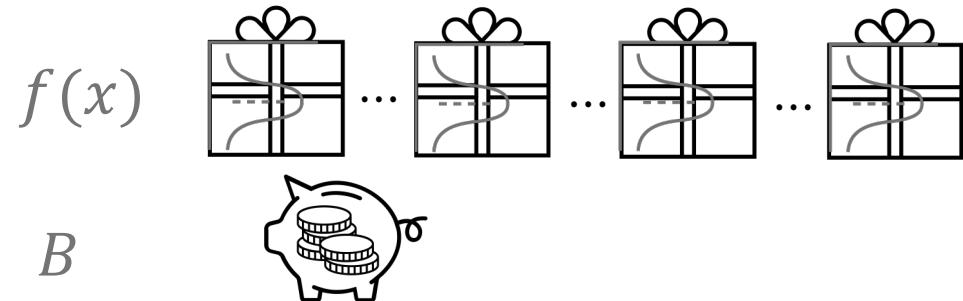
[Aminian, Manshadi, Niazadeh'24]

Pandora's Box

[Weitzman'79]

$f(x)$ B $\Longleftarrow$ $f(x)$ $c(x)$
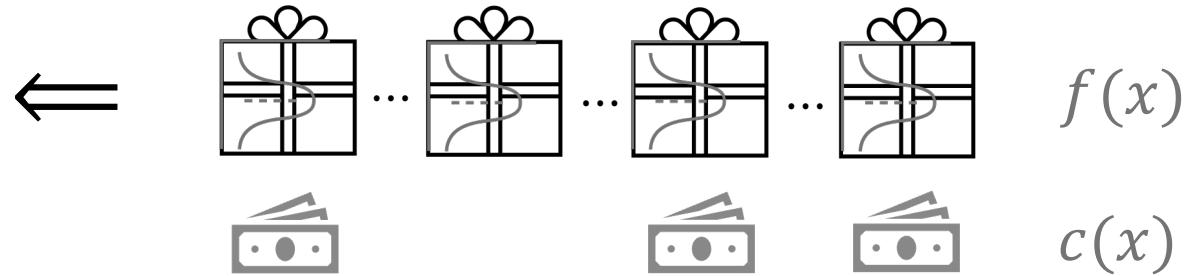
scaling factor

$$\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$
$$\text{s.t. } \mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$$

$$\sup_{\text{policy}} \mathbb{E} \left(\max_{t=1,2,\ldots,T} f(x_t) - \lambda_B \sum_{t=1}^{T} c(x_t) \right)$$

Lagrangian relaxation

Expected-budget-constrained $\Longleftrightarrow$ Cost-per-sample

Budgeted Pandora's Box
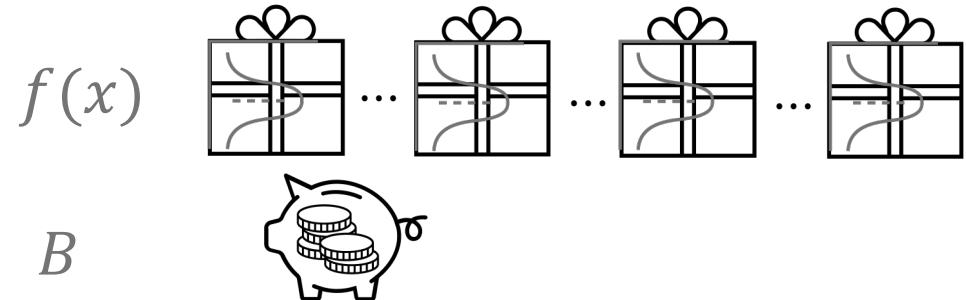
[Aminian, Manshadi, Niazadeh'24]

Pandora's Box

[Weitzman'79]
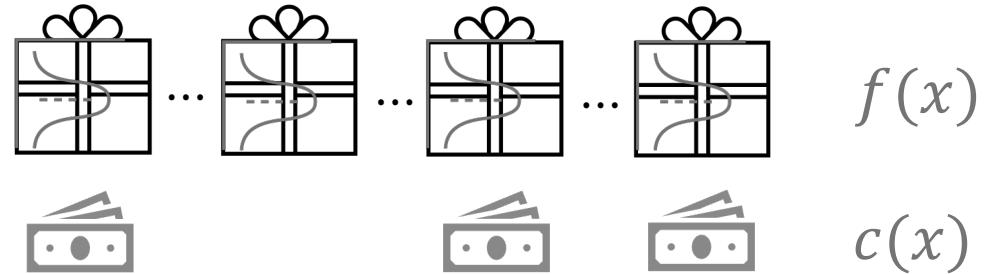
$$\sup_{\text{policy}} \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$
$$\textbf{s.t. } \mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$$

$$\sup_{\text{policy}} \mathbb{E} \left(\max_{t=1,2,\ldots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Expected-budget-constrained $\iff$ Cost-per-sample

How to translate?

$\impliedby$ Optimal policy: $\text{GI}_f(x; c)$

Budgeted Pandora's Box
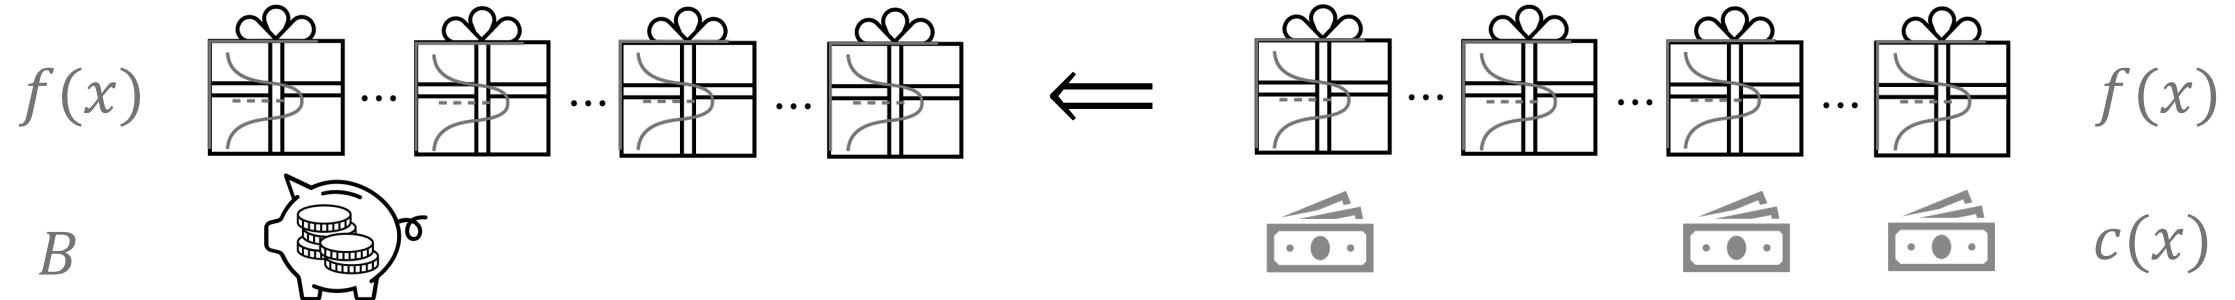
[Aminian, Manshadi, Niazadeh'24]

Pandora's Box

[Weitzman'79]

$$\mathbf{sup}_{\text{policy}} \; \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$

$$\mathbf{s.t.} \; \mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$$

$$\mathbf{sup}_{\text{policy}} \; \mathbb{E} \left(\max_{t=1,2,\ldots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Expected-budget-constrained $\iff$ Cost-per-sample

scale costs

Optimal policy: $\text{GI}_f(x; \lambda_B c(x)) \impliedby$ Optimal policy: $\text{GI}_f(x; c(x))$

Budgeted Pandora's Box
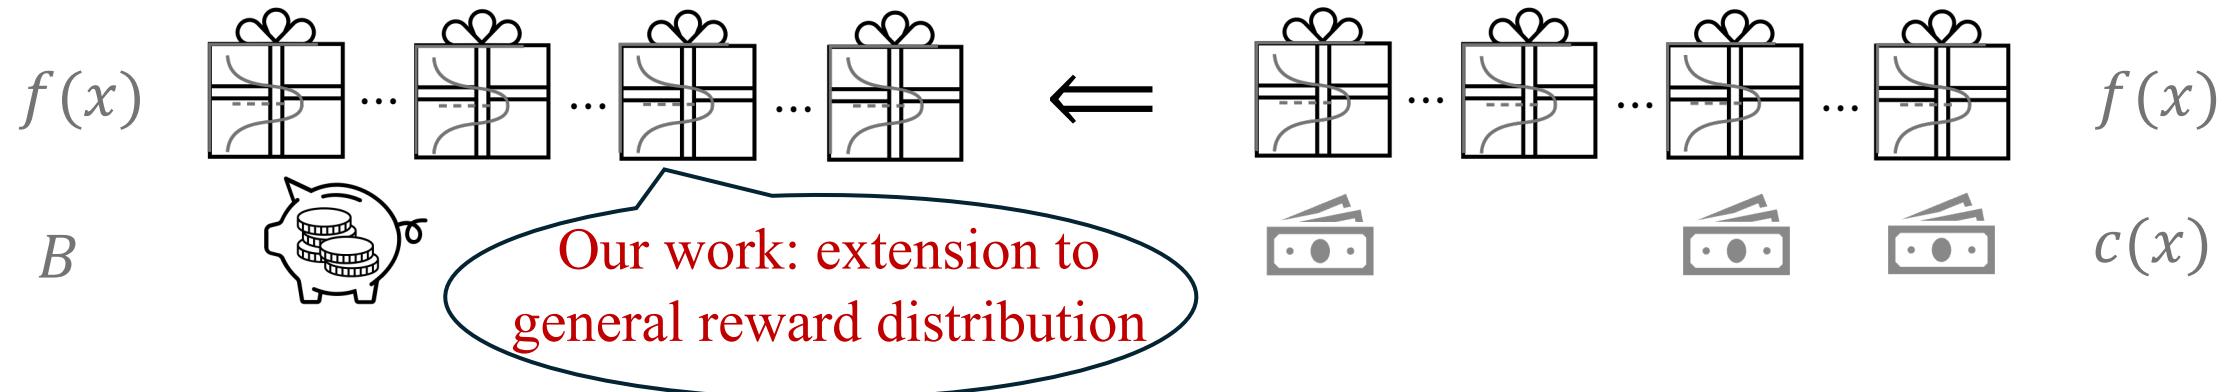
[Aminian, Manshadi, Niazadeh'24]

Pandora's Box

[Weitzman'79]

$f(x)$ ⬅ $f(x)$

B *Our work: extension to general reward distribution* $c(x)$

$$\mathbf{sup}_{\text{policy}} \; \mathbb{E} \max_{t=1,2,\ldots,T} f(x_t)$$

$$\mathbf{s.t.} \; \mathbb{E} \sum_{t=1}^{T} c(x_t) \leq B$$

$$\mathbf{sup}_{\text{policy}} \; \mathbb{E}\left(\max_{t=1,2,\ldots,T} f(x_t) - \sum_{t=1}^{T} c(x_t) \right)$$

Expected-budget-constrained   ⟺   Cost-per-sample

Optimal policy: $\text{GI}_f(x; \lambda_B c(x))$ $\overset{\text{scale costs}}{\Longleftarrow}$ Optimal policy: $\text{GI}_f(x; c(x))$

Cost-aware Bayesian Optimization

Pandora's Box

$f(x)$

$c(x)$

$f(x)$

$c(x)$

Continuous   Discrete

Correlated   Independent

Expected-budget-constrained $\Longleftarrow$ Cost-per-sample   ✔

How to translate?

$\Longleftarrow$ Optimal policy: $\text{GI}_f(x; c(x))$

Cost-aware Bayesian Optimization   Budgeted Pandora's Box

$f(x)$

B

$c(x)$

$f(x)$

B

Continuous

Correlated

Discrete

Independent

scaling factor

How to translate?

$\Longleftarrow$ Optimal policy: $\mathrm{GI}_f(x; \lambda c(x))$

Cost-aware Bayesian Optimization   Budgeted Pandora's Box

| Continuous | $\Longleftarrow$ | Discrete |
| Correlated | $\Longleftarrow$ | Independent |
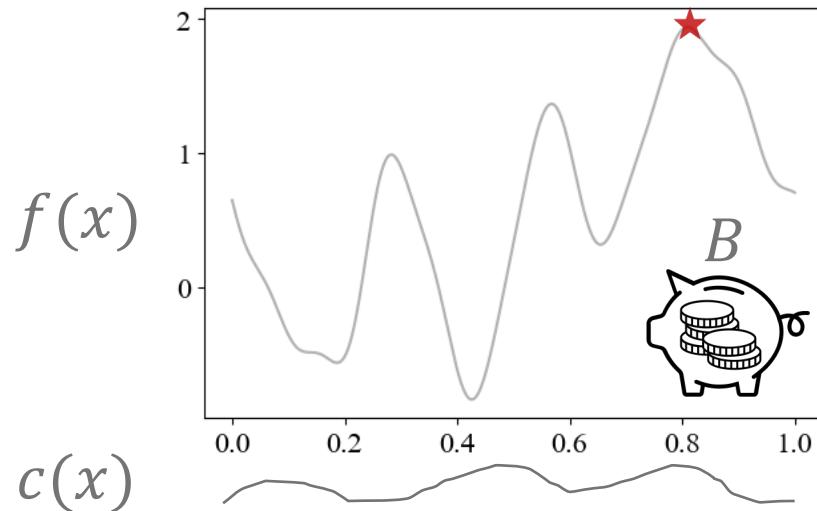
incorporate posterior

scaling factor

Our policy: $\mathrm{GI}_{f|D}(x; \lambda c(x))$ $\Longleftarrow$ Optimal policy: $\mathrm{GI}_f(x; \lambda c(x))$

Cost-aware Bayesian Optimization   Pandora's Box

Continuous	$\Leftarrow$	Discrete ✓
Correlated	$\Leftarrow$	Independent ✓
Expected-budget-constrained	$\Leftarrow$	Cost-per-sample ✓

incorporate posterior

Our policy: $\text{GI}_{f|D}(x; \lambda c(x))$ $\underset{\text{scale costs}}{\Longleftarrow}$ Optimal policy: $\text{GI}_f(x; c(x))$
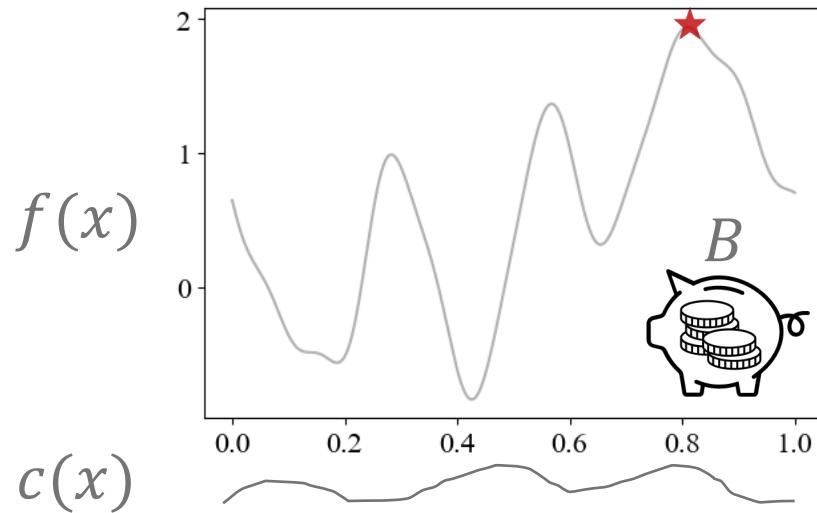
Cost-aware Bayesian Optimization

Pandora's Box

$f(x)$

$c(x)$

$\Longleftarrow$

$f(x)$

$c(x)$

Continuous	Discrete
Correlated	Independent
Expected-budget-constrained	Cost-per-sample

incorporate posterior

Our policy: $\text{GI}_{f|D}(x; \lambda c(x))$ $\Longleftarrow$ Optimal policy: $\text{GI}_f(x; c(x))$

scale costs

How to compute?

How to compute Gittins index?

$f(x)$

$c(x)$

g

$$\downarrow \; \textcolor{red}{\mathbf{?}}$$
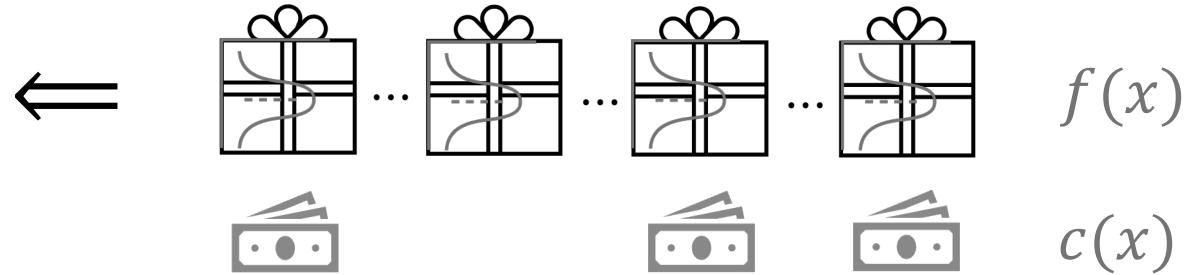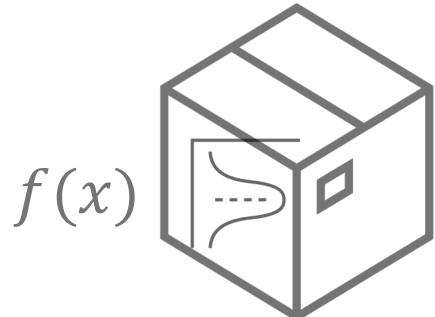
$$\mathrm{GI}_f(x; c(x))$$

$$\downarrow$$

$$g$$

How to compute Gittins index?
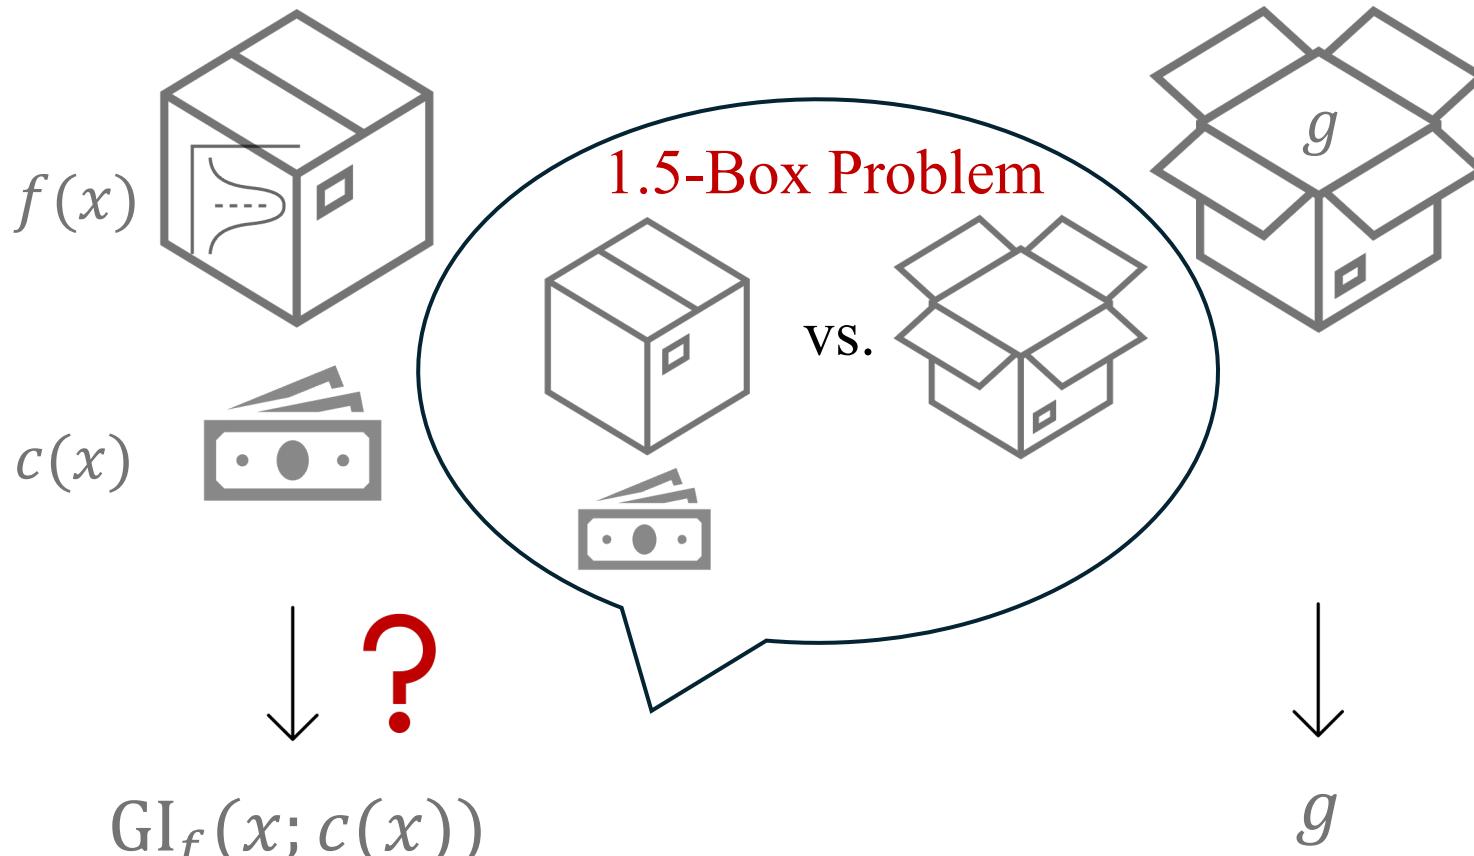
$f(x)$

$c(x)$

1.5-Box Problem

vs.

g

?

$\mathrm{GI}_f(x; c(x))$

g

Whether to open a new box or take current best?

Gittins Index Computation: 1.5-Box Problem

$f(x)$

$c(x)$

vs.

g
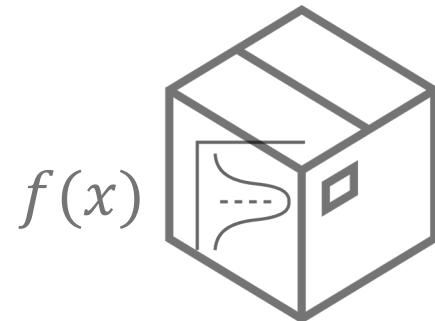
Open closed box
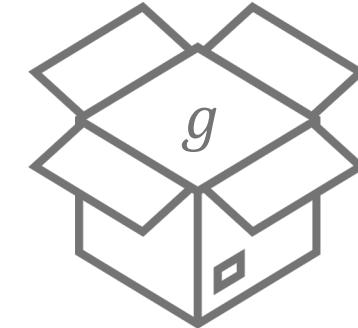
Take opened box

Whether to open a new box or take current best?
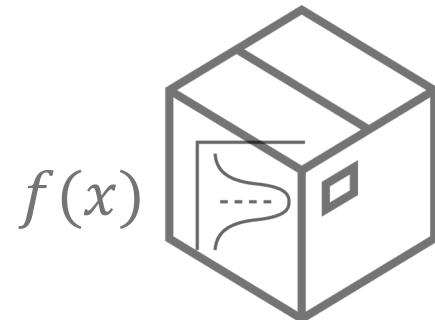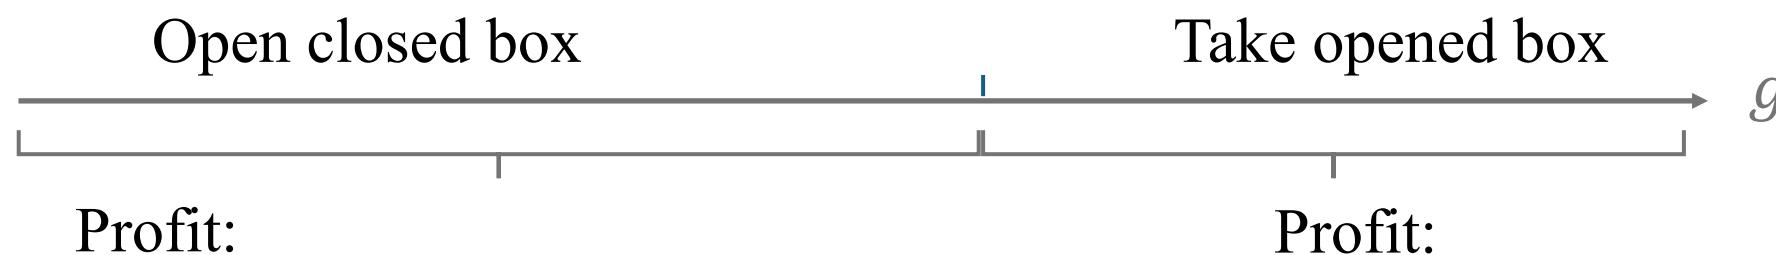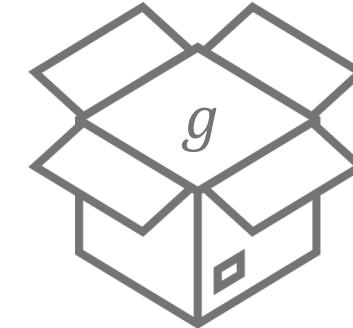
Gittins Index Computation: 1.5-Box Problem

$f(x)$

$c(x)$

vs.

g

Open closed box

Take opened box

g

Profit:

Profit:

$$\begin{cases} \text{take inside value}, f(x) > g \\ \text{take outside option}, f(x) \leq g \end{cases}$$

Gittins Index Computation: 1.5-Box Problem

$f(x)$

$c(x)$

vs.

g

Open closed box

Take opened box

g
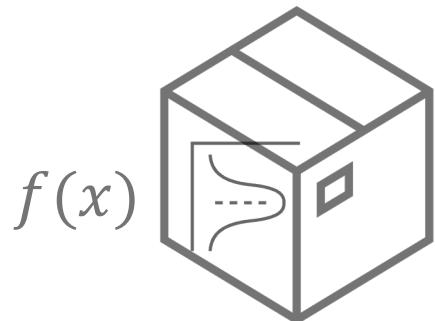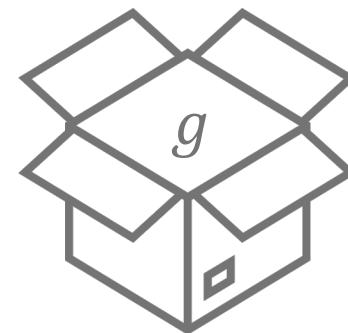
Profit: $\mathbb{E}[\max(f(x), g)] - c(x)$

Profit: g

$\begin{cases} \text{take inside value}, f(x) > g \\ \text{take outside option}, f(x) \le g \end{cases}$

Gittins Index Computation: 1.5-Box Problem

$f(x)$ vs. g

$c(x)$

Gittins index

Open closed box $\mathrm{GI}_f(x; c(x))$ Take opened box

g

Profit: $\mathbb{E}[\max(f(x), g)] - c(x)$ Profit: g
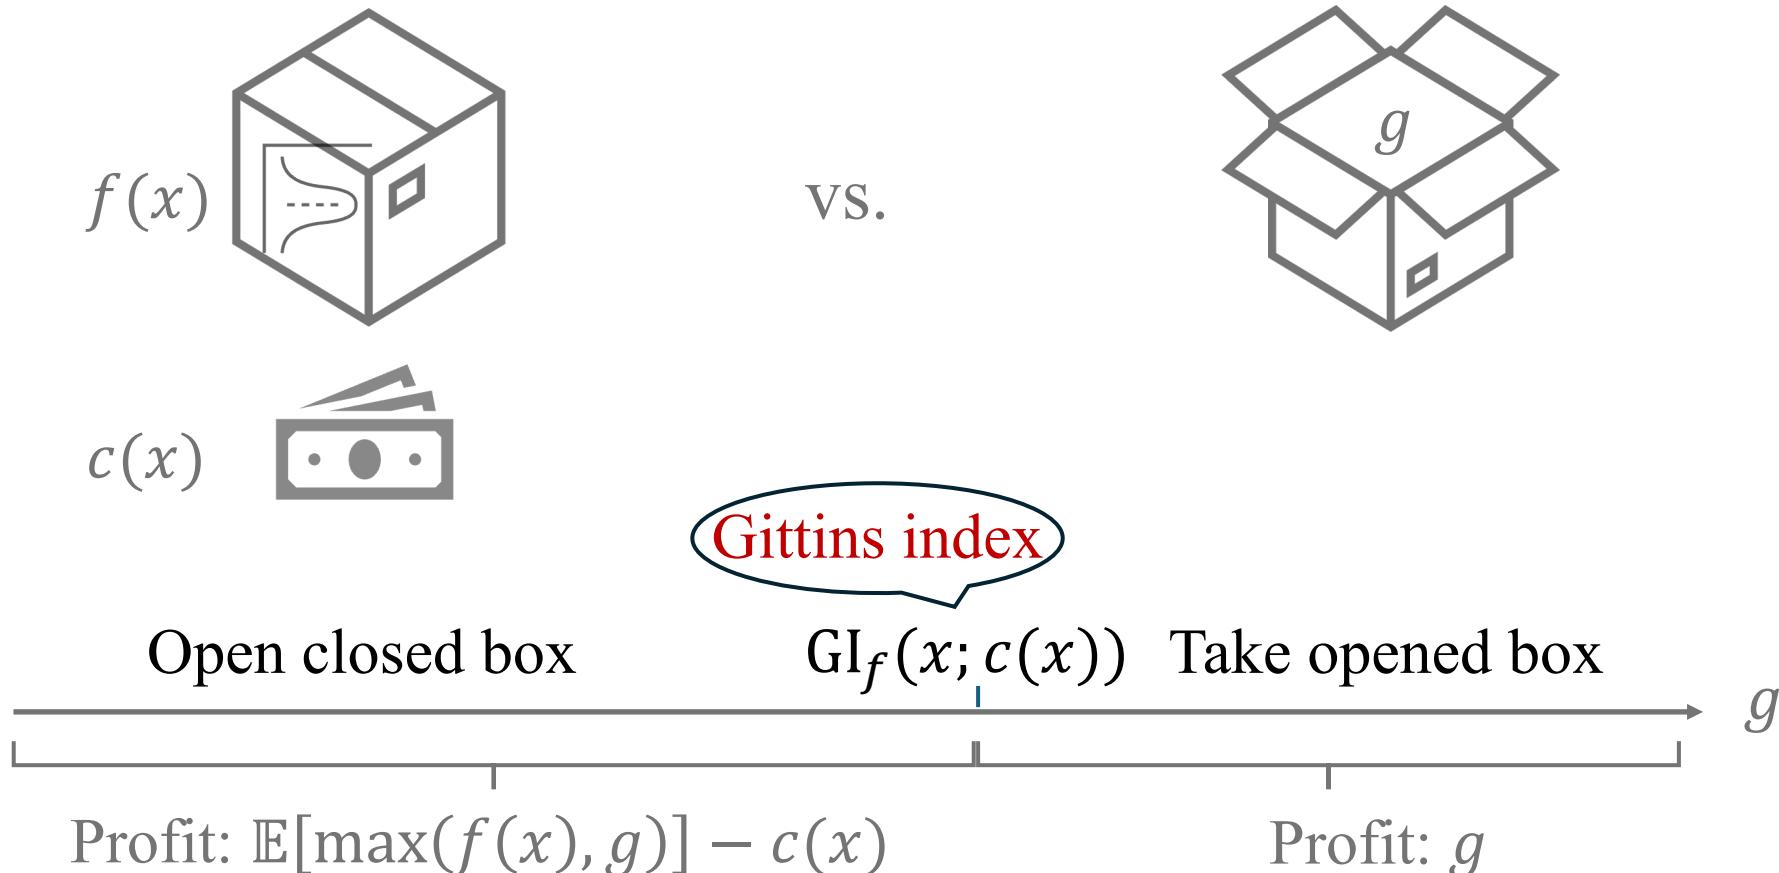
$\mathrm{GI}_f(x; c(x))$: solution g to $\mathbb{E}[\max(f(x), g)] - c(x) = g$

Gittins Index Computation: 1.5-Box Problem

$f(x)$ vs. g

$c(x)$

Gittins index

Open closed box $\text{GI}_f(x; c(x))$ Take opened box g

Profit: $\mathbb{E}[\max(f(x), g)] - c(x)$ Profit: g

$\text{GI}_f(x; c(x))$: solution g to $\mathbb{E}[\max(f(x), g)] - c(x) = g$
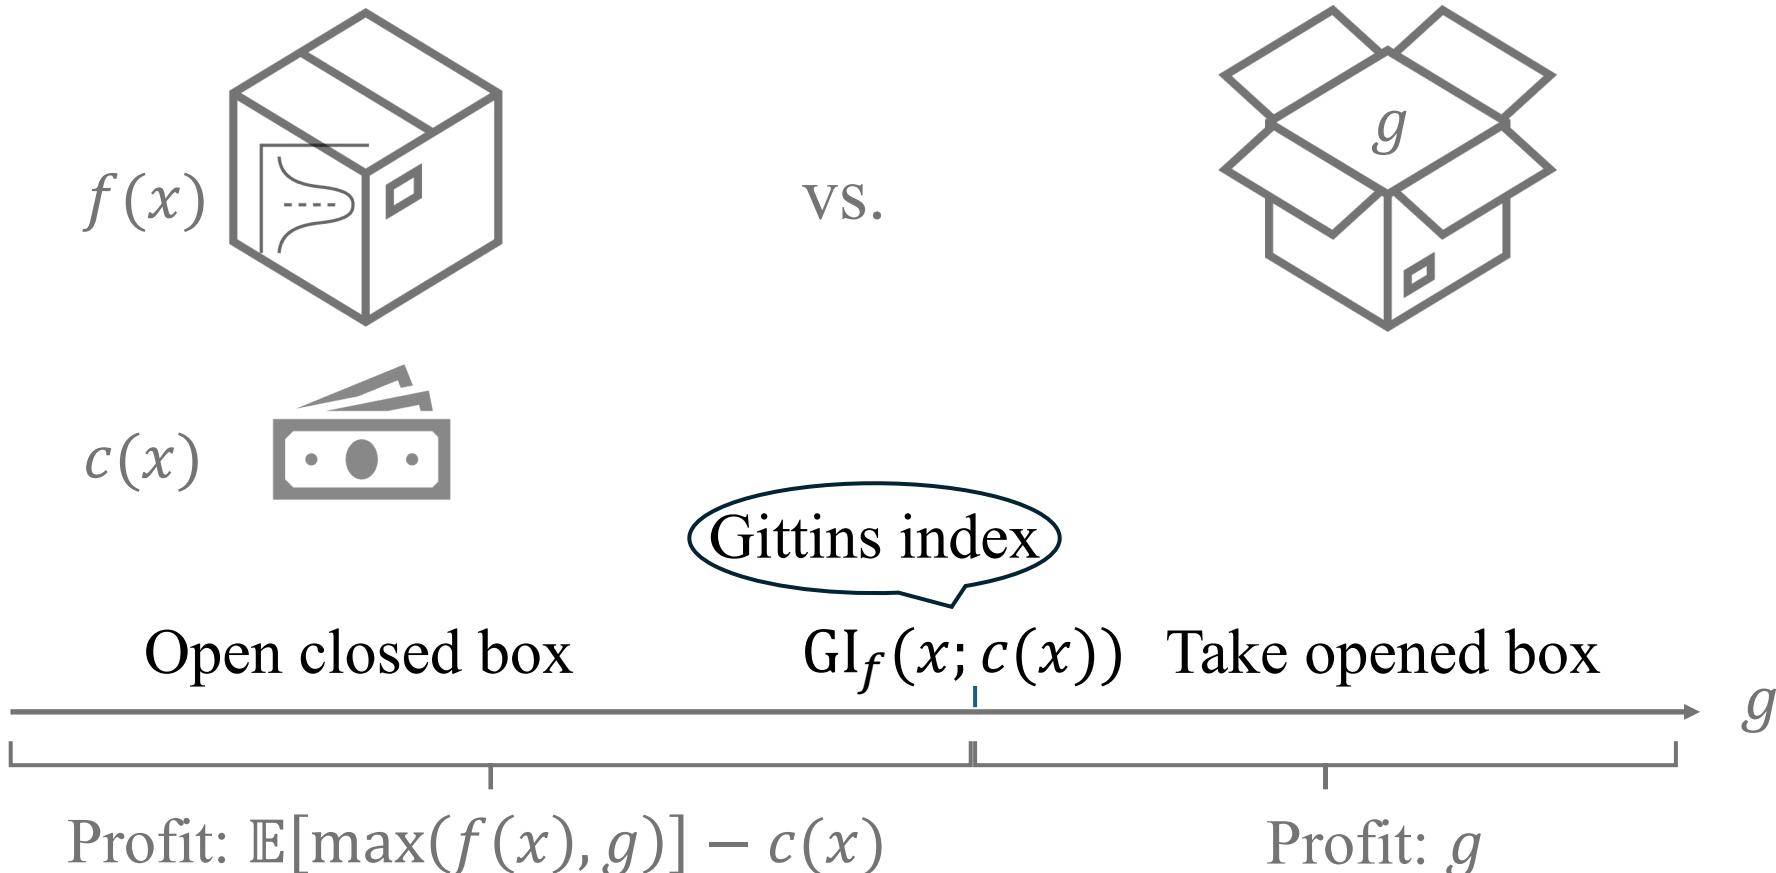
Larger the cost, smaller the Gittins index

Gittins Index Computation: 1.5-Box Problem
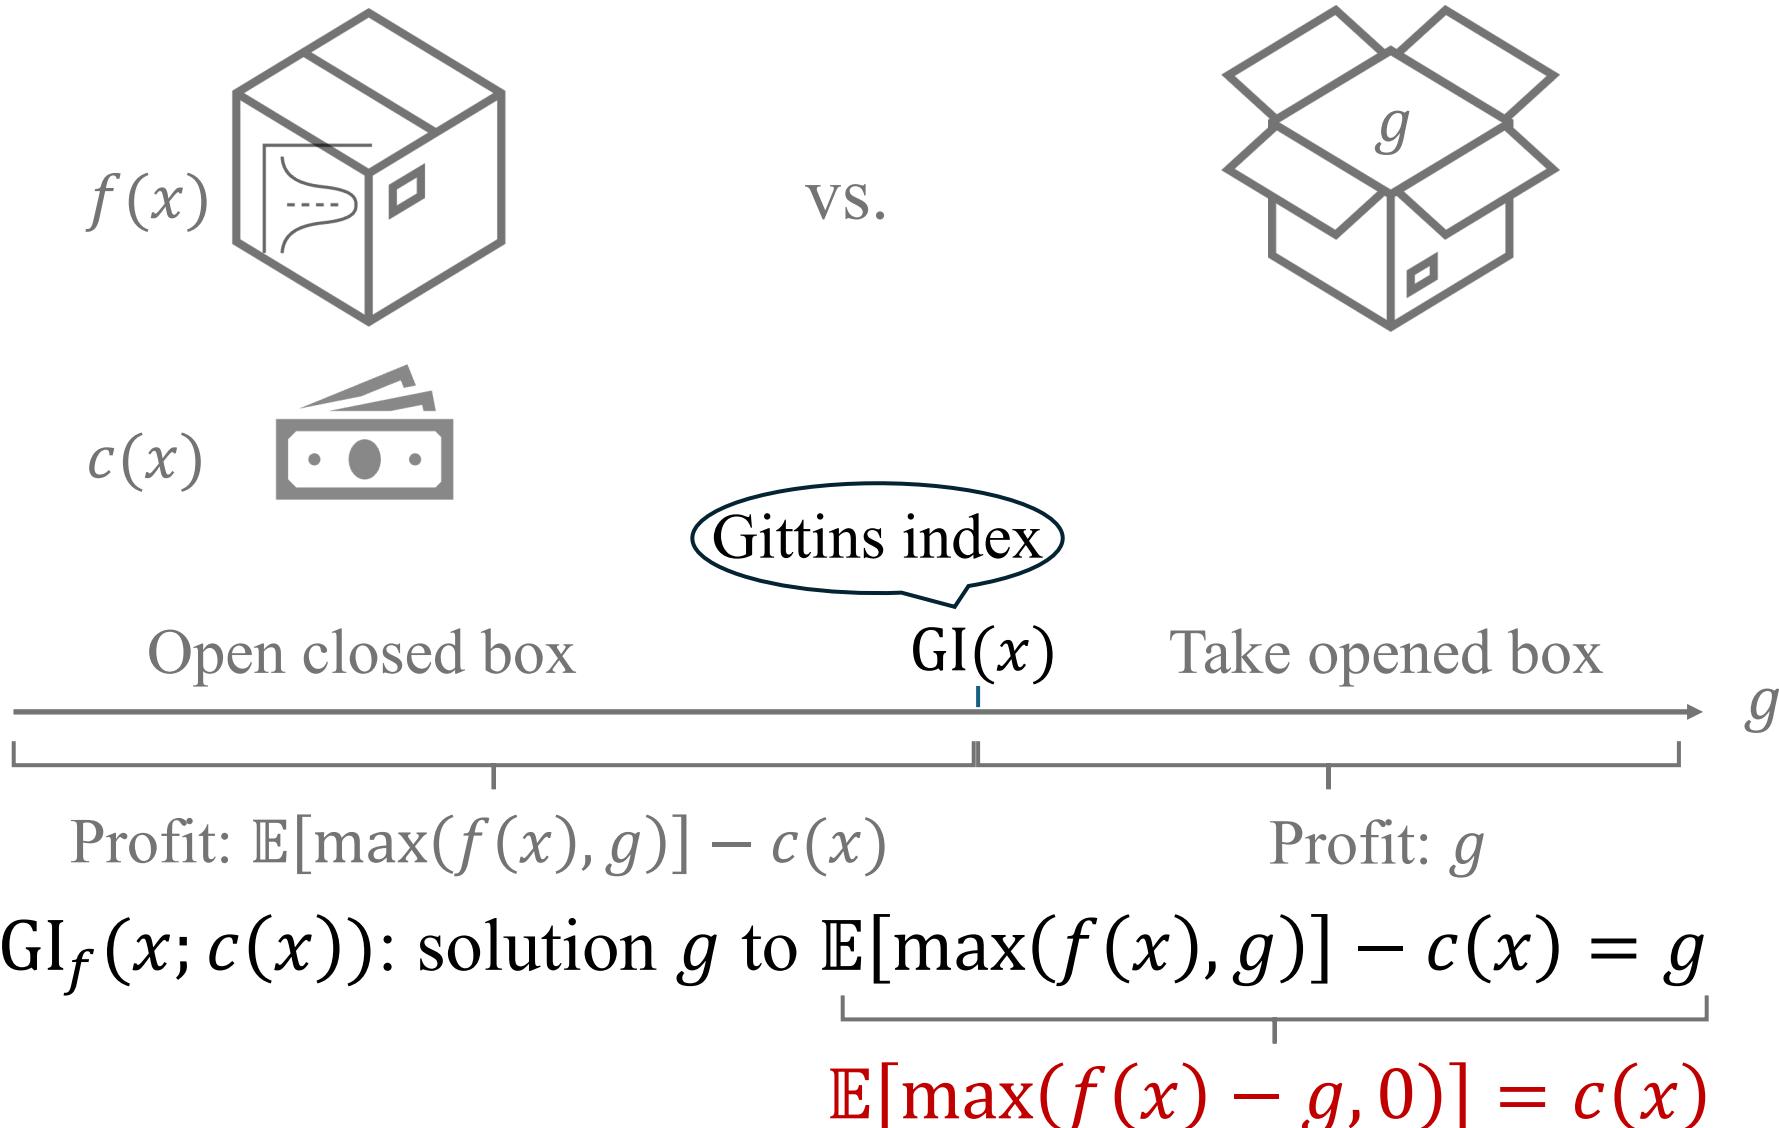
$f(x)$

vs.

g

$c(x)$

Gittins index

Open closed box $\text{GI}(x)$ Take opened box

g

Profit: $\mathbb{E}[\max(f(x), g)] - c(x)$ Profit: g

$\text{GI}_f(x; c(x))$: solution g to $\mathbb{E}[\max(f(x), g)] - c(x) = g$

$\mathbb{E}[\max(f(x) - g, 0)] = c(x)$

Gittins Index Computation: 1.5-Box Problem

$f(x)$ vs. g

$c(x)$

Gittins index

Open closed box $\mathrm{GI}(x)$ Take opened box

g

Profit: $\mathbb{E}[\max(f(x), g)] - c(x)$ Profit: g

$\mathrm{GI}_f(x; c(x))$: solution g to $\mathbb{E}[\max(f(x), g)] - c(x) = g$

$\mathrm{EI}_f(x; g)$ $\mathbb{E}[\max(f(x) - g, 0)] = c(x)$

$$\text{EI}_{f|D}(x; g) := \textcolor{red}{\mathbb{E}[\max(f(x) - g, 0) \mid D]} \qquad \text{GI}_{f|D}(x) = g \text{ s.t. } \text{EI}_{f|D}(x; g) = \textcolor{red}{\lambda c(x)}$$

$$\mathrm{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D] \qquad \mathrm{GI}_{f|D}(x) = g \text{ s.t. } \mathrm{EI}_{f|D}(x; g) = \lambda c(x)$$

Gaussian distribution

$$\text{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D]$$
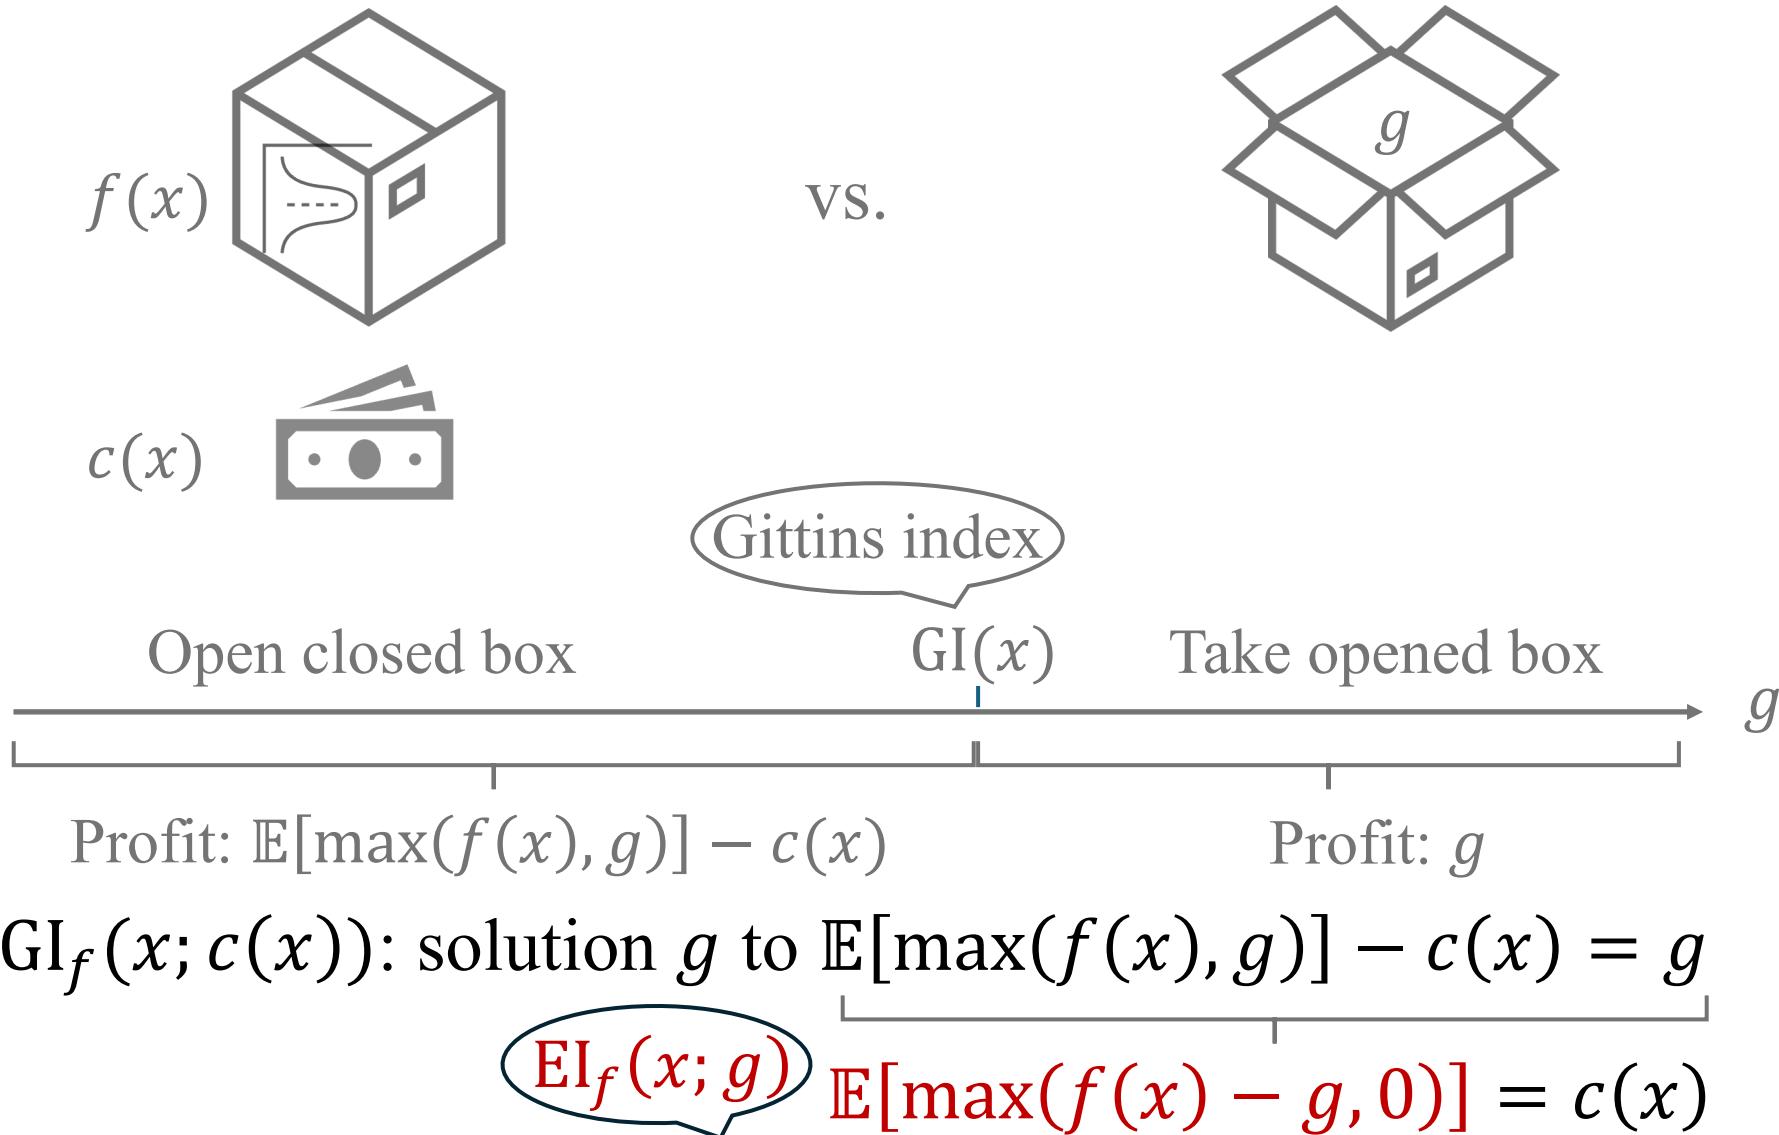
Gaussian distribution — analytical expression
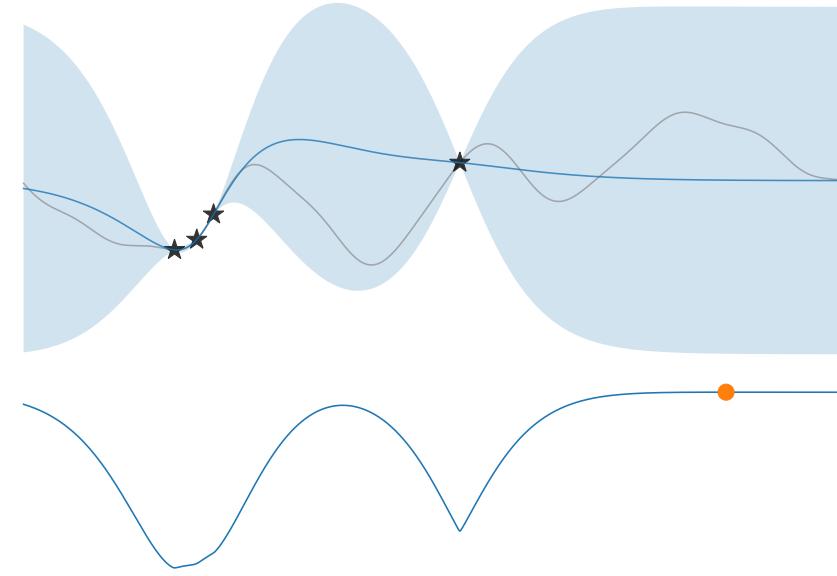
$$\text{GI}_{f|D}(x) = g \text{ s.t. } \text{EI}_{f|D}(x; g) = \lambda c(x)$$

Gittins Index
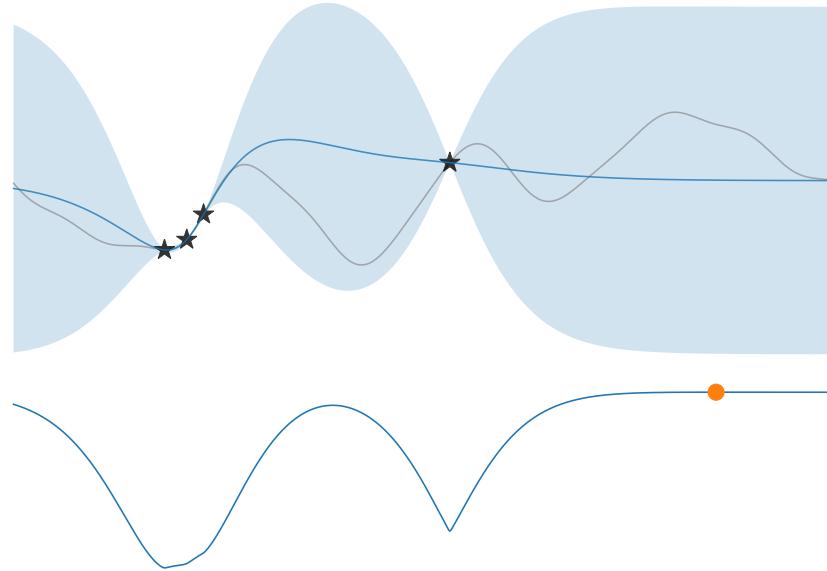
$$\mathrm{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D] \qquad \mathrm{GI}_{f|D}(x) = g \text{ s.t. } \mathrm{EI}_{f|D}(x; g) = \lambda c(x)$$

Gaussian distribution · analytical expression

$\mathrm{GI}_{f|D}$ is easy to compute using $\mathrm{EI}_{f|D}$ + bisection search!

Expected Improvement
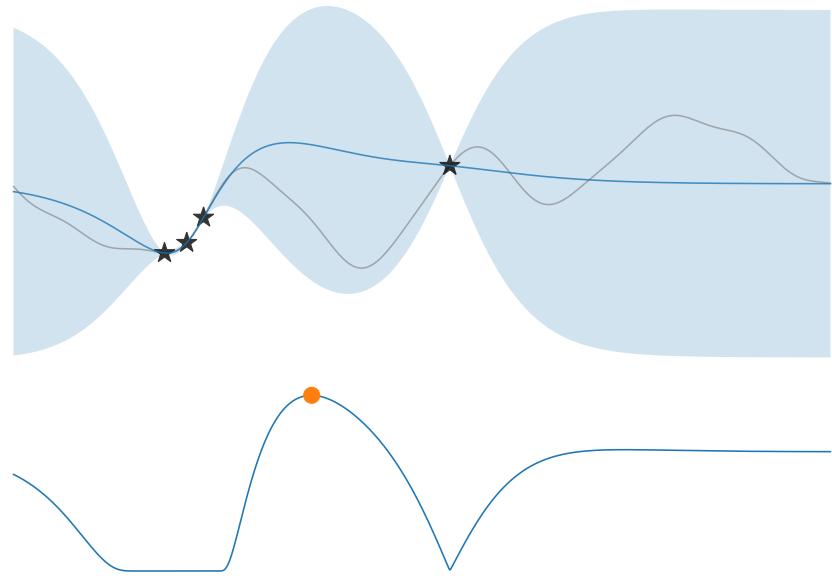
Gittins Index

$$\text{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D]$$

$$\max_x \text{EI}_{f|D}(x; {\color{red}y_{\text{best}}})$$
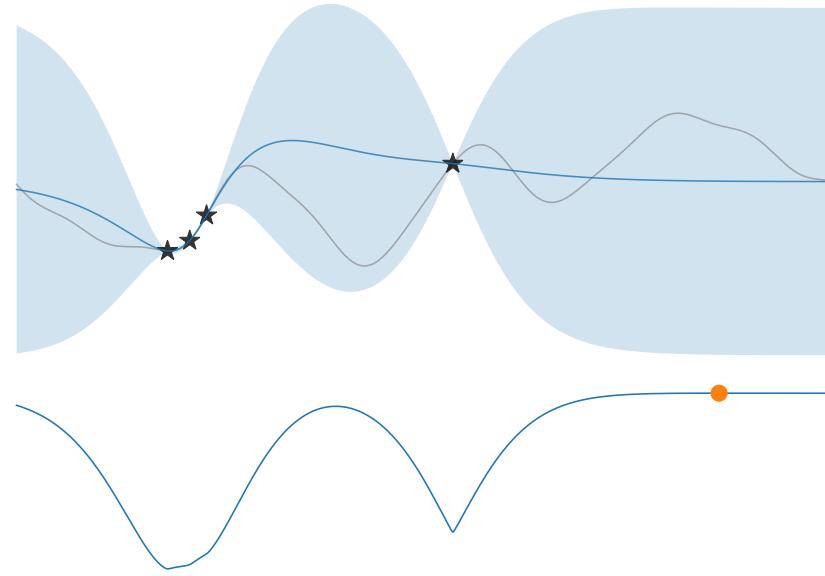
$$\text{GI}_{f|D}(x) = g \text{ s.t. } \text{EI}_{f|D}(x; g) = \lambda c(x)$$

$$\max_x \text{GI}_{f|D}(x; {\color{red}\lambda c(x)})$$

Expected Improvement
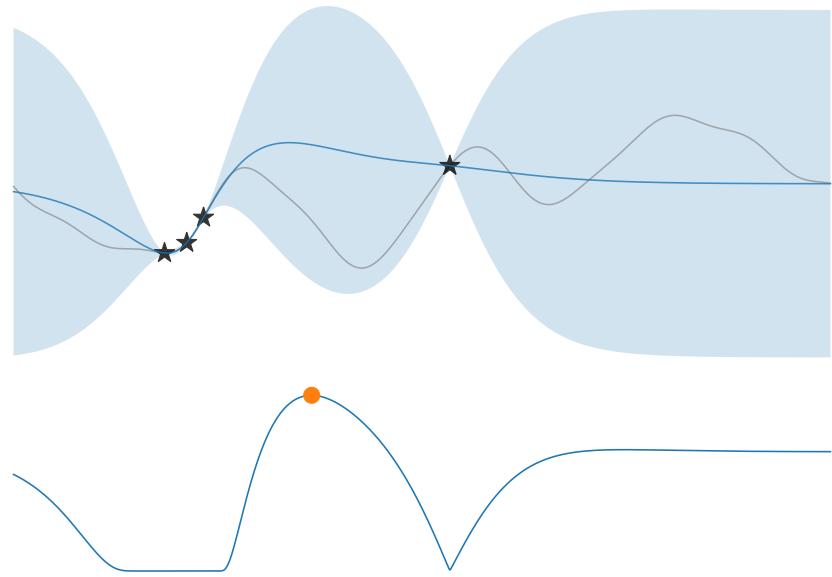
$$\mathrm{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D]$$

$$\max_x \mathrm{EI}_{f|D}(x; \color{red}{y_{\mathrm{best}}})$$

Cost-unaware

Gittins Index
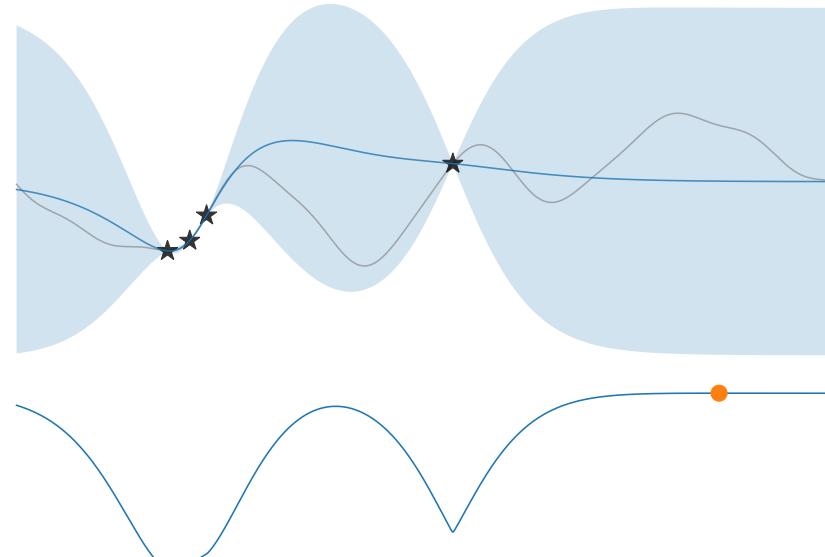
$$\mathrm{GI}_{f|D}(x) = g \text{ s.t. } \mathrm{EI}_{f|D}(x; g) = \lambda c(x)$$

$$\max_x \mathrm{GI}_{f|D}(x; \color{red}{\lambda c(x)})$$

Cost-aware

Expected Improvement
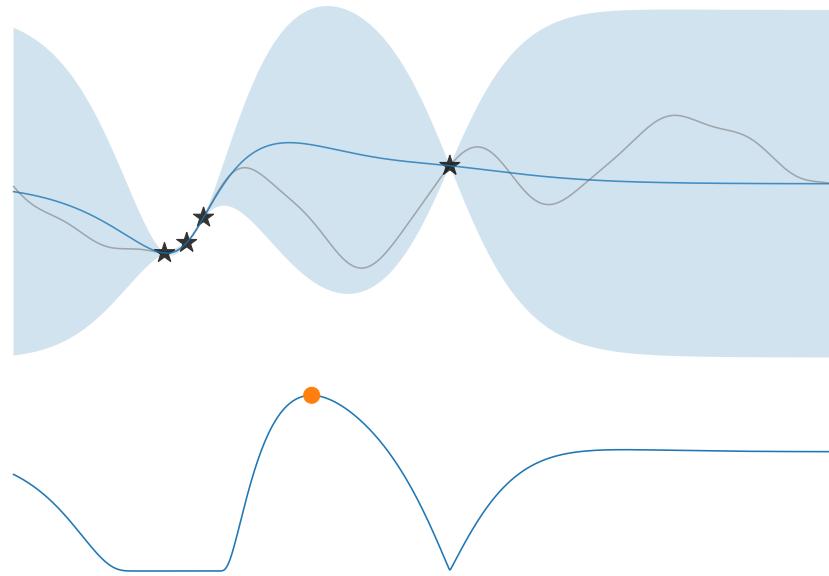
Gittins Index

$$\mathrm{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D]$$

$$\max_x \mathrm{EI}_{f|D}(x; {\color{red}y_{\mathrm{best}}})$$

Cost-unaware

One-step approximation to MDP

$$\mathrm{GI}_{f|D}(x) = g \text{ s.t. } \mathrm{EI}_{f|D}(x; g) = \lambda c(x)$$

$$\max_x \mathrm{GI}_{f|D}(x; \lambda c(x))$$

Cost-aware

Anytime-lookahead

Expected Improvement
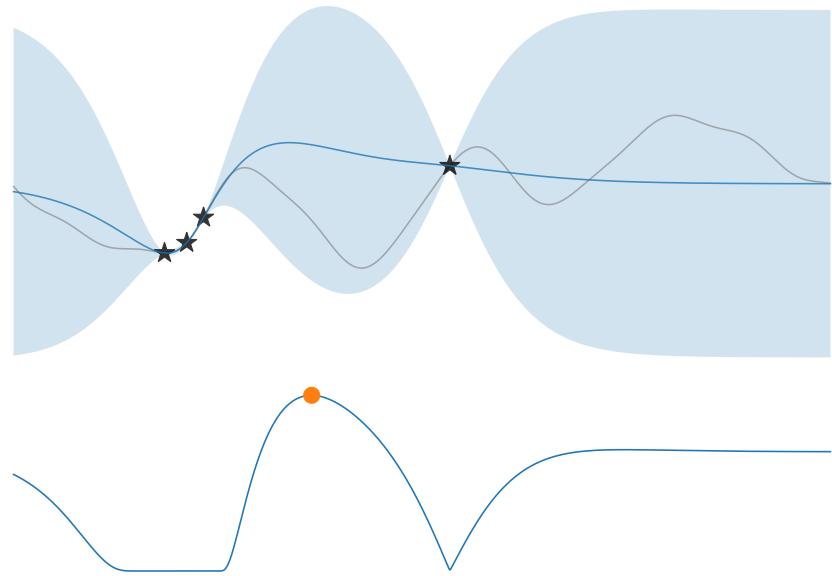
$$\mathrm{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D]$$

$$\max_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$$

Cost-unaware

One-step approximation to MDP

Temporal simplification to MDP

Gittins Index
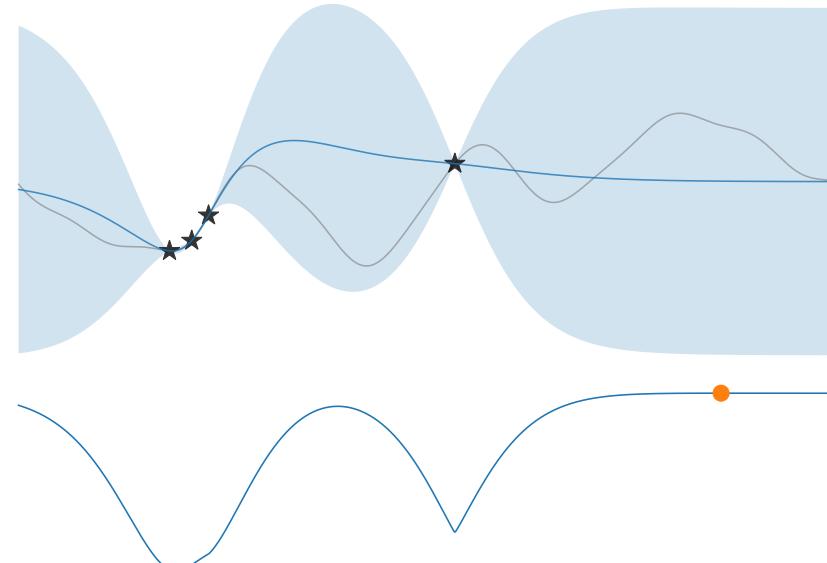
$$\mathrm{GI}_{f|D}(x) = g \text{ s.t. } \mathrm{EI}_{f|D}(x; g) = \lambda c(x)$$

$$\max_x \mathrm{GI}_{f|D}(x; \lambda c(x))$$

Cost-aware

Anytime-lookahead

Spatial simplification to MDP

Cost-aware Bayesian Optimization

Pandora's Box

$f(x)$

$c(x)$

B

$f(x)$

$c(x)$

Continuous

Discrete
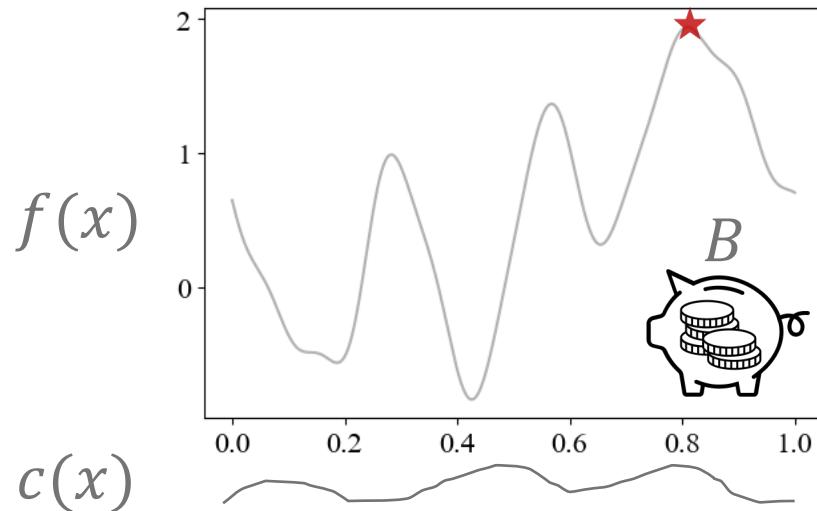
Correlated

Independent

Expected-budget-constrained

Cost-per-sample

incorporate posterior

Is Gittins index good? $\Longleftarrow$ Gittins index is optimal

scale costs

Cost-aware Bayesian Optimization

$f(x)$

B

$c(x)$

Continuous

Correlated

Expected-budget-constrained

Is Gittins index **empirically** good?

this work

Pandora's Box

$f(x)$

$c(x)$

Discrete

Independent

Cost-per-sample

Gittins index is optimal

Experiment Setup: Objective Functions

Samples from prior

Synthetic

Figure from *infpy*
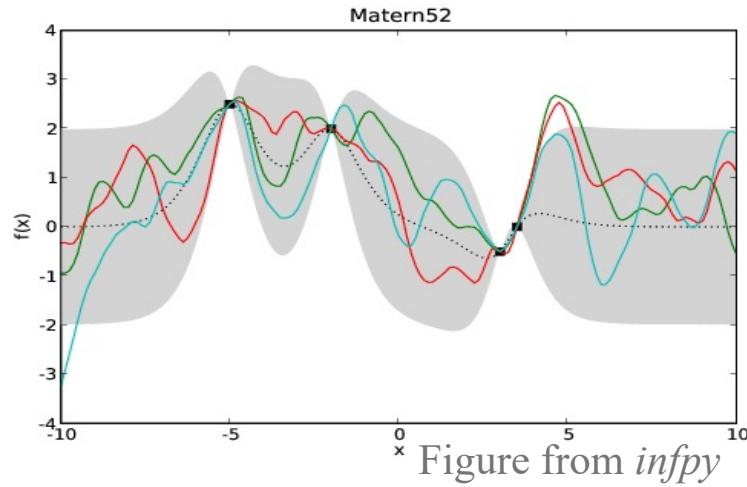
Ackley function
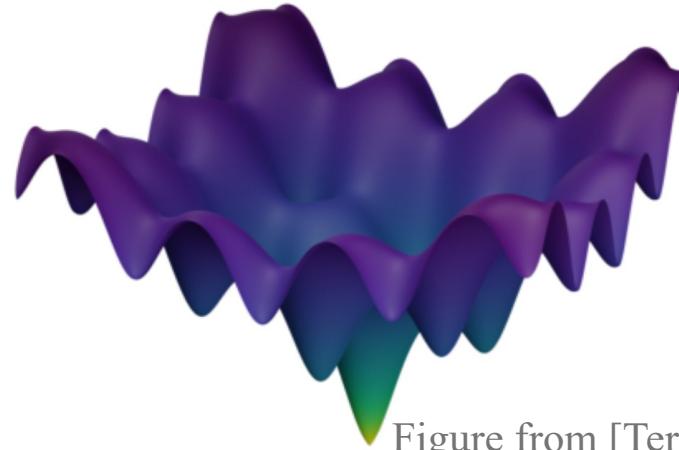
Figure from [Terenin'22]

Pest Control

Empirical

Figure from ChatGPT

Lunar Lander
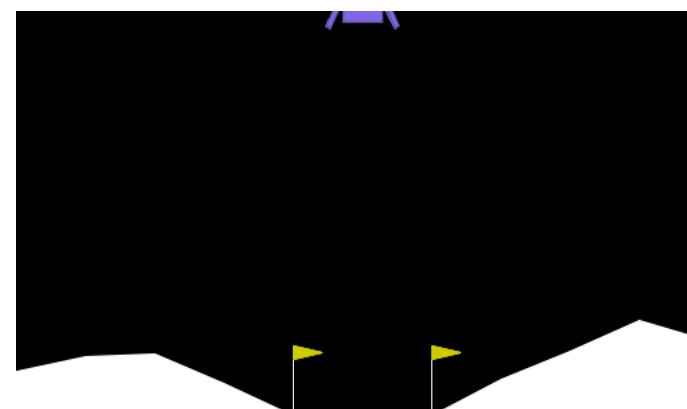
Figure from OpenAI Gym

Experiment Results: PBGI vs Baselines

Synthetic
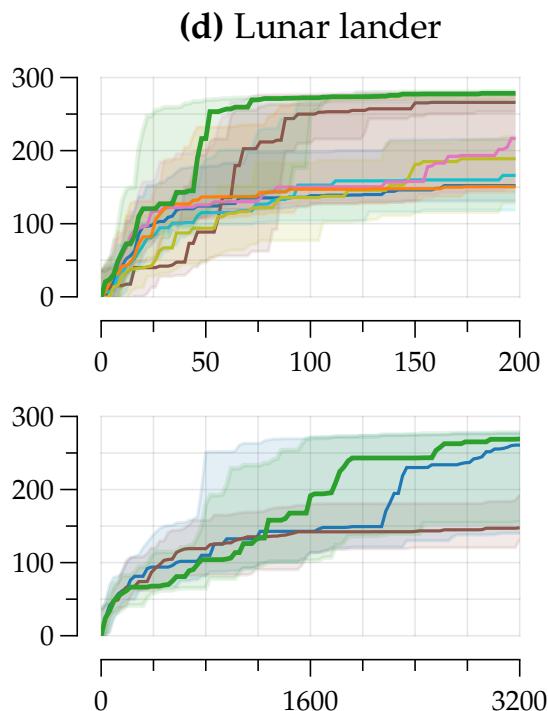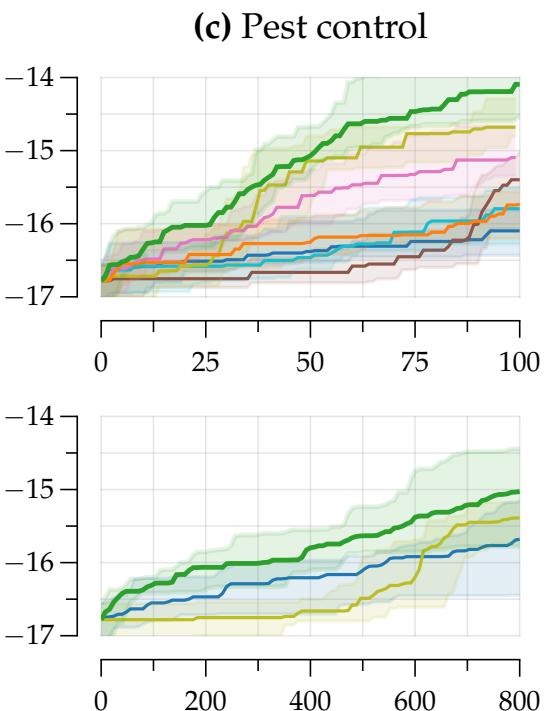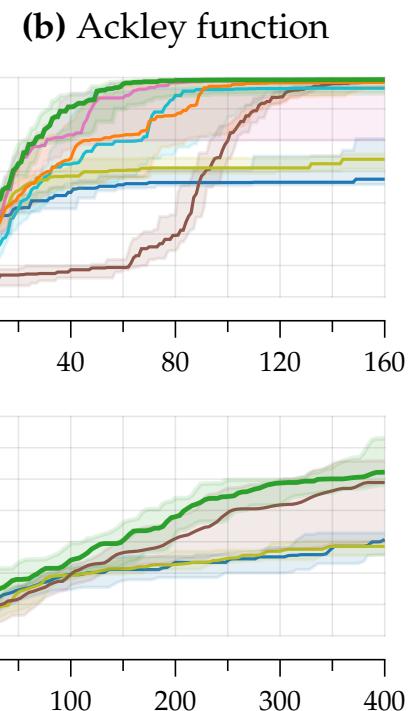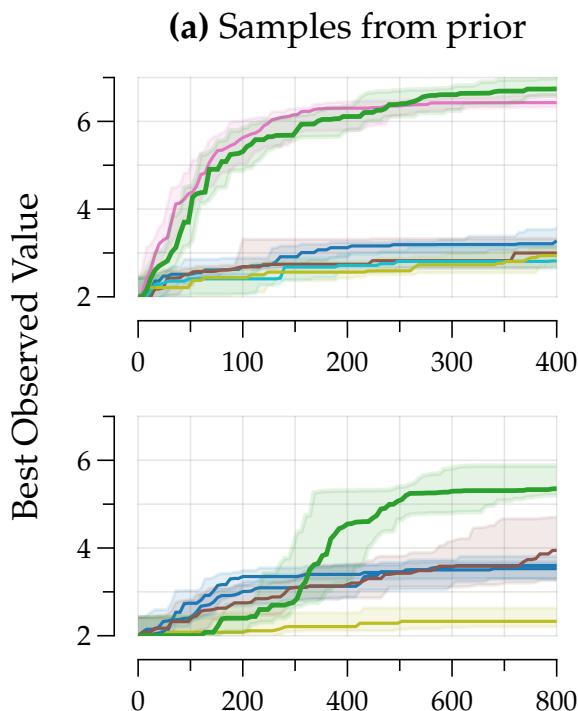
Empirical

(a) Samples from prior   **(b)** Ackley function   **(c)** Pest control   **(d)** Lunar lander

Best Observed Value

Uniform Costs

Varying Costs

Cumulative Cost

— Pandora's Box Gittins Index (PBGI)
— Upper Confidence Bound (UCB)
— Expected Improvement (EI)
— EI Per Unit Cost (EIPC)
— Thompson Sampling
— Knowledge Gradient
— (Budgeted) Multi-Step EI
— (Multi-fidelity) Max-value Entropy Search

Cost-aware Bayesian Optimization

$f(x)$

B

$c(x)$
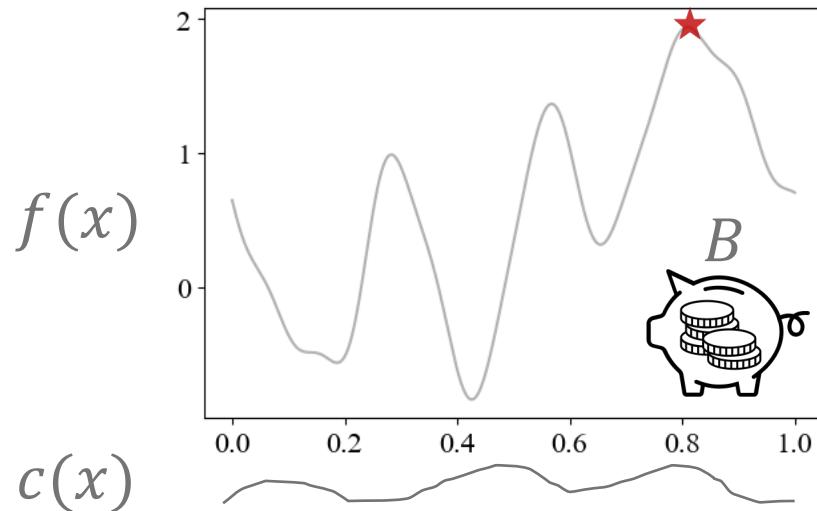
Continuous

Correlated

Expected-budget-constrained

Is Gittins index theoretically good?

Pandora's Box

$f(x)$

$c(x)$
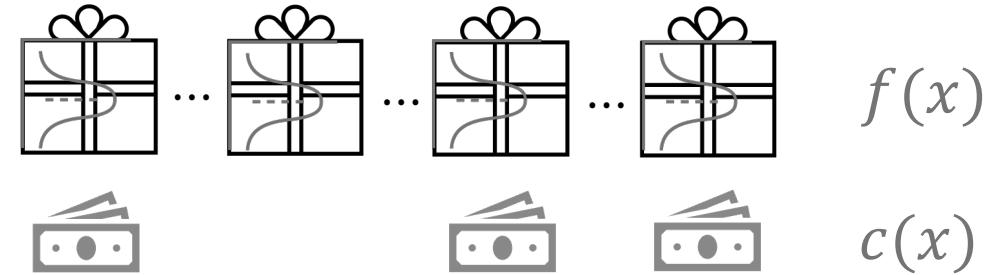
Discrete

Independent

Cost-per-sample

Gittins index is optimal

Cost-aware Bayesian Optimization
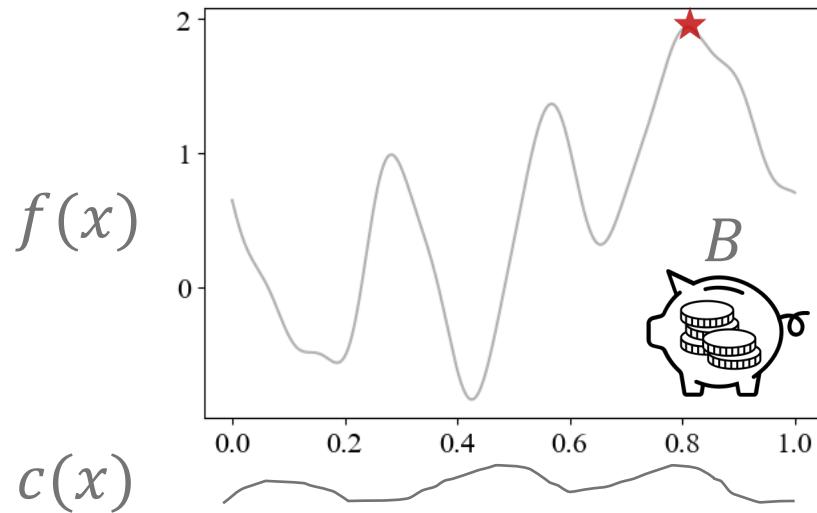
Pandora's Box

$f(x)$

B

$c(x)$

$f(x)$

$c(x)$

Continuous

Correlated

Expected-budget-constrained

Is Gittins index **theoretically** good?

ongoing work

Discrete

Independent

Cost-per-sample

Gittins index is optimal

Timing Experiment

$d = 4$ $d = 8$ $d = 16$

Runtime

Cumulative Cost

—— EI   —— MSEI   —— PBGI   —— TS   —— KG
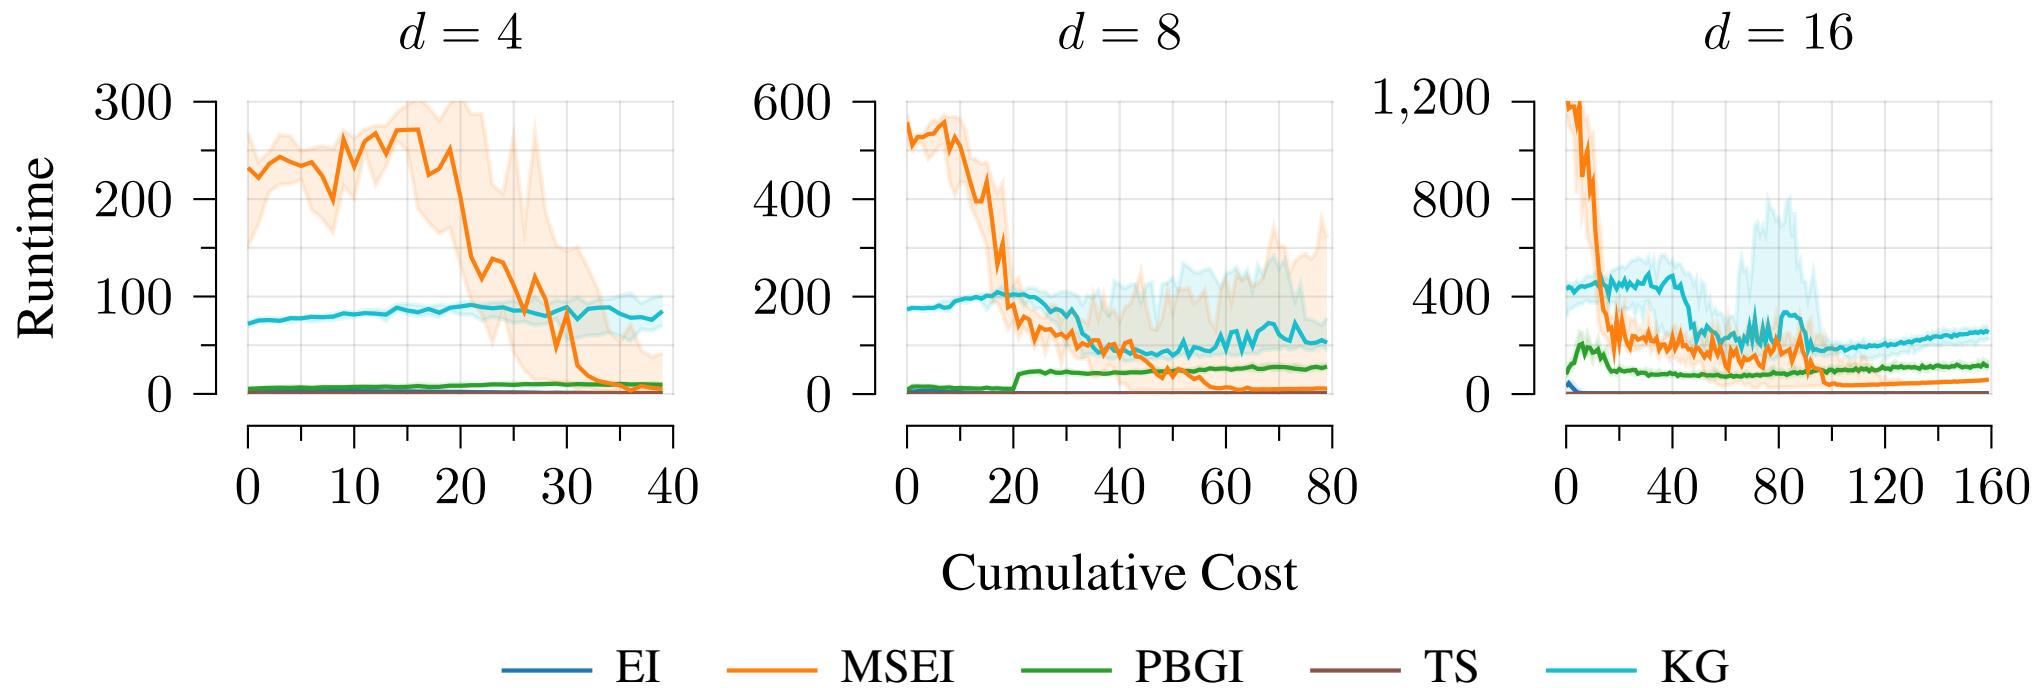
$$\mathrm{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D] \quad \mathrm{GI}_{f|D}(x) = g \text{ s.t. } \mathrm{EI}_{f|D}(x; g) = \lambda c(x)$$

PBGI is easy to compute using EI + bisection search!

Effect of λ
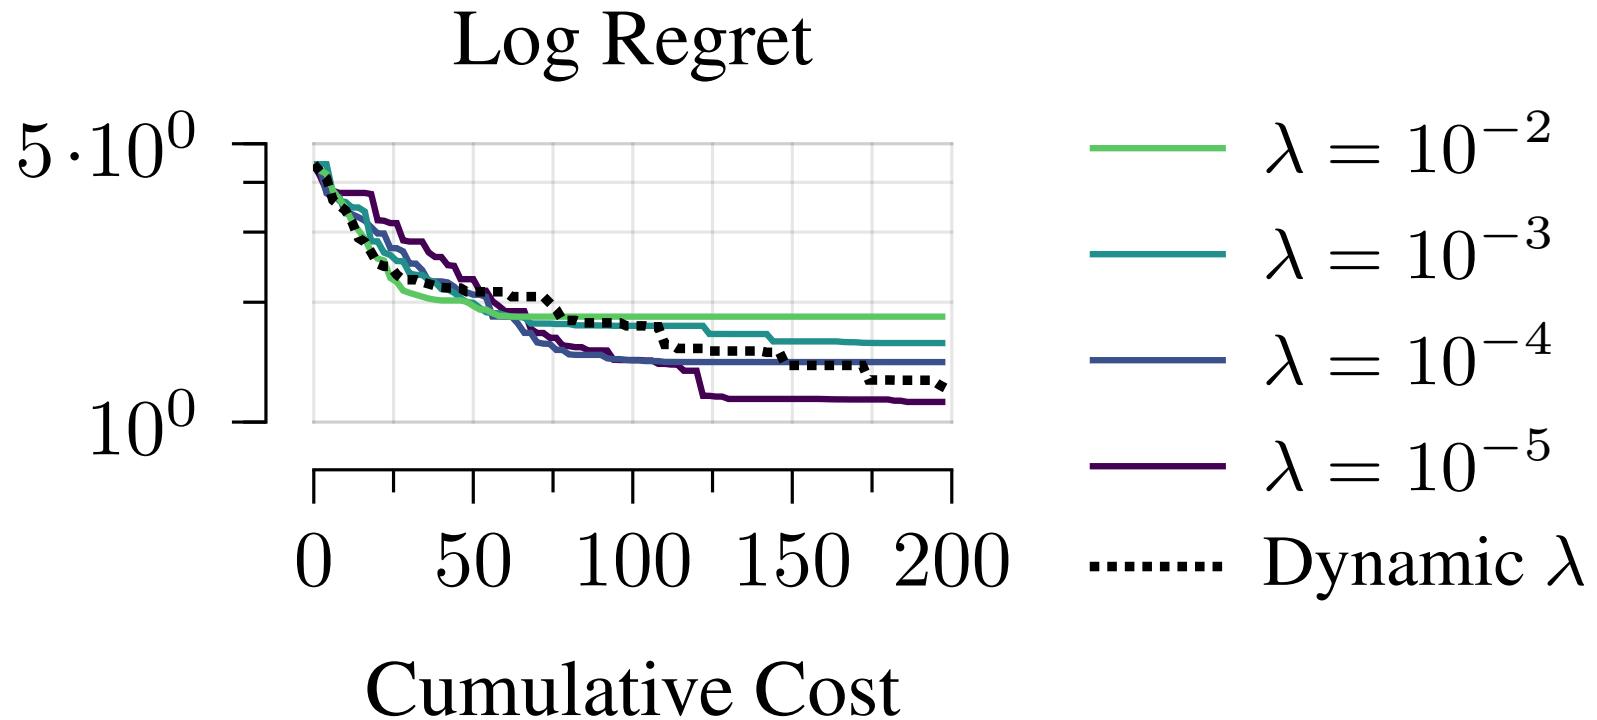

Larger budget, smaller λ, higher exploration

EI vs PBGI

EI

PBGI (large λ)

PBGI (small λ) $\approx$UCB

std
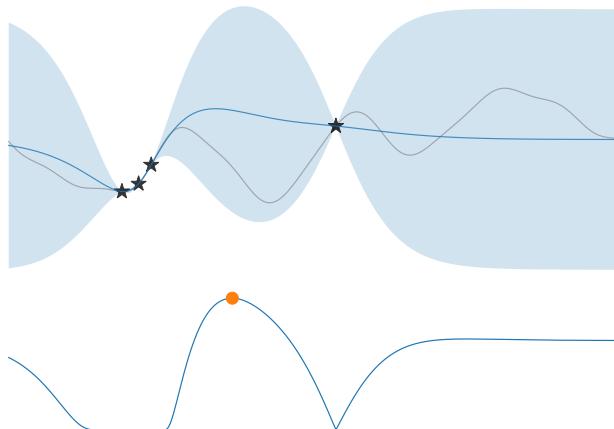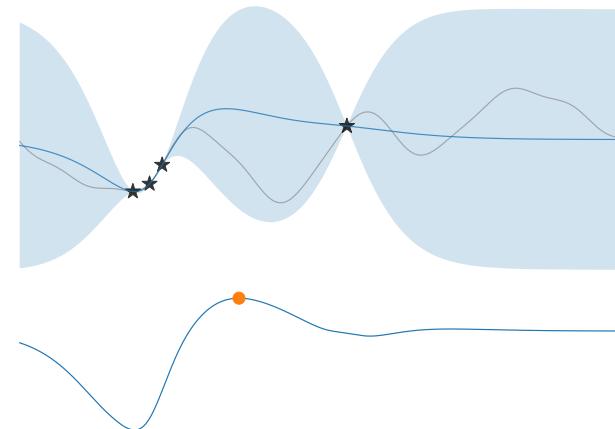
mean

mean

mean

Part II: Observable Partial Feedback

Hyperparameter tuning:
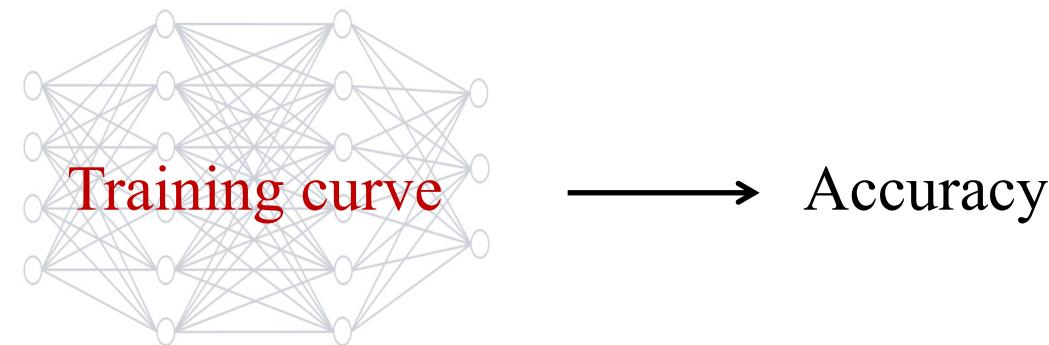
Training parameters $\longrightarrow$ Training curve $\longrightarrow$ Accuracy

Plasma physics:

Reactor parameters $\longrightarrow$ Convergence curve $\longrightarrow$ Stability

Practical Factor: Observable Partial Feedback

Hyperparameter tuning:
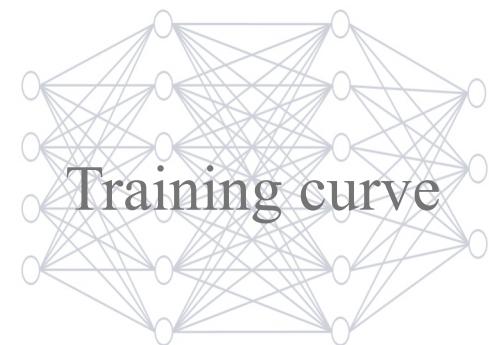
Training parameters $\longrightarrow$ Training curve $\longrightarrow$ Accuracy

Input parameters $\longrightarrow$ Performance metric

Plasma physics:

Reactor parameters $\longrightarrow$ Convergence curve $\longrightarrow$ Stability

Practical Factor: Observable Partial Feedback

Input parameters $\longrightarrow$ Performance curve $\longrightarrow$ Performance metric

Practical Factor: Observable Partial Feedback

Black-box process:

$x \longrightarrow$ $\{f(x,t)\}_{t \leq T_x}$ $\longrightarrow f(x, T_x)$

full evaluation time of hyperparameter x

Optimizing Black-box Processes

Black-box process:

$$x \longrightarrow \boxed{\{f(x,t)\}_{t \leq T_x}} \longrightarrow f(x, T_x)$$

full evaluation time of hyperparameter x
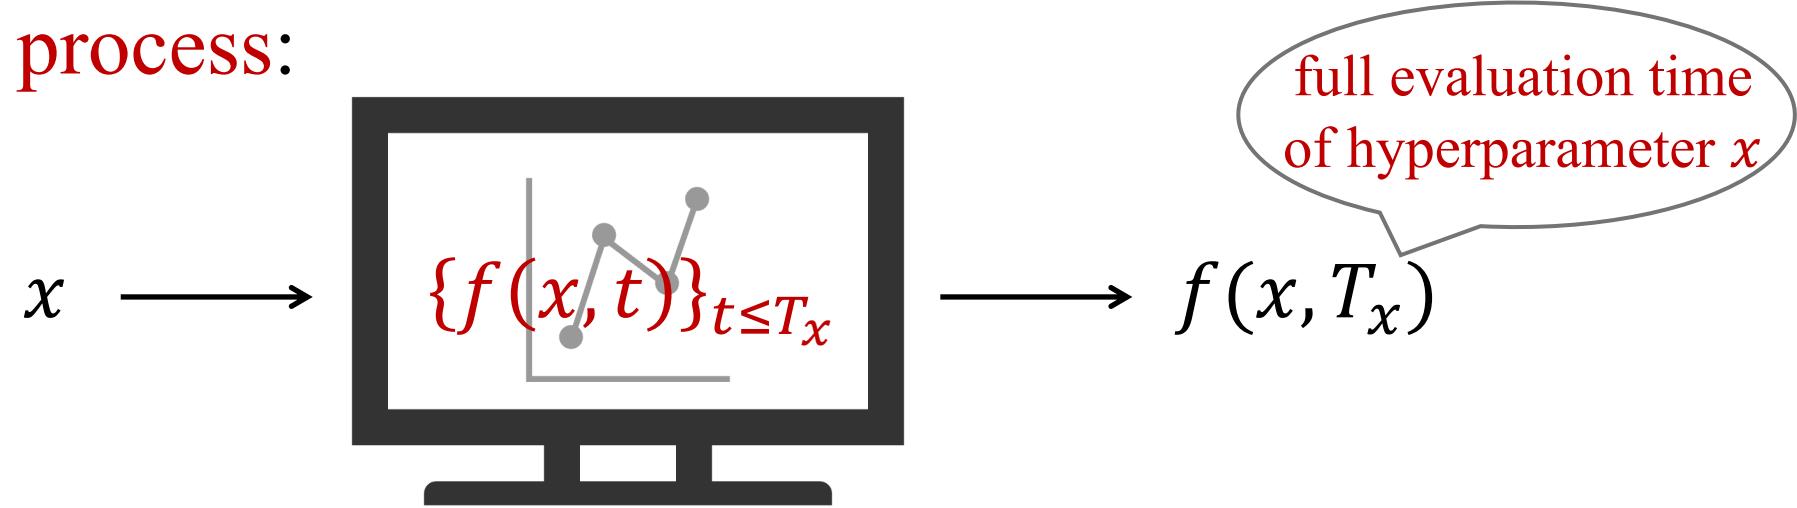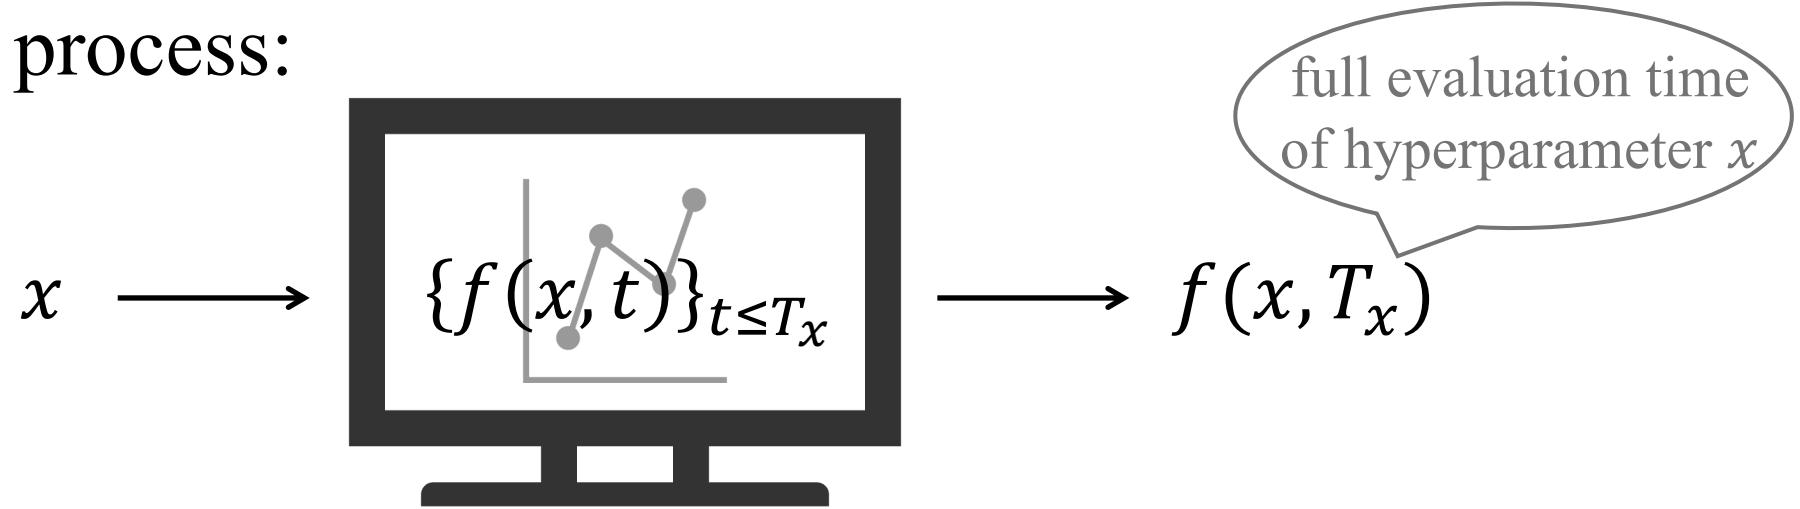
Goal: $\max_{x \in \mathcal{X}} f(x, T_x)$

Optimizing Black-box Processes

Black-box process:

$x \longrightarrow \{f(x,t)\}_{t \leq T_x} \longrightarrow f(x, T_x)$

full evaluation time of hyperparameter x

expensive-to-evaluate

Goal: $\max_{x \in \mathcal{X}} f(x, T_x)$

Freeze-thaw Bayesian Optimization

[Swersky, Snoek, Adams'14]

$$x \longrightarrow \boxed{\quad} \longrightarrow f(x, T_x)$$

full evaluation time of hyperparameter x

expensive-to-evaluate

Goal: $\max_{x \in X} f(x, T_x)$

Allow pauses and resumes of hyperparameter tests

Freeze-thaw Bayesian Optimization

[Swersky, Snoek, Adams'14]

full evaluation time
of hyperparameter x

$x \longrightarrow$ $\longrightarrow f(x, T_x)$

Goal: $\sup_{\text{policy}} \mathbb{E} \max_{\text{completed } x \in \mathcal{X}} f(x, T_x)$

s.t. total evaluation time within budget

Allow pauses and resumes of hyperparameter tests

Freeze-thaw Bayesian Optimization

[Swersky, Snoek, Adams'14]

~Temporal GP

full evaluation time
of hyperparameter x

$$x \longrightarrow \boxed{\quad} \longrightarrow f(x, T_x)$$

Goal: $\sup_{\text{policy}} \mathbb{E} \max_{\text{completed } x \in \mathcal{X}} f(x, T_x)$

s.t. total evaluation time within budget

Allow pauses and resumes of hyperparameter tests

Freeze-thaw Bayesian Optimization

[Swersky, Snoek, Adams'14]

entropy-based

~Temporal GP

full evaluation time of hyperparameter x

$$x \longrightarrow \boxed{\quad} \longrightarrow f(x, T_x)$$

Goal: $\sup_{\text{policy}} \mathbb{E} \max_{\text{completed } x \in \mathcal{X}} f(x, T_x)$

s.t. total evaluation time within budget

Allow pauses and resumes of hyperparameter tests

Freeze-thaw Bayesian Optimization

[Swersky, Snoek, Adams'14]

Our view: lack of a guidance to incorporate partial feedback

Freeze-thaw Bayesian Optimization

[Swersky, Snoek, Adams'14]

entropy-based

~Temporal GP

$$x \longrightarrow \boxed{} \longrightarrow f(x, T_x)$$
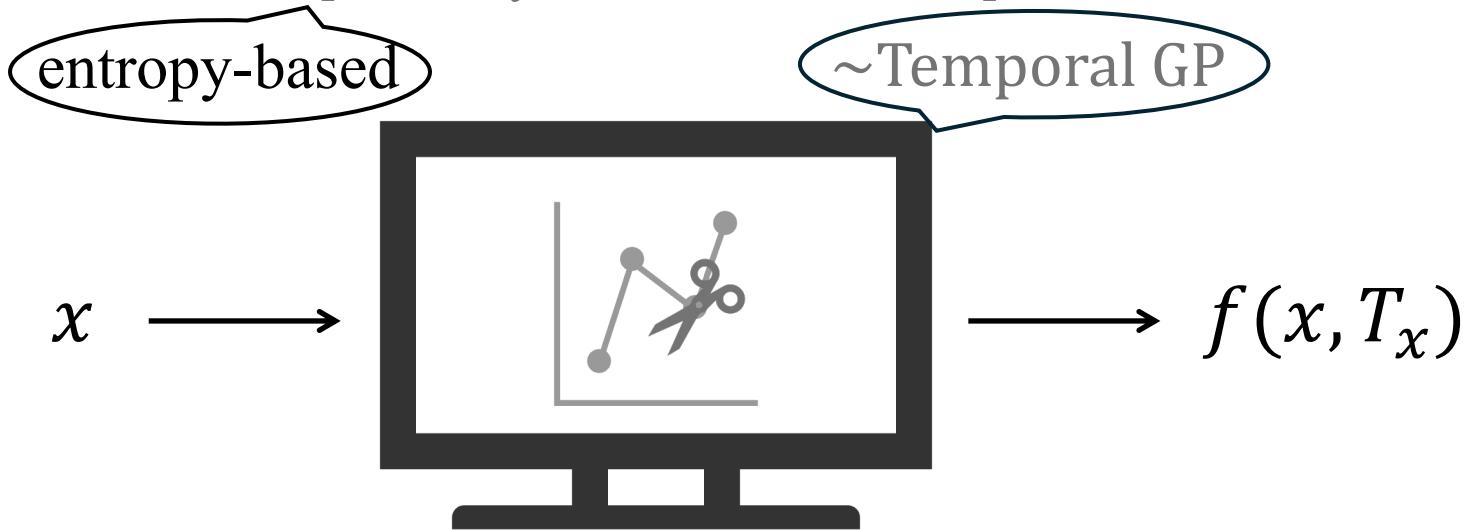
New design principle: Gittins Index

inherently feedback-aware

New Design Principle: Gittins Index

Studied Problem
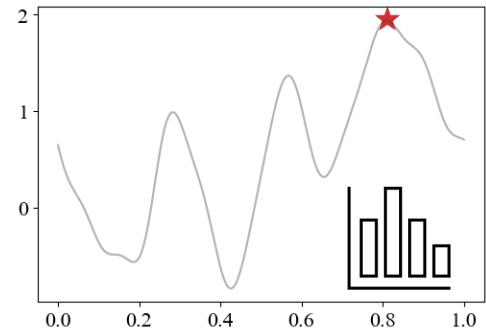
Cost-aware Bayesian optimization for black-box function with varying costs

Key idea

vs.

Link to Pandora's box and Gittins index theory

Impact

Competitive empirical performance

Ongoing work
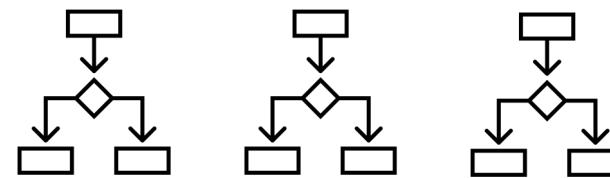
Freeze-thaw Bayesian optimization for black-box process with partial feedback

Find our paper on ArXiv!

"Cost-aware Bayesian Optimization via the Pandora's Box Gittins Index."