The Gittins-Index Design Principle for Cost-aware Bayesian Decision-Making under Uncertainty

Qian Xie 谢倩

B-exam @ Cornell ORIE

Dissertation Work Collaborators

Ziv Scully
Cornell, advisor

Alexander Terenin
Cornell

Peter Frazier
Cornell

Linda Cai
UC Berkeley

Nairen Cao
NYU

Yueli He
NYU

Theo Brown
UCL

Raul Astudillo
MBZUAI

Motivation: Design Selection under Uncertainty

ML model training:

Training hyperparameters
(e.g., learning rate, # layers) →

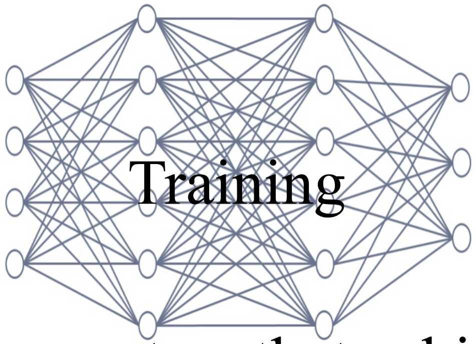

→ Accuracy

Goal: find hyperparameters that achieve high accuracy

Motivation: Design Selection under Uncertainty

ML model training:

Training hyperparameters (e.g., learning rate, # layers) → Training → Accuracy

Goal: find hyperparameters that achieve high accuracy

Adaptive experimentation:

Decision/design variables (e.g., UI design, pricing) → Deployment → Revenue

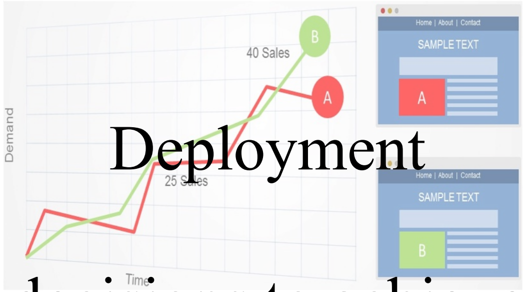

Goal: make decisions to achieve high revenue

Motivation: Design Selection under Uncertainty

Design choices x ⟶  ⟶ Performance metric $f(x)$

non-analytical & no gradient info

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

ML model training:

Training hyperparameters (e.g., learning rate, # layers) ⟶ 

⟶ Accuracy

Adaptive experimentation:

Decision/design variables (e.g., UI design, pricing) ⟶ 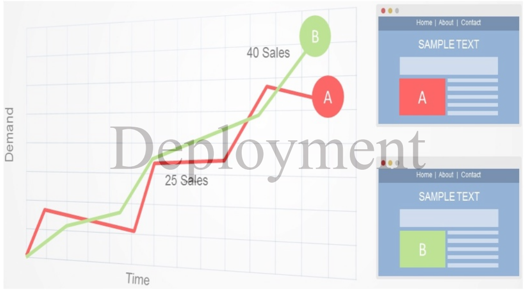

⟶ Revenue

Motivation: Design Selection under Uncertainty

Black-box optimization:

(gradient-based methods not applicable)

Input x ⟶

non-analytical & no gradient info

⟶ Observed outcome $f(x)$

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

ML model training:

Training hyperparameters (e.g., learning rate, # layers) ⟶

⟶ Accuracy

Adaptive experimentation:

Decision/design variables (e.g., UI design, pricing) ⟶

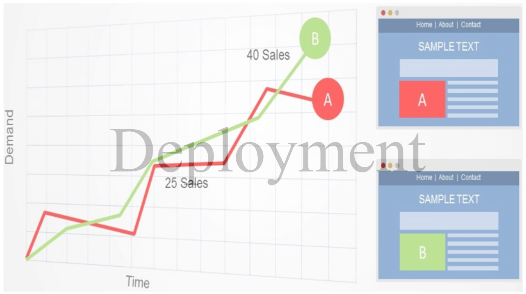

⟶ Revenue

Motivation: Costly Evaluation for Design Selection

Black-box optimization:

expensive-to-evaluate

Input x ⟶  ⟶ Observed outcome $f(x)$

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

ML model training:

Training hyperparameters (e.g., learning rate, # layers) ⟶ 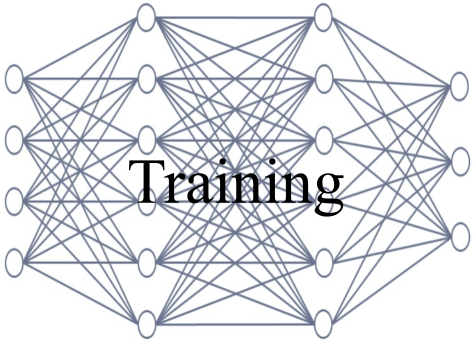

⟶ Accuracy

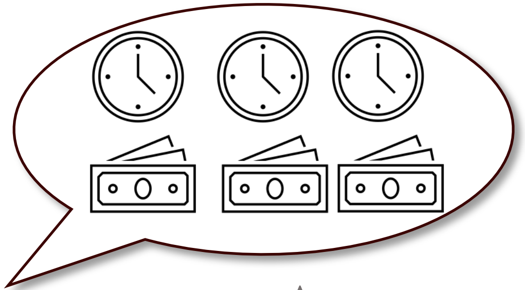

Training time

Compute credits

Adaptive experimentation:

Decision/design variables (e.g., UI design, pricing) ⟶ 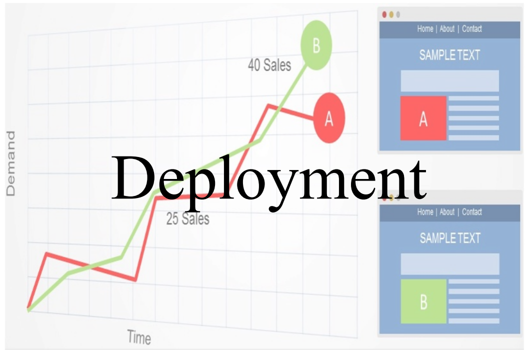

⟶ Revenue

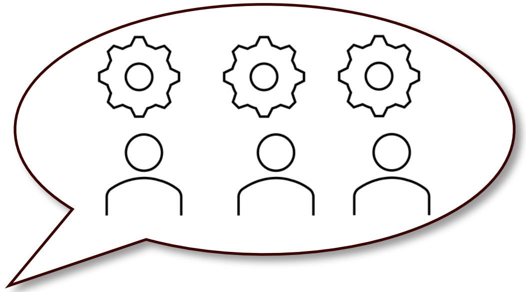

Operational cost

User experience

Motivation: Costly Evaluation for Design Selection

Black-box optimization:

Input x ⟶  ⟶ Observed outcome $f(x)$

expensive-to-evaluate

limited #evaluations

Goal: $\mathbb{E}\max_{i=1,2,\dots,N} f(x_i)$

ML model training:

Training hyperparameters (e.g., learning rate, # layers) ⟶ 

⟶ Accuracy

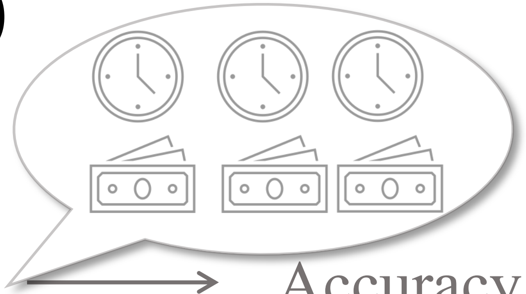

Training time

Compute credits

Adaptive experimentation:

Decision/design variables (e.g., UI design, pricing) ⟶ 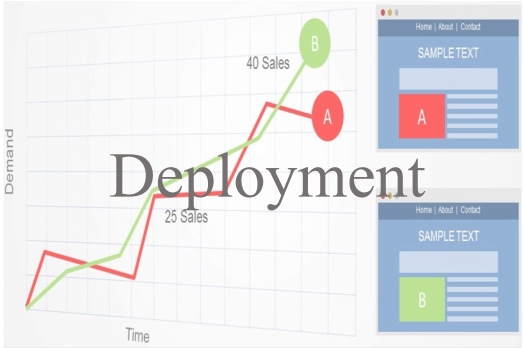

⟶ Revenue

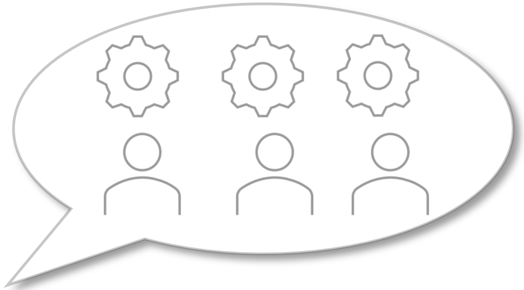

Operational cost

User experience

Naïve Approach: Grid Search

ML model training:

Training hyperparameters → 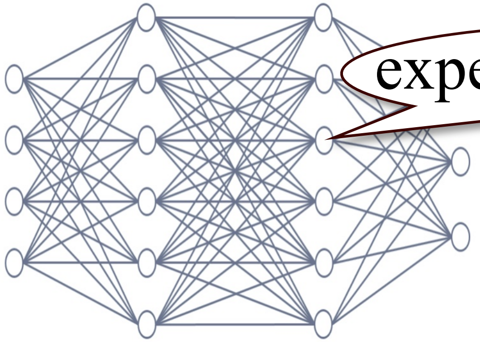 → Accuracy

expensive-to-evaluate

Goal: find hyperparameters that achieve high accuracy

Training hyperparameter	Range	Number of Options
Batch size	[16, 512]	10
Learning rate	[1e-4, 1e-1]	10
Momentum	[0.1, 0.99]	10
Weight decay	[1e-5, 1e-1]	10
Number of layers	{1, 2, 3, 4}	4
Max units per layer	[64, 1024]	10
Dropout	[0.0, 1.0]	10

40,000,000 combinations!

Sample-Efficient Decision-Making

Design choices
(e.g., learning rate, price)

→

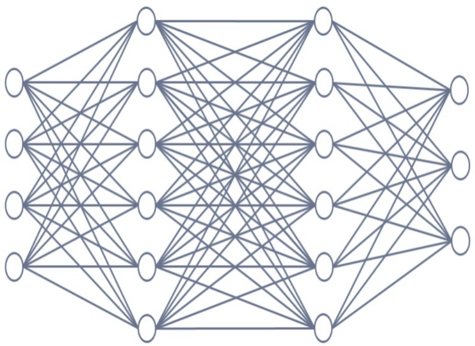

Training/Experimentation

→

Performance metric
(e.g., accuracy, revenue)

Goal: $\mathbb{E} \max_{i=1,2,\ldots,N} f(x_i)$

As few as possible

Sample-Efficient Adaptive Decision-Making

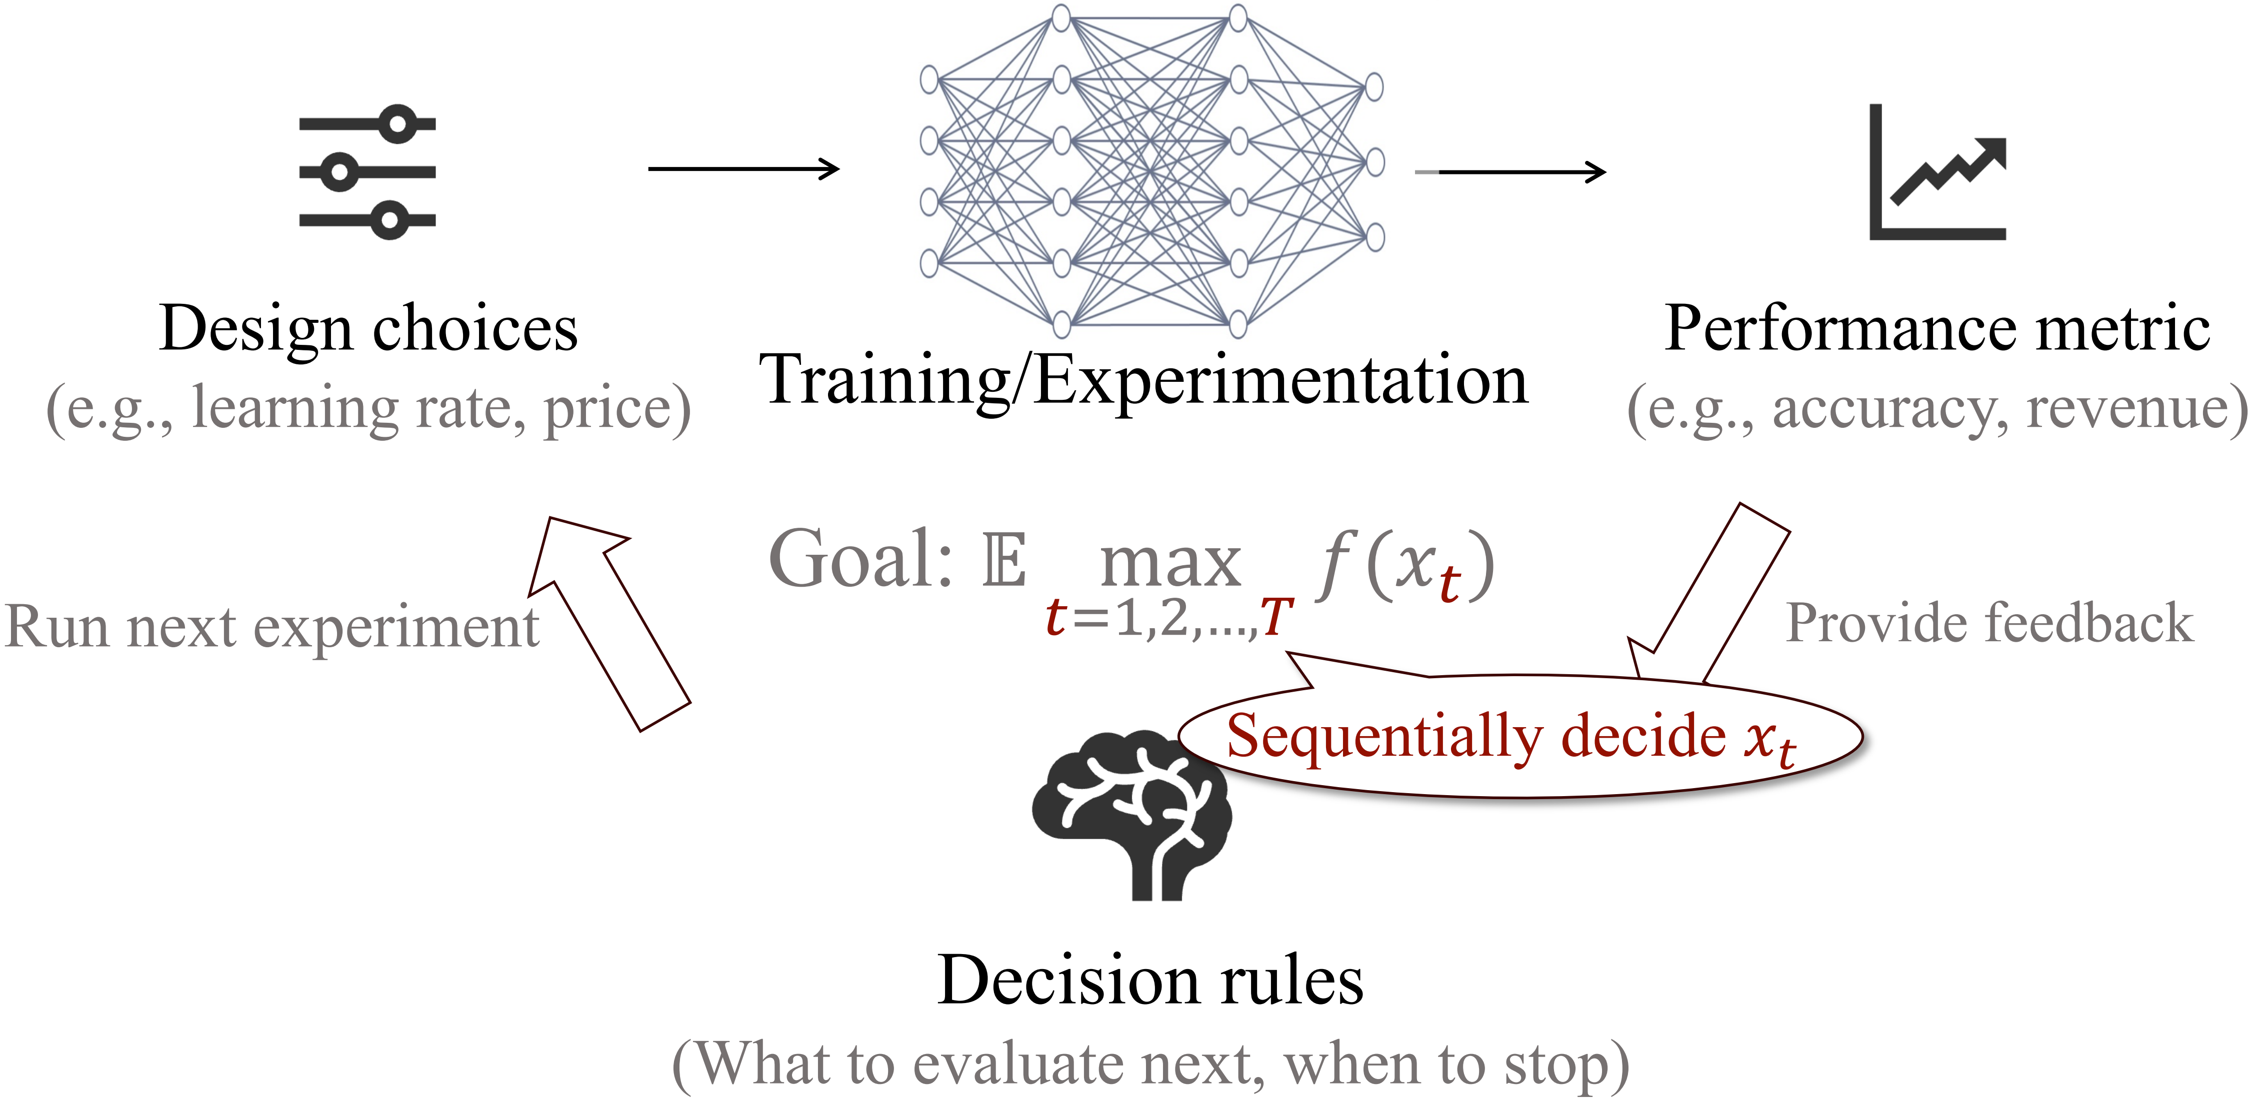

Sample-Efficient Adaptive Decision-Making

Manual tuning:

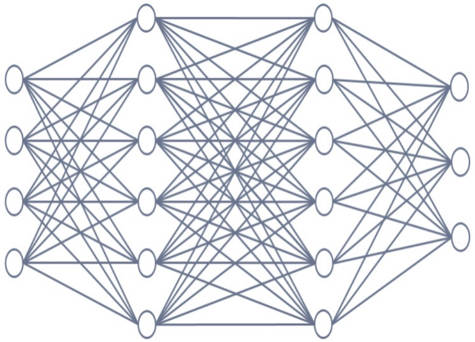

Design choices
(e.g., learning rate, price)

→ **Training/Experimentation** →

Performance metric
(e.g., accuracy, revenue)

Run next experiment

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

Provide feedback

Sequentially decide x_t

Experience-based human decision rules
(What to evaluate next, when to stop)

Sample-Efficient Adaptive Decision-Making

Automated machine learning:

Design choices
(e.g., learning rate, price)

→

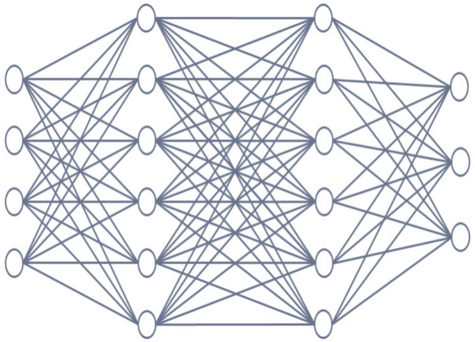

Training/Experimentation

→

Performance metric
(e.g., accuracy, revenue)

Run next experiment

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

Provide feedback

Sequentially decide x_t

Smart automated decision rules
(What to evaluate next, when to stop)

Sample-Efficient Bayesian Decision-Making

Automated machine learning:

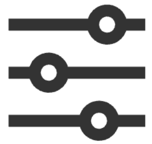

Design choices
(e.g., learning rate, price)

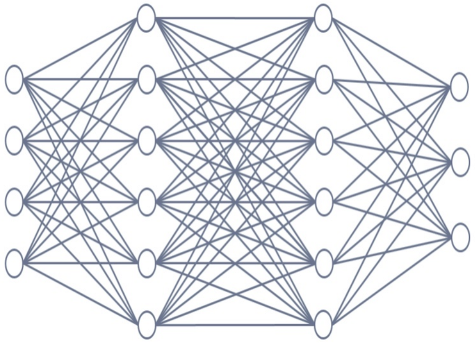

Training/Experimentation

Performance metric
(e.g., accuracy, revenue)

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

Run next experiment

Update belief
(Bayes' rule)

Sequentially decide x_t

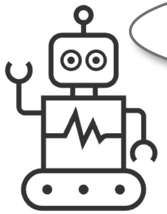

Smart automated decision rules
(What to evaluate next, when to stop)

Under-explored Factors for Better Decisions

Varying evaluation costs

Smart stopping time

Observable multi-stage feedback

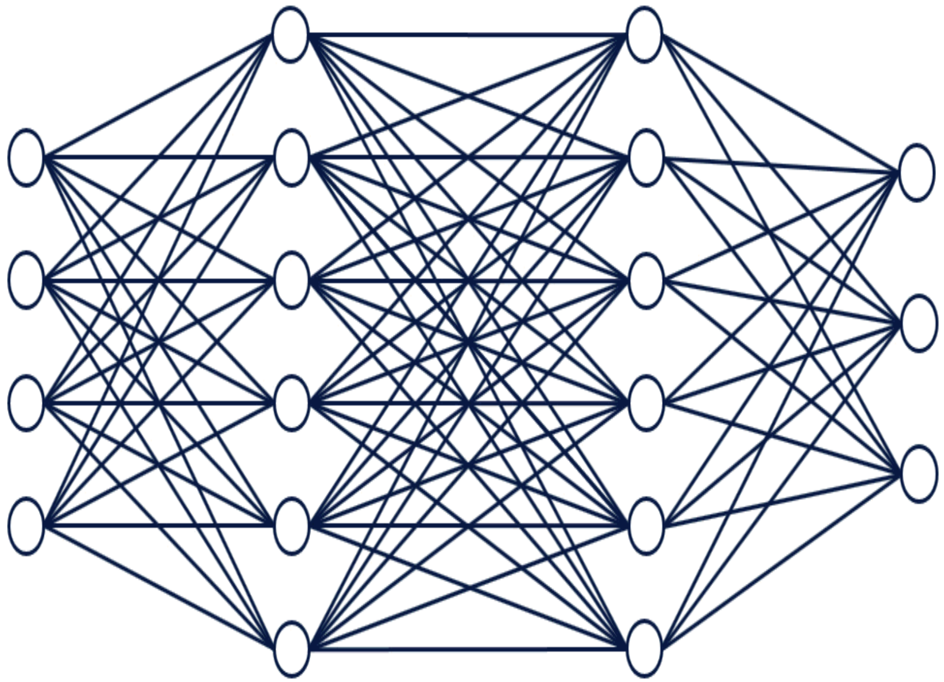

Under-explored Factors for Better Decisions

- Varying evaluation costs
- Smart stopping time
- Observable multi-stage feedback

New design principle:
Gittins index

Why Gittins index?

Varying evaluation costs

Smart stopping time

Observable multi-stage feedback

Featured in related sequential decision problems

Why Gittins index?

Varying evaluation costs

Featured in Pandora's box

Part I

Smart stopping time

Featured in Pandora's box

Observable multi-stage feedback

New design principle:
Gittins index

Featured in related sequential decision problems

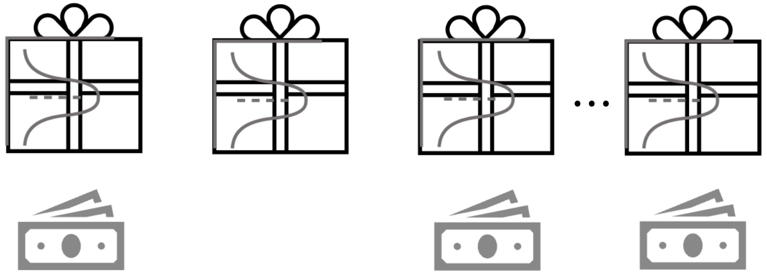

"Cost-aware Bayesian Optimization via the Pandora's Box Gittins Index." NeurIPS'24.

"Cost-aware Stopping for Bayesian Optimization." ICML'26.

Why Gittins index?

Varying evaluation costs

Featured in Pandora's box

Part I

Smart stopping time

Featured in Pandora's box

Observable multi-stage feedback

New design principle: Gittins index

Optimal in related sequential decision problems

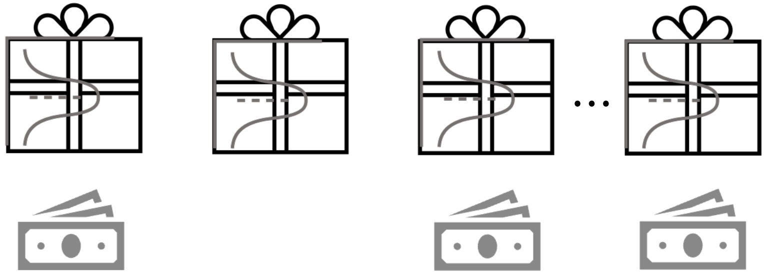

"Cost-aware Bayesian Optimization via the Pandora's Box Gittins Index." NeurIPS'24.

"Cost-aware Stopping for Bayesian Optimization." ICML'26.

Why Gittins index?

Varying evaluation costs

Features in Pandora's box

Part I

Smart stopping time

Features in Pandora's box

Observable multi-stage feedback

Featured in Markov chain selection

Part II

New design principle: Gittins index

Optimal in related sequential decision problems

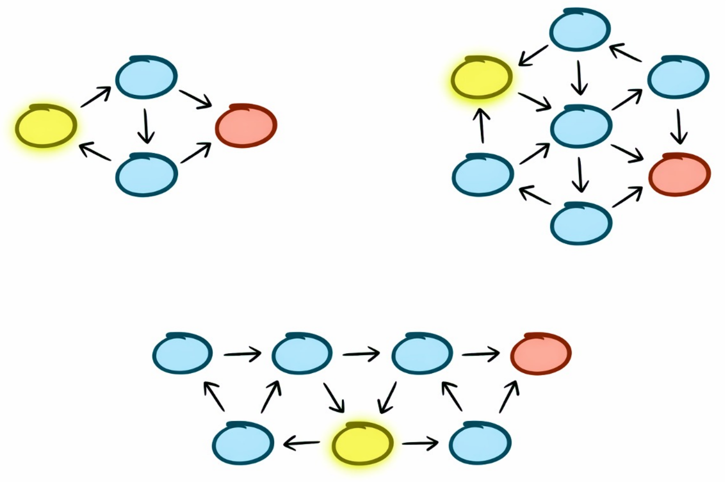

"Efficient Cost-Aware LLM Evaluation via Bayesian Bandit Gittins Indices." ICML'26 Workshop DEMO.

Talk Outline

Part I:

Cost-aware Bayesian Optimization with Adaptive Stopping via Pandora's Box Gittins Indices

Part II:

Towards Bayesian Optimization with Multi-Stage Feedback via Markov Chain Gittins Indices

Part I:

Cost-aware Bayesian Optimization with Adaptive Stopping via Pandora's Box Gittins Indices

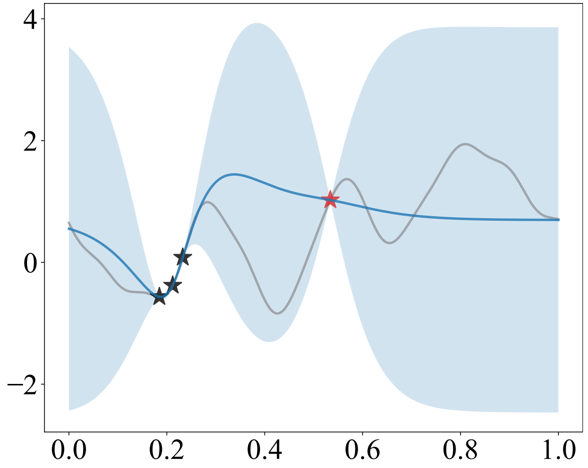

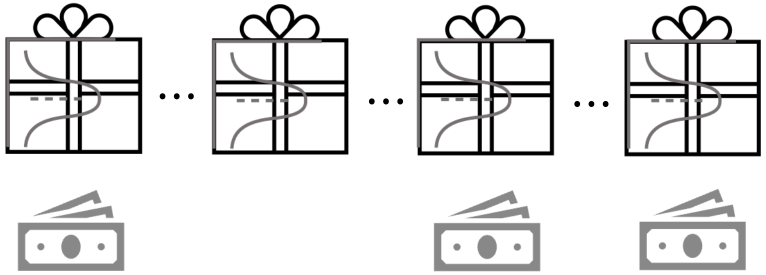

Bayesian Optimization

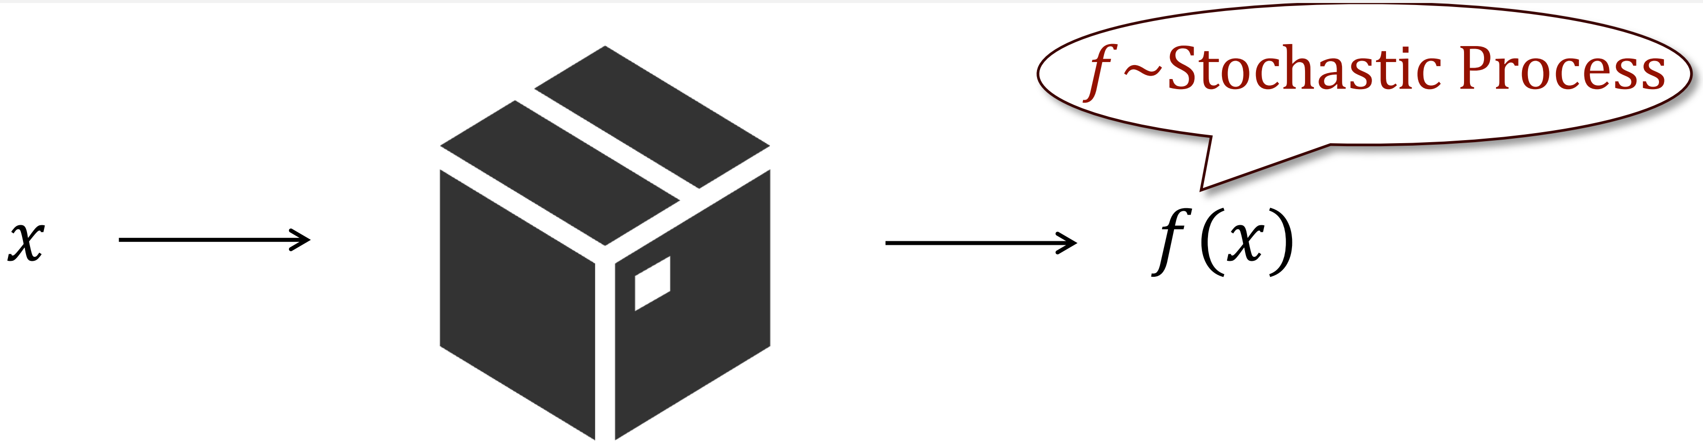

Black-box function

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

Bayesian Optimization

Time 0

$x \longrightarrow$

$\longrightarrow f(x)$

$f \sim$Stochastic Process

Model belief

Black-box function

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

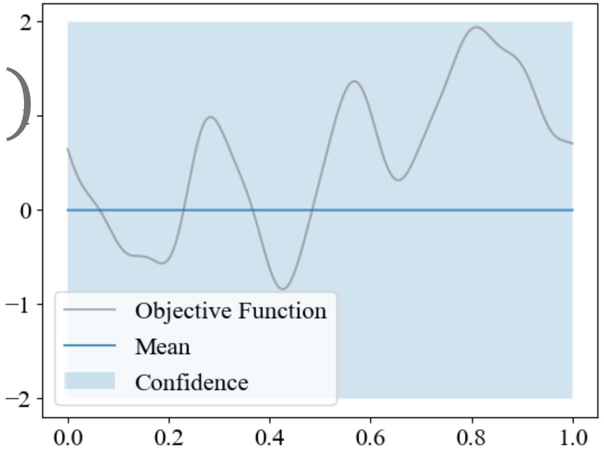

Probabilistic model

(e.g., Gaussian process)

Bayesian Optimization

Time t

$x_1, \ldots, x_t$ →  → $f(x_1), \ldots, f(x_t)$

$f \sim$ Stochastic Process

Black-box function

Model belief

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

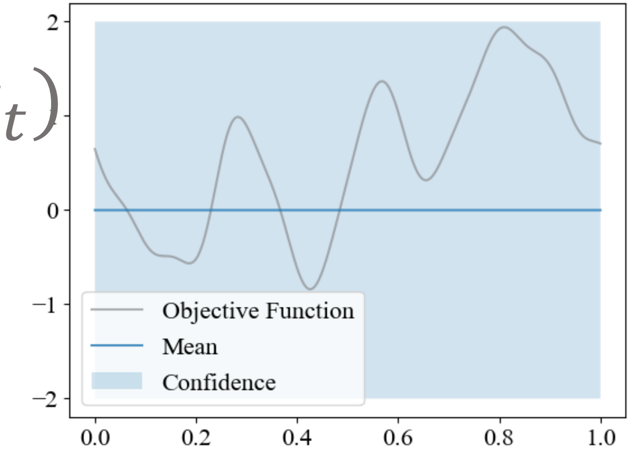

Probabilistic model
(e.g., Gaussian process)

Bayesian Optimization

Time t

$f \sim$Stochastic Process

$x_1, \dots, x_t \longrightarrow$

$\longrightarrow f(x_1), \dots, f(x_t)$

Black-box function

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

Update belief
(Bayes' rule)

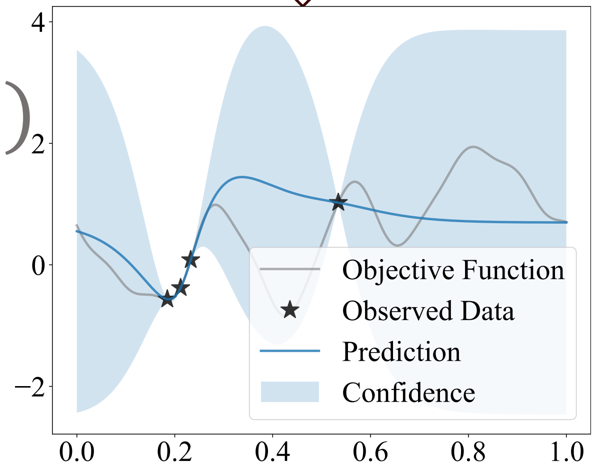

Probabilistic model
(e.g., Gaussian process)

Bayesian Optimization

Time t

$x_1, \dots, x_t$ →

Black-box function

→ $f(x_1), \dots, f(x_t)$

Update belief

(Bayes' rule)

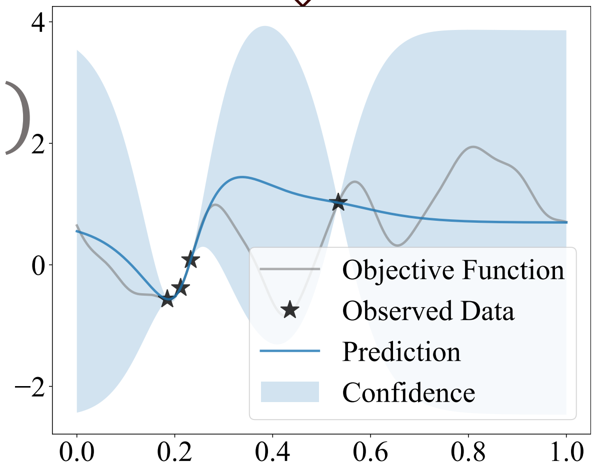

Probabilistic model

(e.g., Gaussian process)

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

scoring (worth of each point

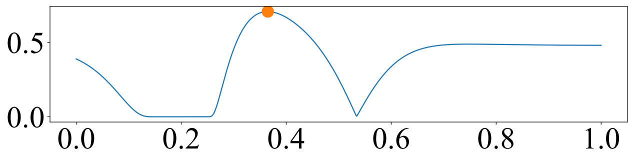

Decision rule

(What to evaluate next)

Bayesian Optimization

Time t

$x_1, \dots, x_t \longrightarrow$

$\longrightarrow f(x_1), \dots, f(x_t)$

Black-box function

Update belief
(Bayes' rule)

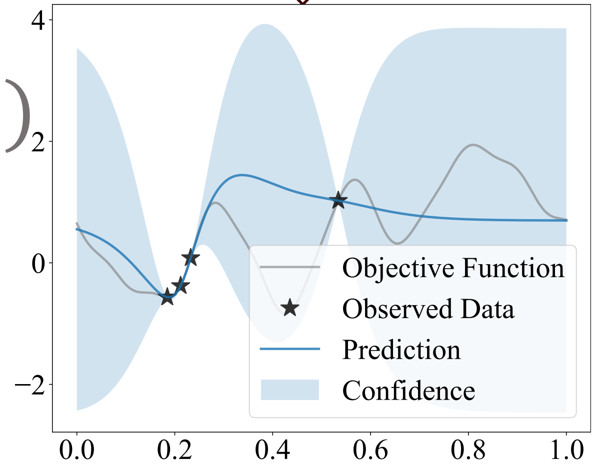

Probabilistic model
(e.g., Gaussian process)

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

scoring (worth of each point

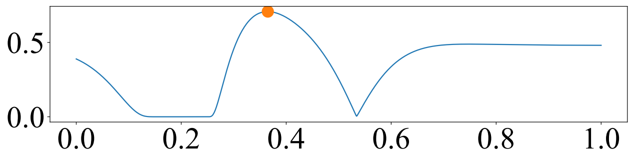

Decision rule
(What to evaluate next)

Decide next input
(gradient-based optimization)

Bayesian Optimization

Time $t + 1$

$x_1, \ldots, x_{t+1}$ →  → $f(x_1), \ldots, f(x_{t+1})$

Black-box function

Decide next input
(gradient-based optimization)

Update belief
(Bayes' rule)

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

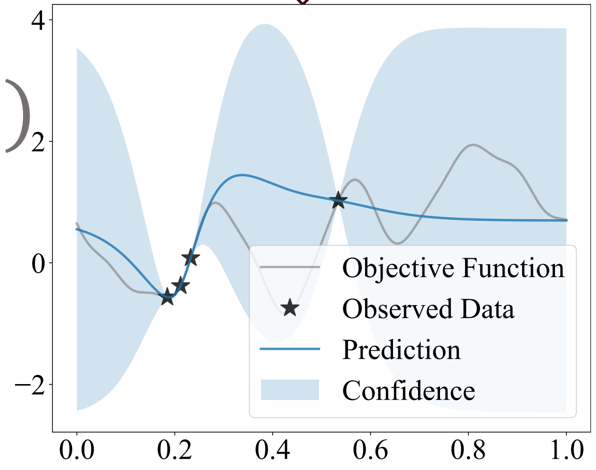

Probabilistic model
(e.g., Gaussian process)

scoring (worth of each point)

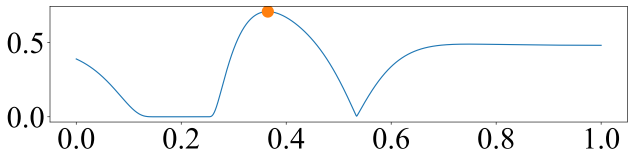

Decision rule
(What to evaluate next)

Bayesian Optimization

Time $t + 1$

$x_1, \dots, x_{t+1}$ →  → $f(x_1), \dots, f(x_{t+1})$

Black-box function

Update belief
(Bayes' rule)

Probabilistic model
(e.g., Gaussian process)

Decide next input
(gradient-based
optimization)

My focus

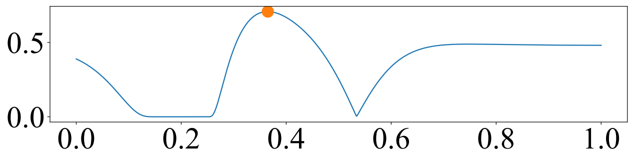

Decision rule
(What to evaluate next)

scoring (worth
of each point

Challenges in Decision Rule Design

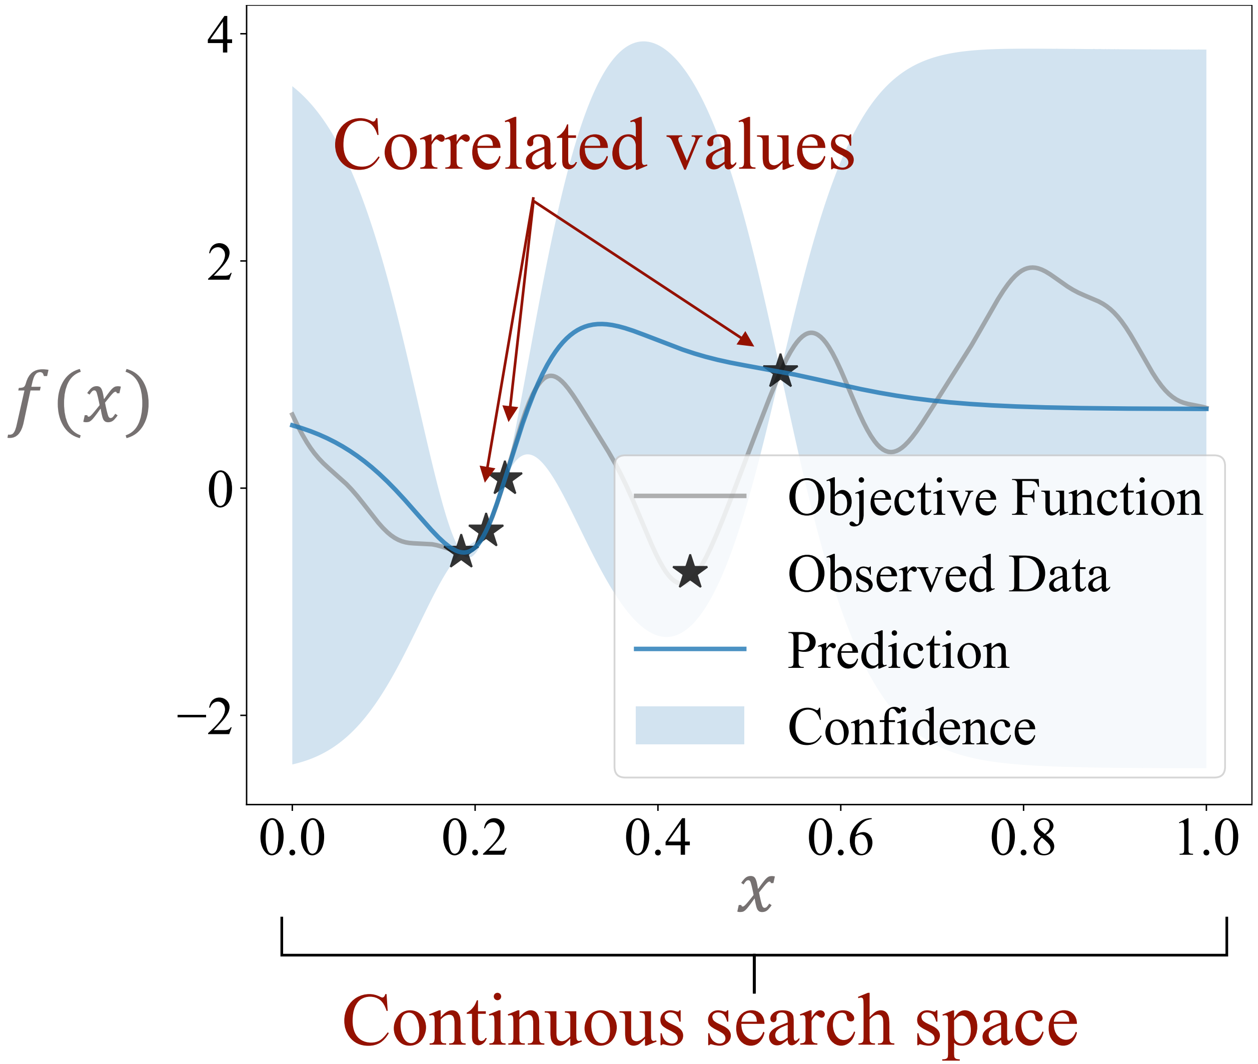

Correlation & continuity $\Longrightarrow$ Intractable MDP $\Longrightarrow$ Optimal policy unknown

Popular Decision Rule: Expected Improvement

current best observed

data D

$$\mathrm{EI}(x; y_{\text{best}}) := \mathbb{E}[\underbrace{\max(f(x) - y_{\text{best}}, 0)}_{\text{“improvement”}} \mid x_1, \ldots, x_t]$$

$$x_{t+1} = \operatorname{argmax}_x \mathrm{EI}_{f|D}(x; y_{\text{best}})$$

posterior distribution

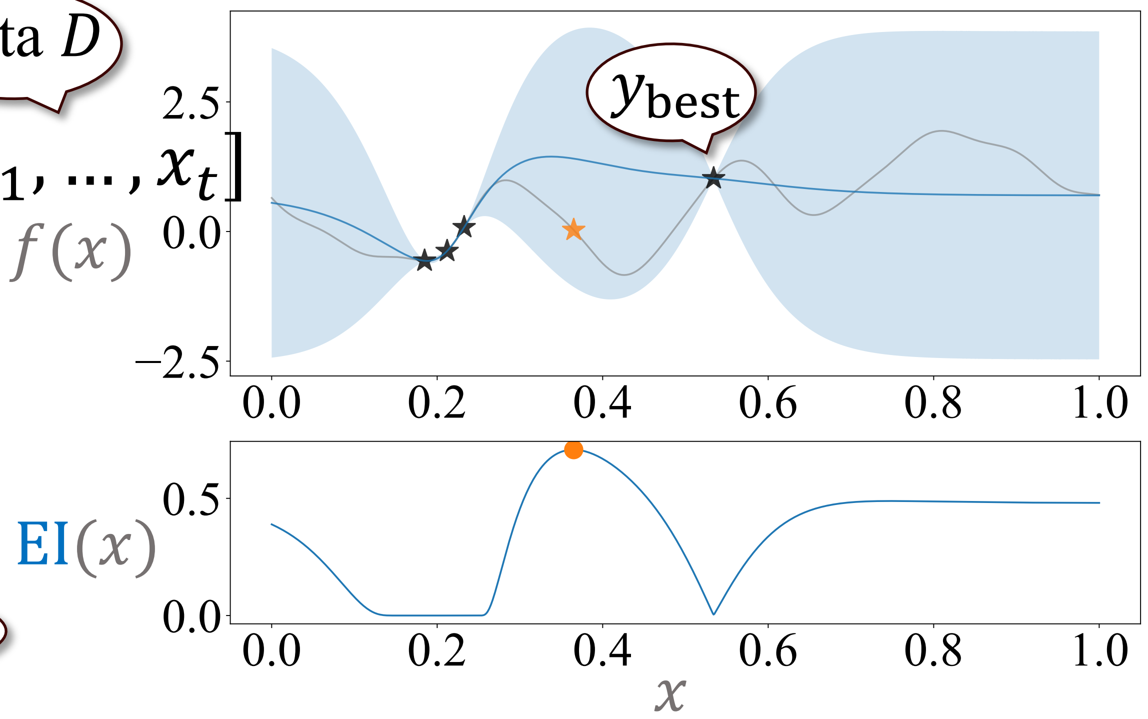

One-step approximation to MDP

Expected improvement

How does EI incorporate costs and stopping?

Varying evaluation costs

$$\mathrm{EIPC}_c(x; y_{\mathrm{best}}) = \mathrm{EI}(x; y_{\mathrm{best}})/c(x)$$

[Snoek et al.'12]

Arbitrarily bad

[Astudillo et al.'21]

Smart stopping time

$$\tau\colon\ \mathrm{EI}(x_\tau; y_{\mathrm{best}}) \le \theta$$

[Locatelli'97, Nguyen et al.'17, Ishibashi et al.'23]

Which threshold?

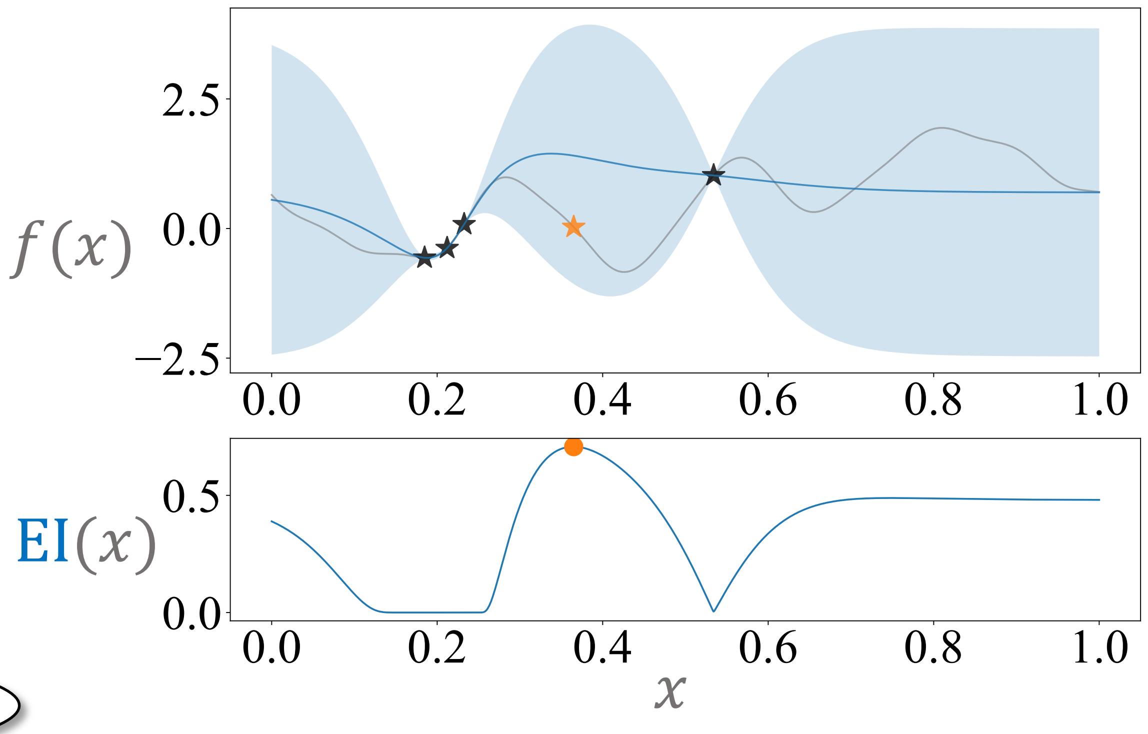

Expected improvement

Expected Improvement vs Gittins Index

Varying evaluation costs

$\mathrm{EIPC}_c(x; y_{\mathrm{best}}) = \mathrm{EI}(x; y_{\mathrm{best}})/c(x)$

Arbitrarily bad

$\mathrm{GI}_c(x)$

Costs built-in

Smart stopping time

$\tau\colon\ \mathrm{EI}(x_\tau; y_{\mathrm{best}}) \le \theta$

Which threshold?

$\tau\colon\ \mathrm{GI}_c(x_\tau) \le y_{\mathrm{best}}$

Built-in threshold

Expected Improvement vs Gittins Index

Varying evaluation costs

$\mathrm{EIPC}_c(x; y_{\mathrm{best}}) = \mathrm{EI}(x; y_{\mathrm{best}})/c(x)$

Arbitrarily bad

$\mathrm{GI}_c(x)$

Costs built-in

Smart stopping time

$\tau\colon\ \mathrm{EI}(x_\tau; y_{\mathrm{best}}) \leq \theta$

Which threshold?

$\tau\colon\ \mathrm{GI}_c(x_\tau) \leq y_{\mathrm{best}}$

Built-in threshold

Where do Gittins indices come from?

Bayesian Optimization

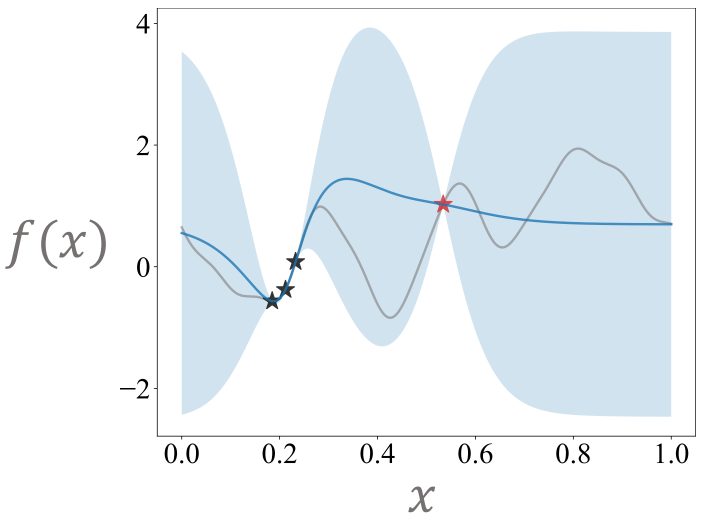

Continuous search space

Correlated function values

Bayesian Optimization

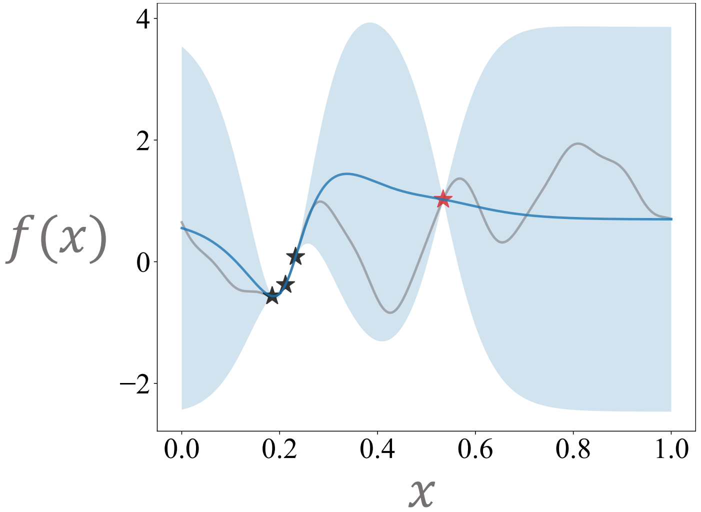

Continuous search space ⟹ Discrete

Correlated function values ⟹ Independent

Bayesian Optimization

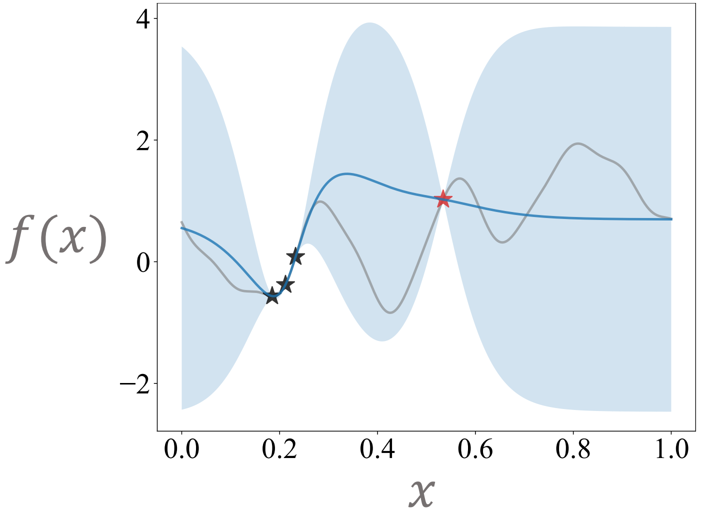

Continuous search space

Correlated function values

Pandora's Box

[Weitzman'79]

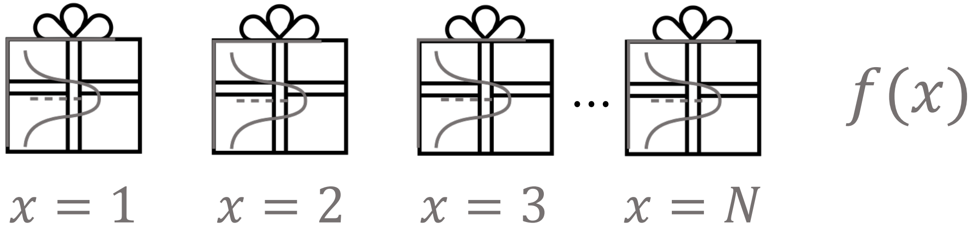

Discrete

Independent

Bayesian Optimization

Continuous search space

Correlated function values

Pandora's Box

[Weitzman'79]

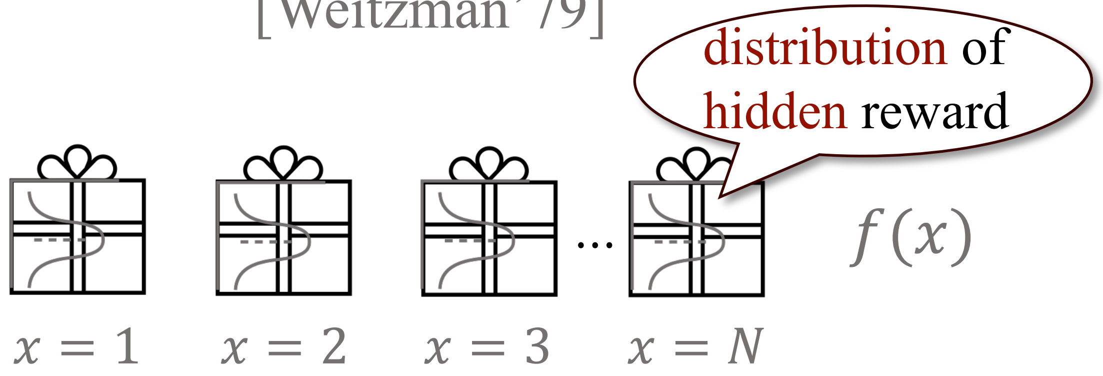

Discrete

Independent

Bayesian Optimization

Continuous search space

Correlated function values

Pandora's Box

[Weitzman'79]

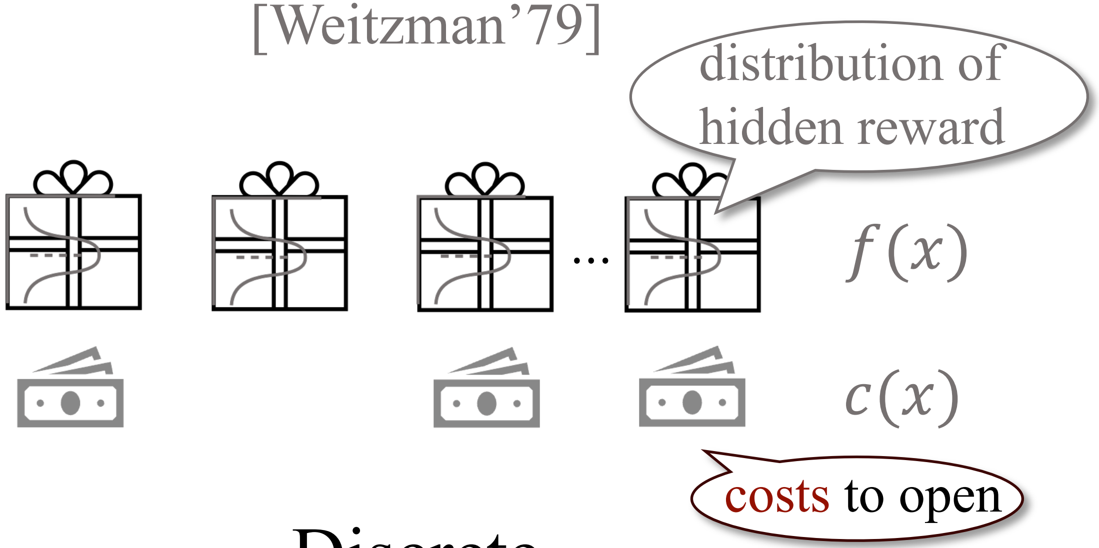

Discrete

Independent

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

Bayesian Optimization

Continuous search space

Correlated function values

Pandora's Box

[Weitzman'79]

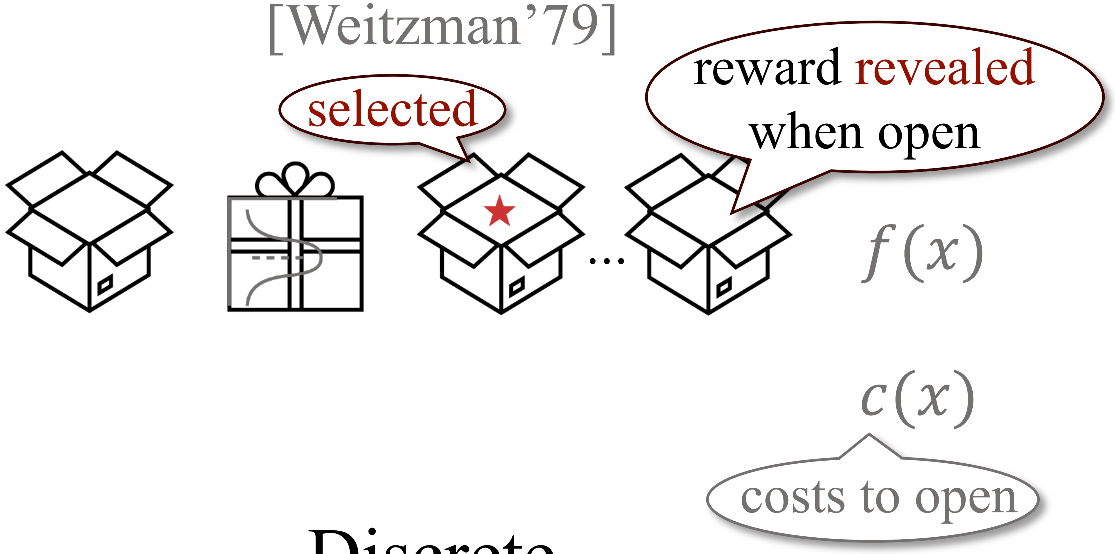

Discrete

Independent

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

Bayesian Optimization

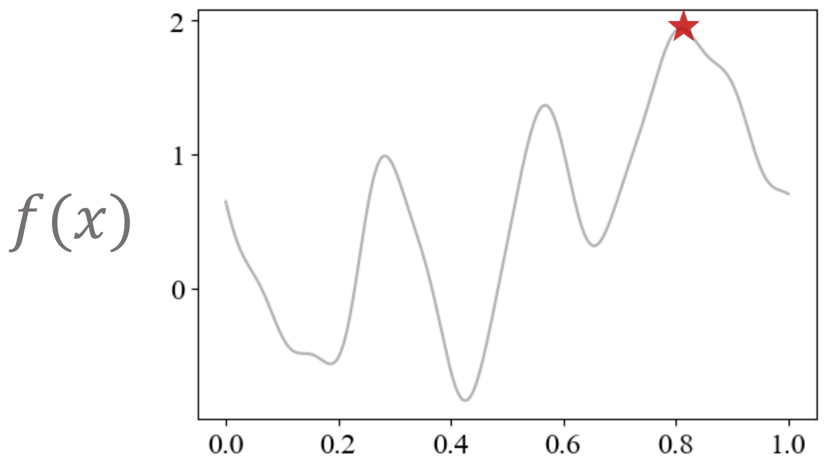

Continuous

Correlated

Cost-unaware

Fixed-iteration

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

Pandora's Box

[Weitzman'79]

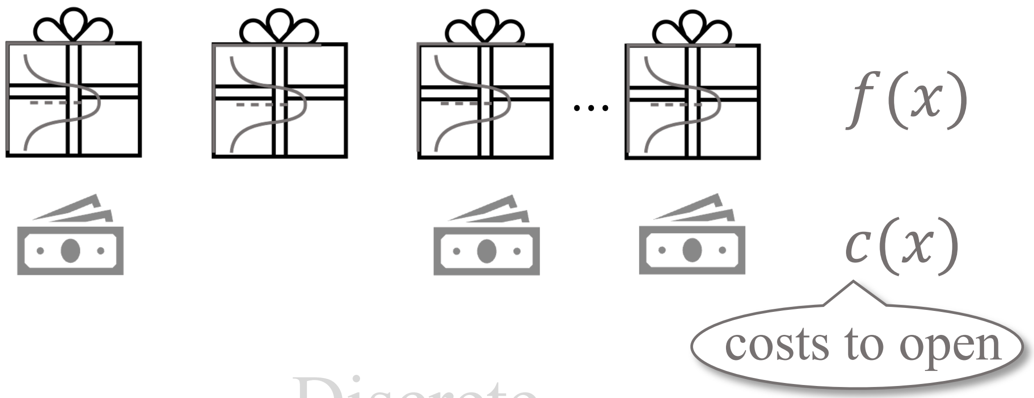

Discrete

Independent

Cost-aware

Flexible-stopping

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

Cost-aware Bayesian Optimization | Pandora's Box

[Weitzman'79]

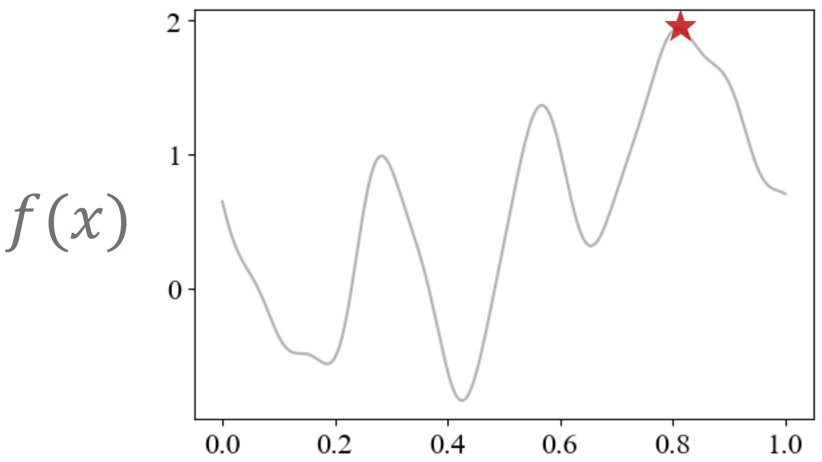

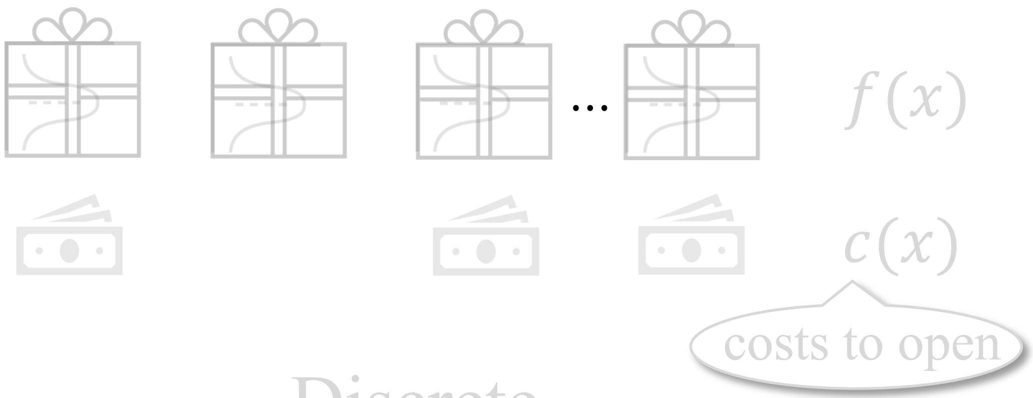

Cost-aware Bayesian Optimization	Pandora's Box
Continuous	Discrete
Correlated	Independent
Cost-aware	Cost-aware
Flexible-stopping	Flexible-stopping
$\mathbb{E} \max_{t=1,2,\ldots,T} f(x_t) - \mathbb{E} \sum_{t=1}^{T} c(x_t)$	$\mathbb{E} \max_{t=1,2,\ldots,T} f(x_t) - \mathbb{E} \sum_{t=1}^{T} c(x_t)$

Cost-aware Bayesian Optimization

Continuous

Correlated

Pandora's Box

[Weitzman'79]

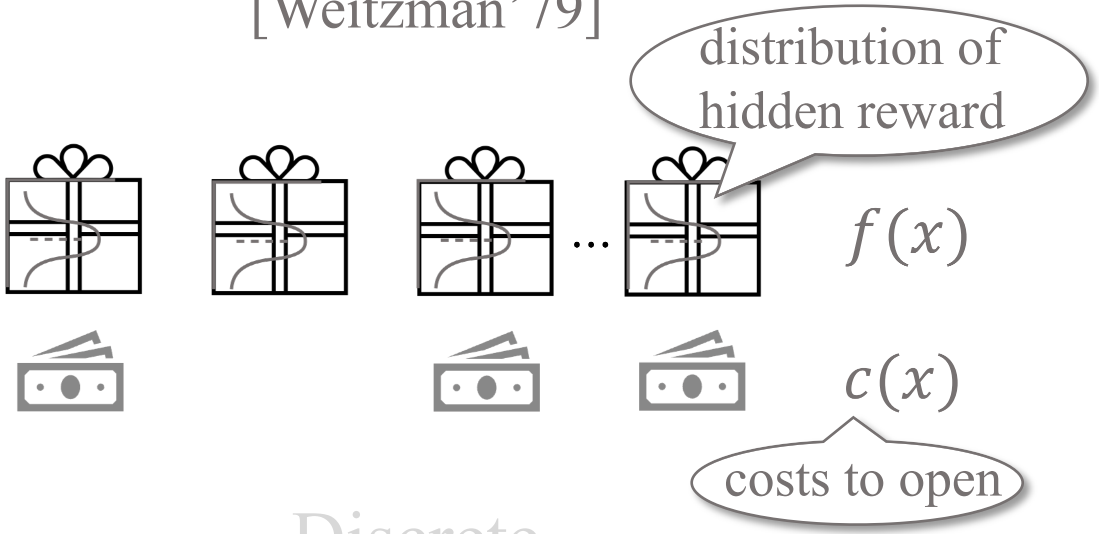

Discrete

Independent

Optimal policy: Gittins index

Cost-aware Bayesian Optimization

Continuous

Correlated

Pandora's Box

[Weitzman'79]

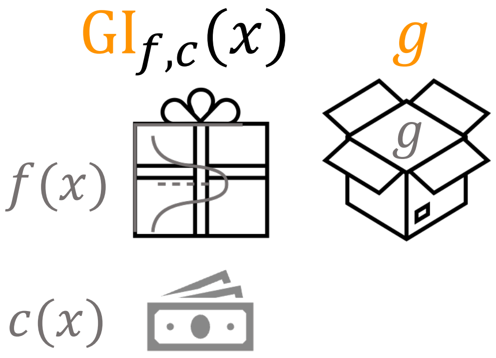

Step 1: Assign each box a Gittins index

Discrete

Independent

Optimal policy: Gittins index

Cost-aware Bayesian Optimization

Continuous

Correlated

Pandora's Box

[Weitzman'79]

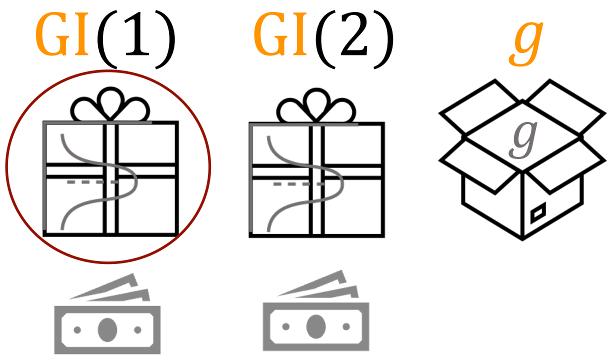

Step 2: Act on the box with the highest index

- *Closed*: open it
- *Opened*: select & stop

Discrete

Independent

Optimal policy: Gittins index

Cost-aware Bayesian Optimization

Continuous

Correlated

Pandora's Box

[Weitzman'79]

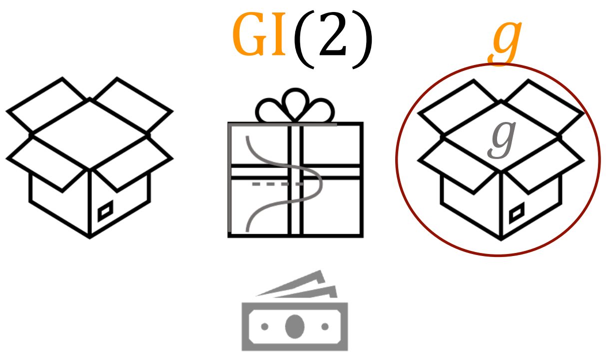

Step 2': Act on the box with the highest index

- *Closed*: open it
- *Opened*: select & stop

Discrete

Independent

Optimal policy: Gittins index

Cost-aware Bayesian Optimization | Pandora's Box

[Weitzman'79]

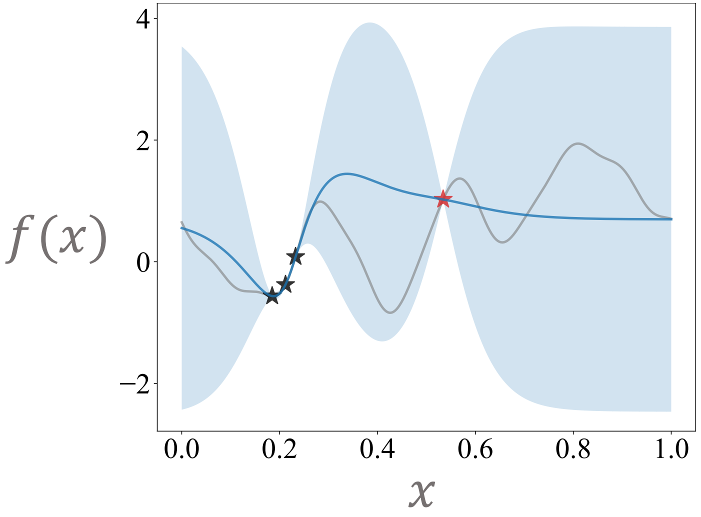

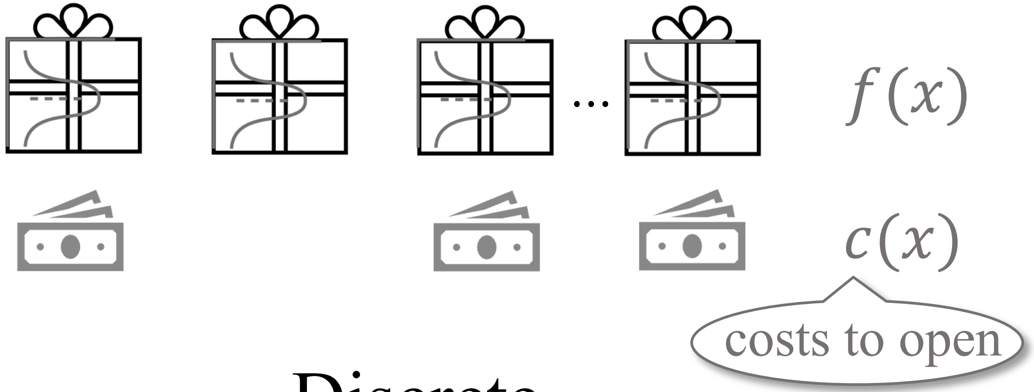

Continuous search space | Discrete

Correlated function values | Independent

How to translate?

⟸ Optimal policy: Gittins index

Cost-aware Bayesian Optimization

Pandora's Box

[Weitzman'79]

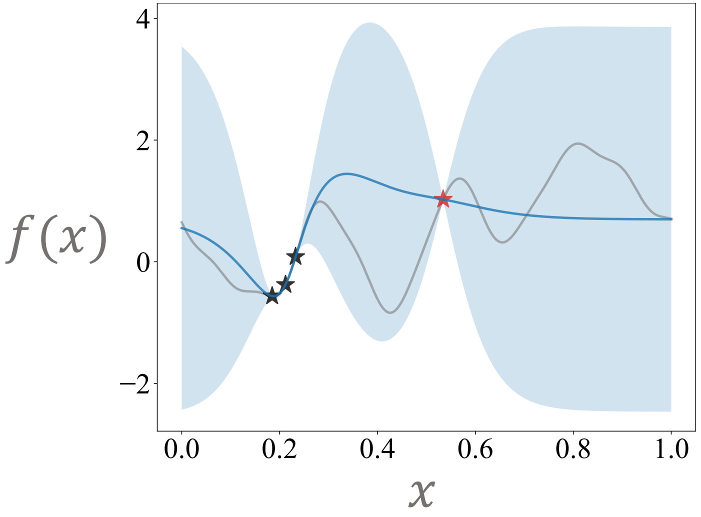

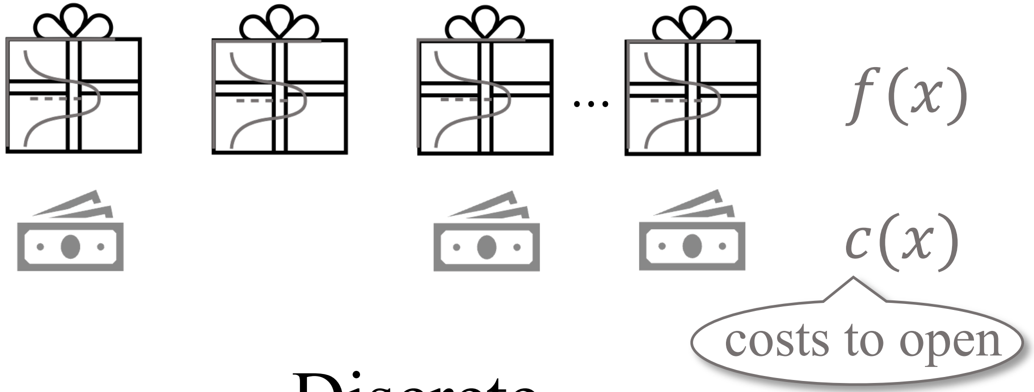

Continuous search space | Discrete

Correlated function values | Independent

Our policy: $\mathrm{GI}_{f|D,c}(x)$ ⟸ Optimal policy: $\mathrm{GI}_{f,c}(x)$

incorporate posterior

take continuum limit

Data

New!

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

Cost-aware Bayesian Optimization

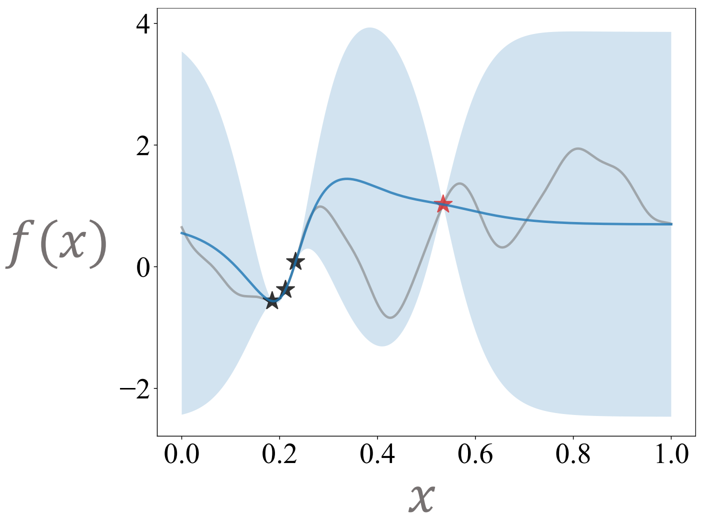

Continuous search space

Correlated function values

Our policy: $\mathrm{GI}_{f|D,c}(x)$

How to implement?

incorporate posterior

$\Longleftarrow$

take continuum limit

New!

Pandora's Box

[Weitzman'79]

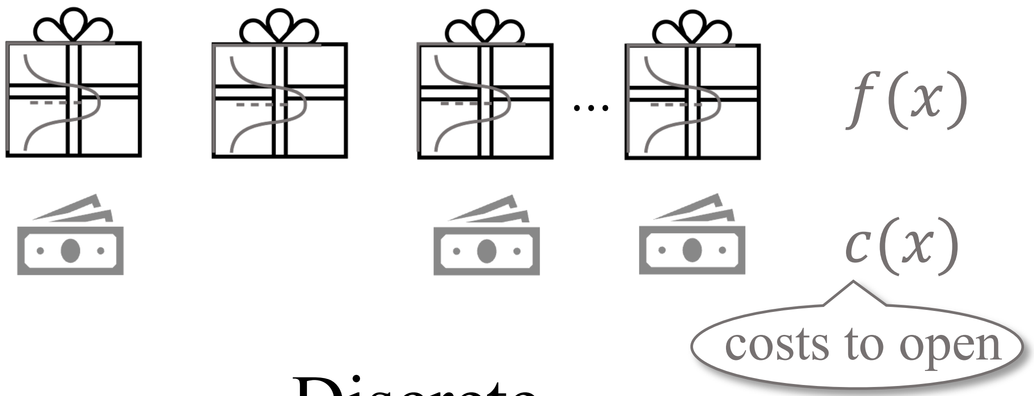

Discrete

Independent

Optimal policy: $\mathrm{GI}_{f,c}(x)$

How to implement Gittins Indices?

$\mathrm{GI}_{f|D,c}(x) := \text{solution } g \text{ s.t.}$

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

Data

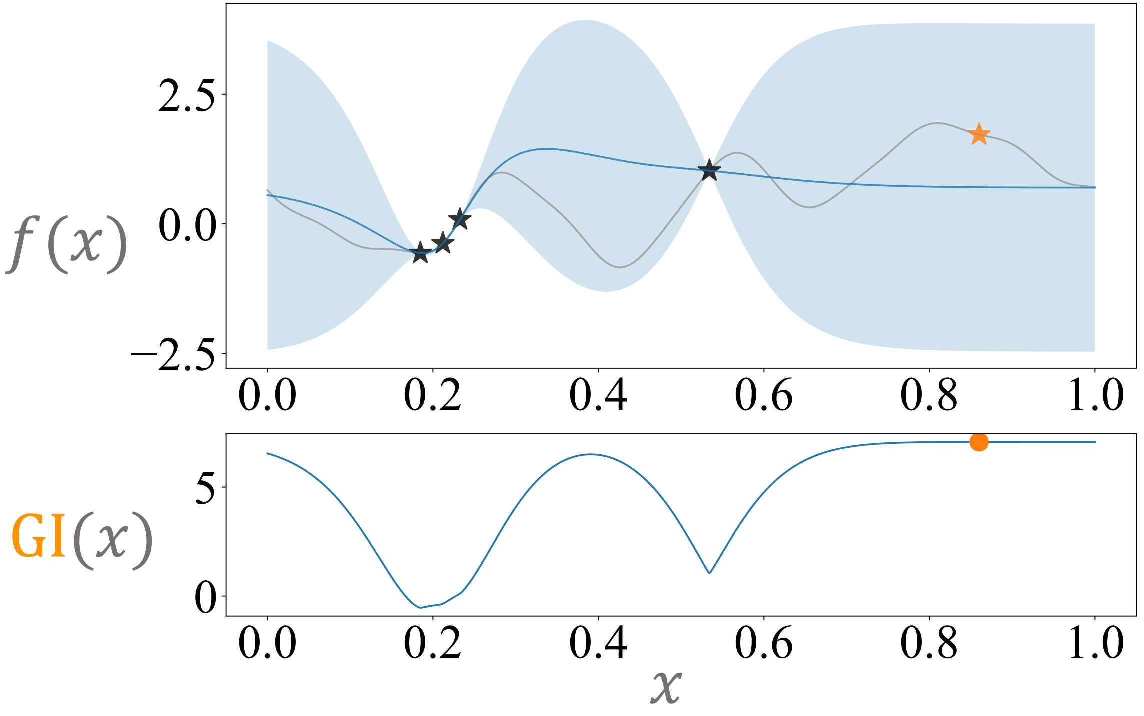

Gittins Indices

How to implement Gittins Indices?

$\mathrm{GI}_{f|D,c}(x) := \text{solution } g \text{ s.t.}$

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

Data

analytical expression
& monotonicity in g

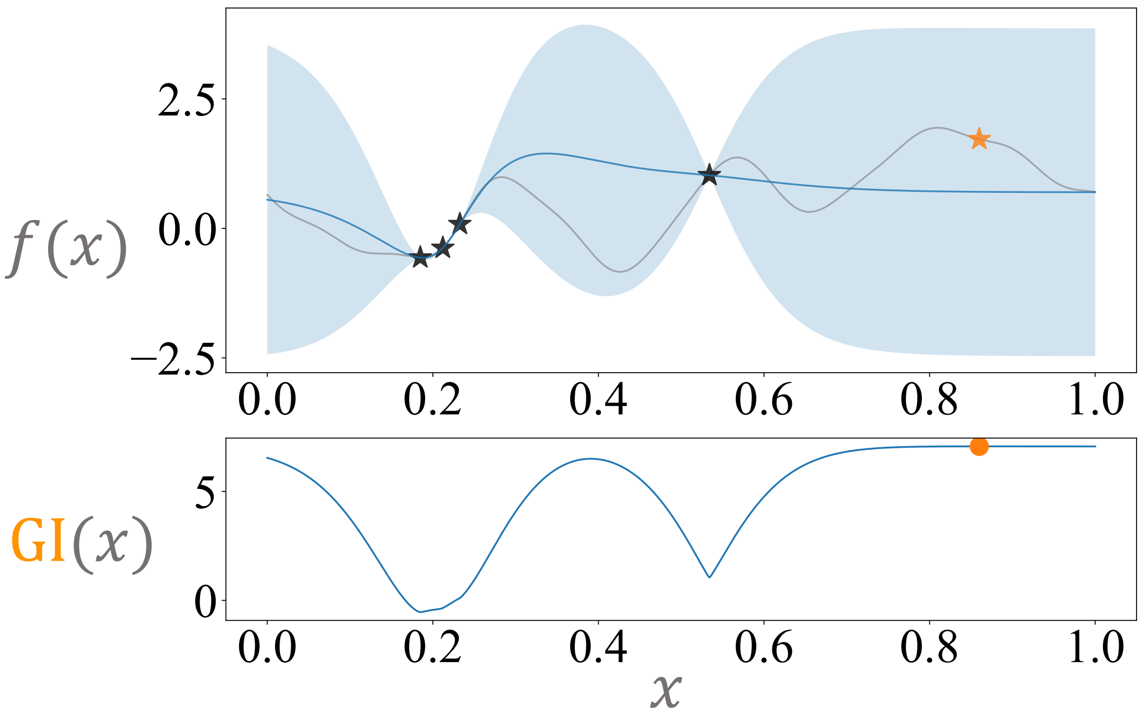

Gittins Indices

How to implement Gittins Indices?

bisection search

$\text{GI}_{f|D,c}(x) := \text{solution } g \text{ s.t.}$

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

Data

analytical expression
& monotonicity in g

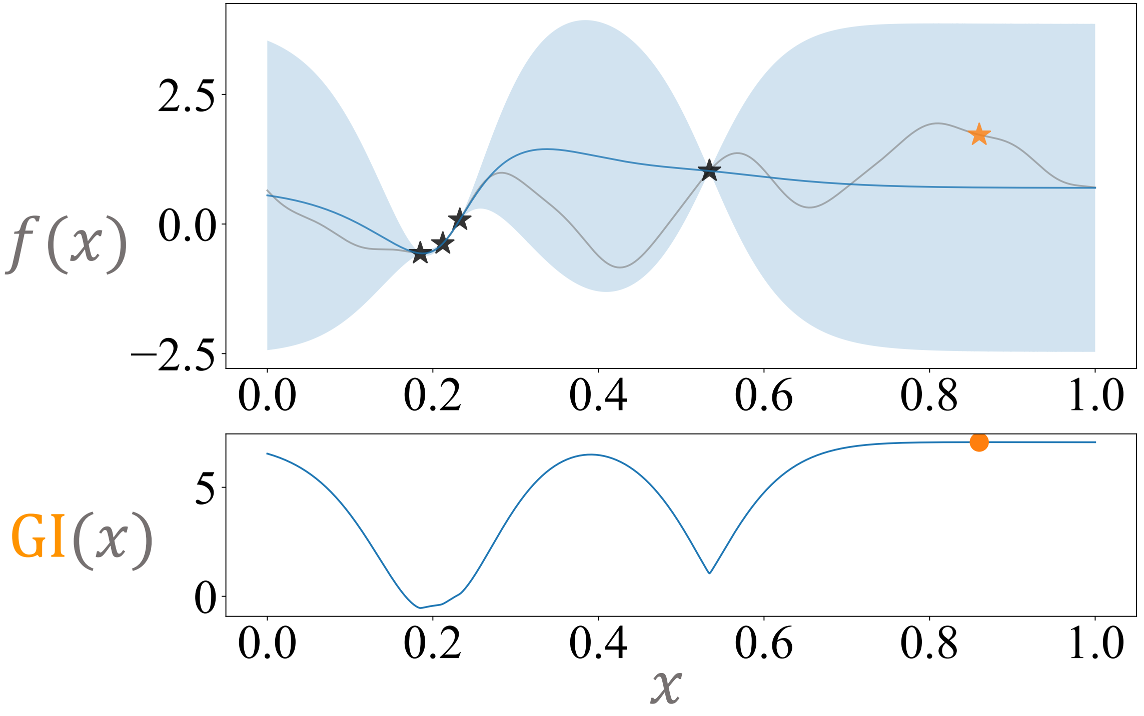

Gittins Indices

Easy to compute!

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

How to implement Gittins Indices?

bisection search

$\mathrm{GI}_{f|D,c}(x) := $ solution g s.t.

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

Data

analytical expression
& monotonicity in g

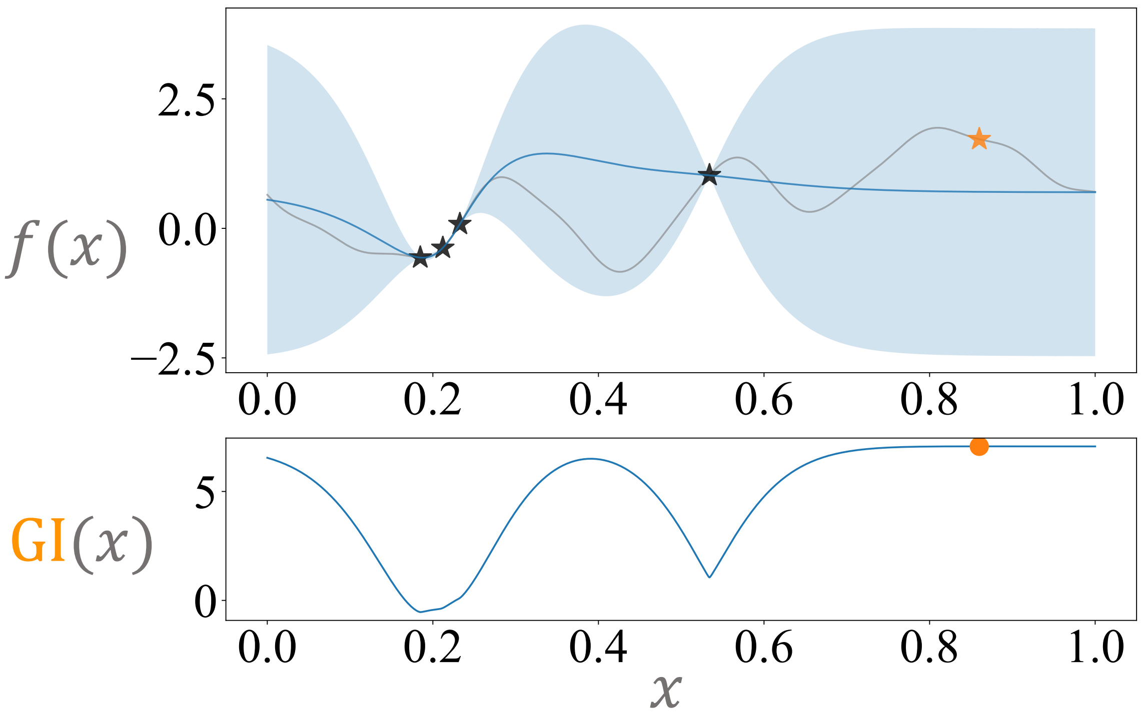

Gittins Indices

$$x_{t+1} = \mathrm{argmax}_x \, \mathrm{GI}_{f|D,c}(x)$$

Analytical gradient

Easy to compute and optimize!

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

How to implement Gittins Indices?

bisection search

$\mathrm{GI}_{f|D,c}(x) := \text{solution } g \text{ s.t.}$

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

$\mathrm{EI}_{f|D}(x; g)$

Data

& monotonicity in g

$$x_{t+1} = \mathrm{argmax}_x \, \mathrm{GI}_{f|D,c}(x)$$

Analytical gradient

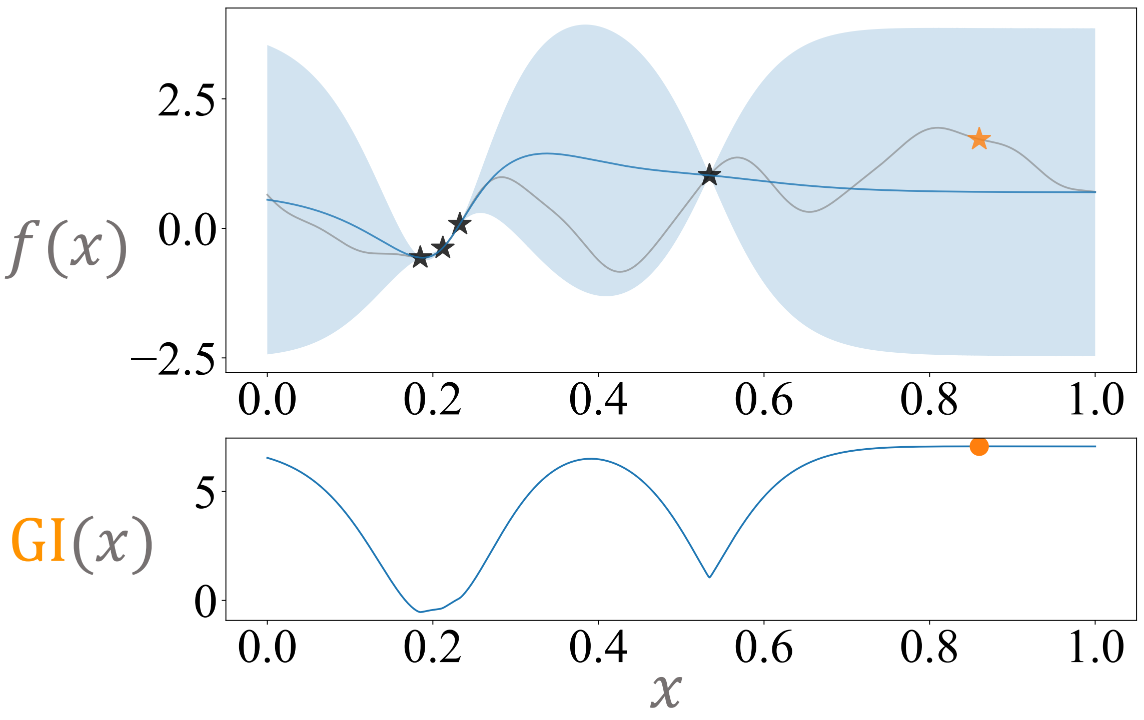

Gittins Indices

Easy to compute and optimize!

Expected Improvement

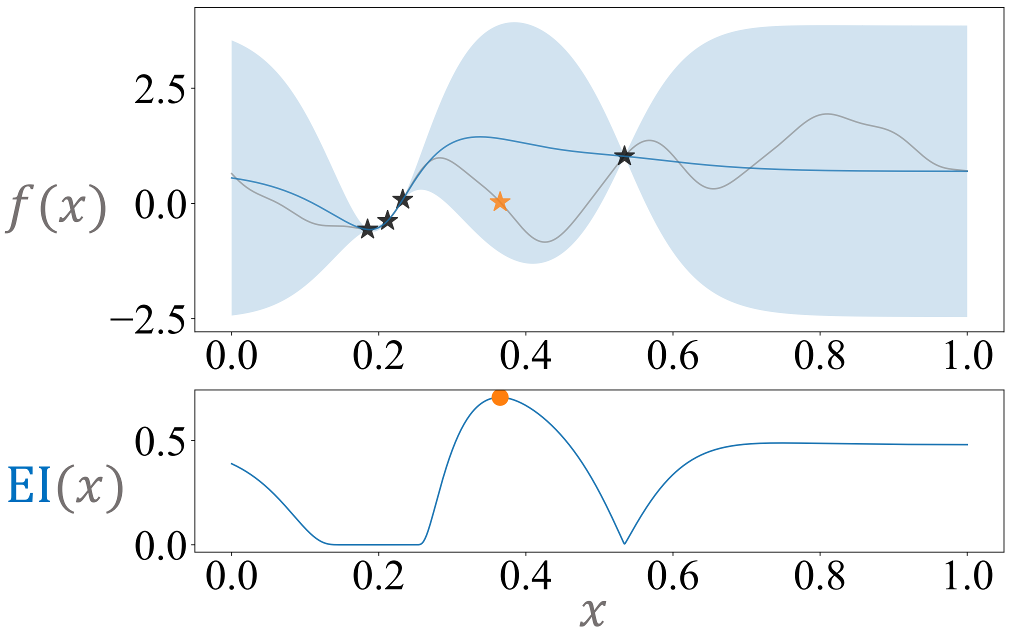

$$\mathrm{EI}_{f|D}(x; y_{\text{best}}) := \mathbb{E}[\max(f(x) - y_{\text{best}}, 0) \mid D]$$

Acquisition rule: $\operatorname{argmax}_x \mathrm{EI}_{f|D}(x; y_{\text{best}})$

Gittins Index

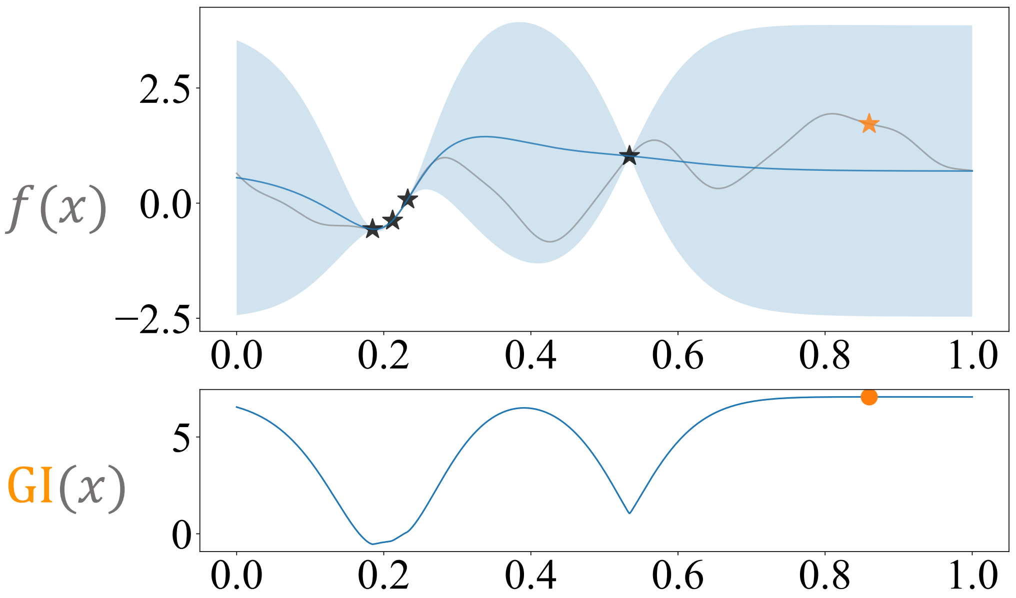

$$\mathrm{GI}_{f|D,c}(x) := \text{solution } g \text{ s.t. } \mathrm{EI}_{f|D}(x; g) = c(x)$$

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

$\operatorname{argmax}_x \mathrm{GI}_{f|D,c}(x)$

Expected Improvement

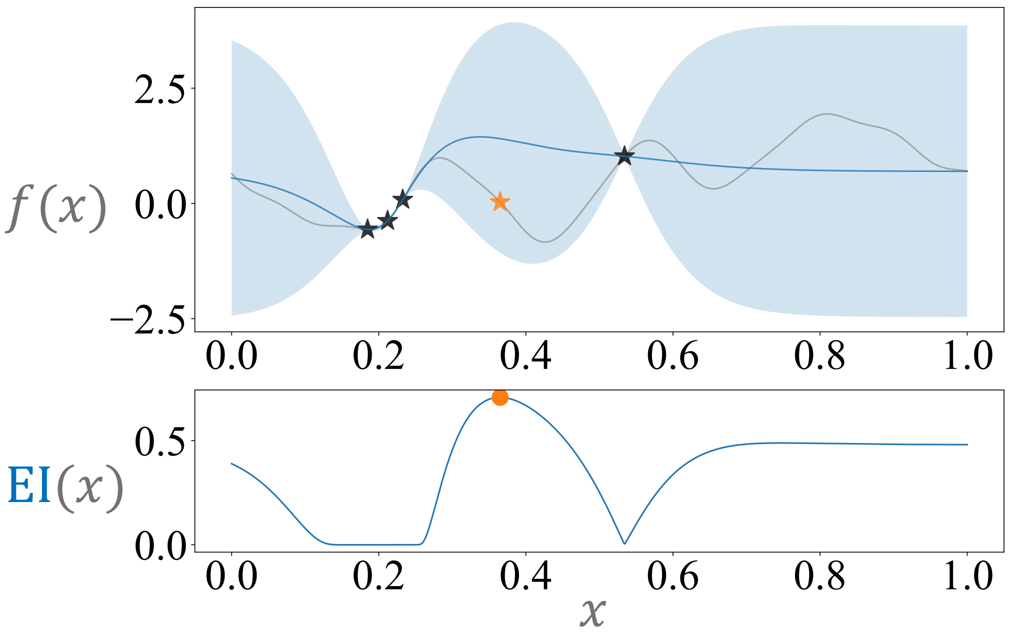

$\mathrm{EI}_{f|D}(x; y_{\mathrm{best}}) := \mathbb{E}[\max(f(x) - y_{\mathrm{best}}, 0) \mid D]$

Gittins Index

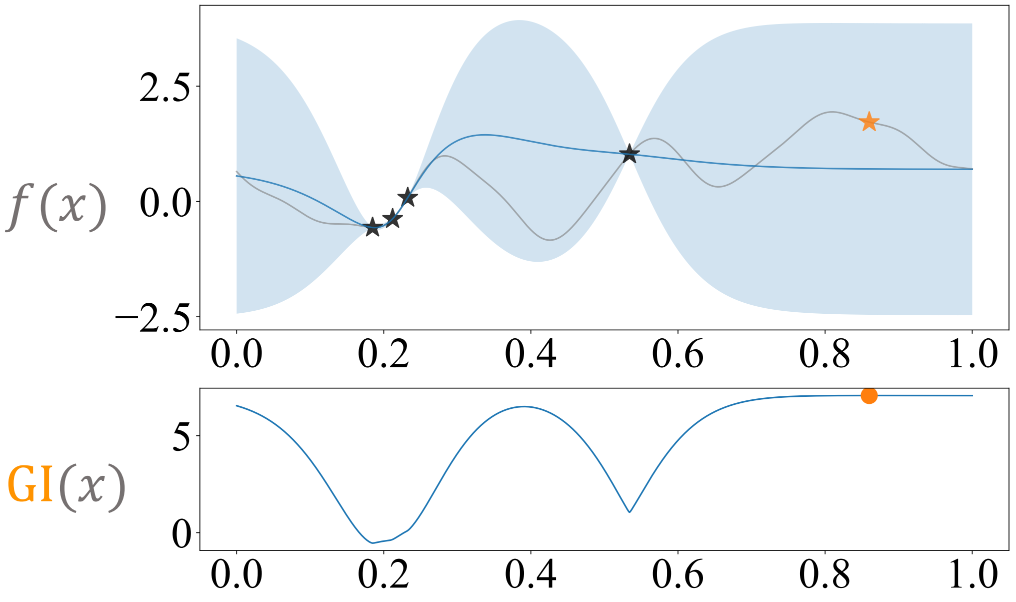

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t. $\mathrm{EI}_{f|D}(x; g) = c(x)$
where $\mathrm{EI}_{f|D}(x; g) := \mathbb{E}[\max(f(x) - g, 0) \mid D]$

Acquisition rule: $\operatorname{argmax}_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$ | $\operatorname{argmax}_x \mathrm{GI}_{f|D,c}(x)$

Stopping rule: | $\tau: \mathrm{GI}_{f|D,c}(x_\tau) \le y_{\mathrm{best}}$

Expected Improvement

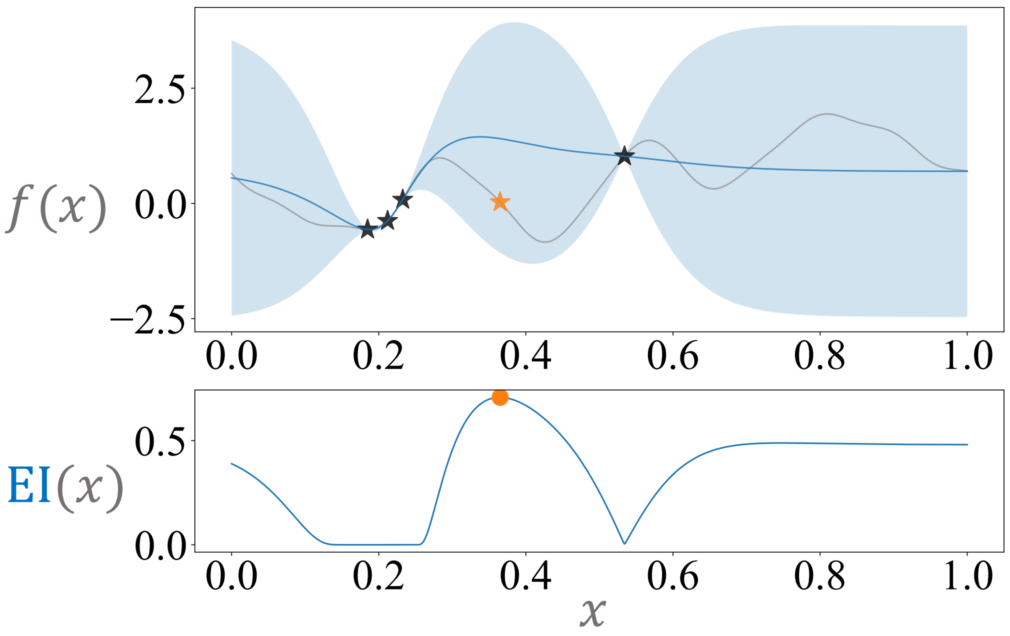

Gittins Index

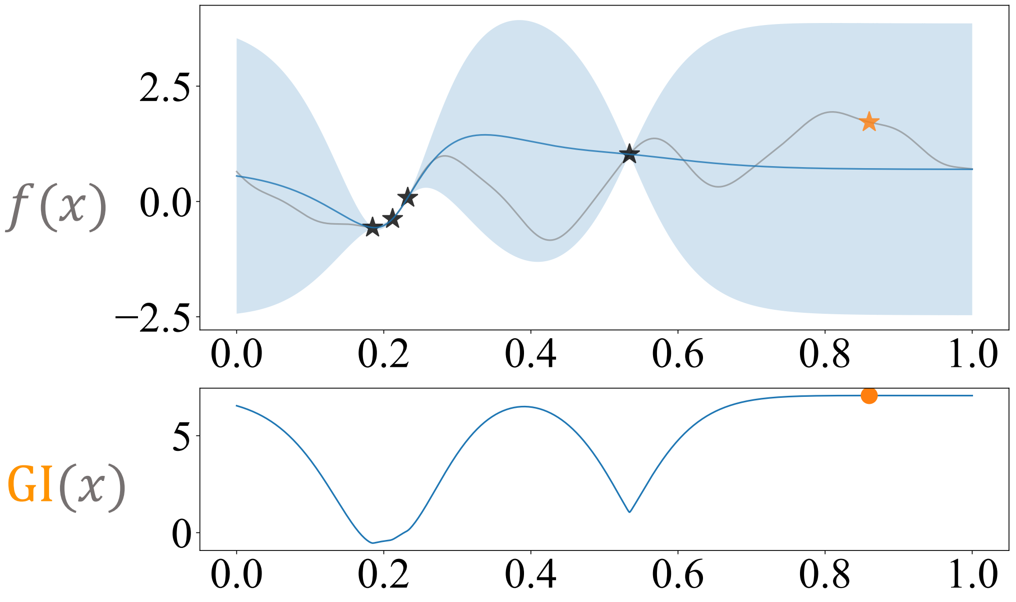

$\mathrm{EI}_{f|D}(x; y_{\mathrm{best}}) := \mathbb{E}[\max(f(x) - y_{\mathrm{best}}, 0) \mid D]$

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

Acquisition rule: $\operatorname{argmax}_x \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$ $\qquad$ $\operatorname{argmax}_x \mathrm{GI}_{f|D,c}(x)$

Stopping rule: $\tau: \mathrm{GI}_{f|D,c}(x_\tau) \le y_{\mathrm{best}}$

$\Leftrightarrow \tau: \mathrm{EI}_{f|D}(x; y_{\mathrm{best}}) \le \mathrm{EI}_{f|D}\left(x; \mathrm{GI}_{f|D,c}(x_\tau)\right) = c(x_\tau)$

Expected Improvement

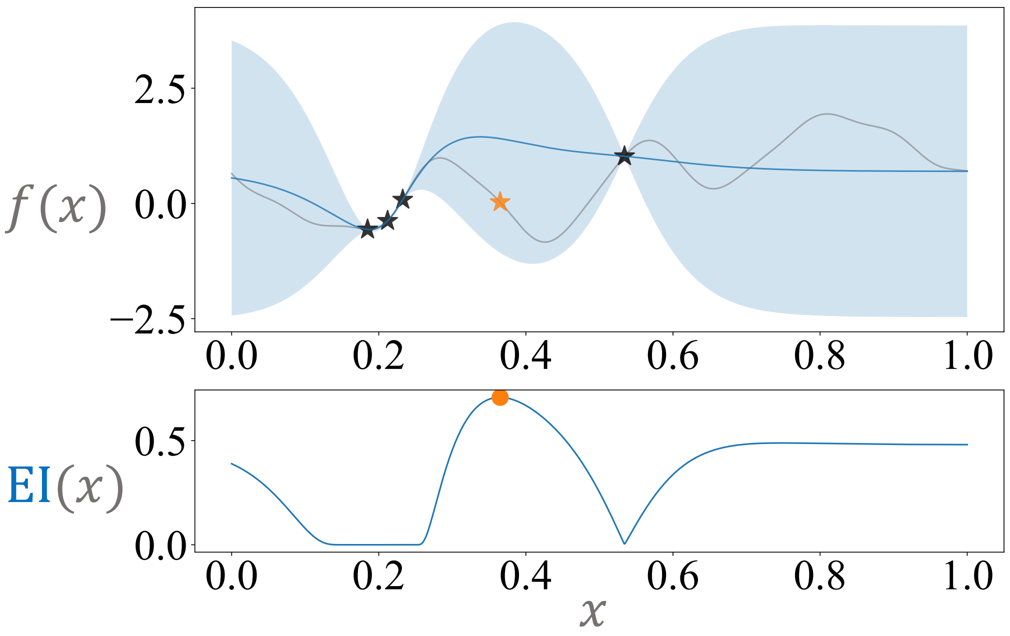

Gittins Index

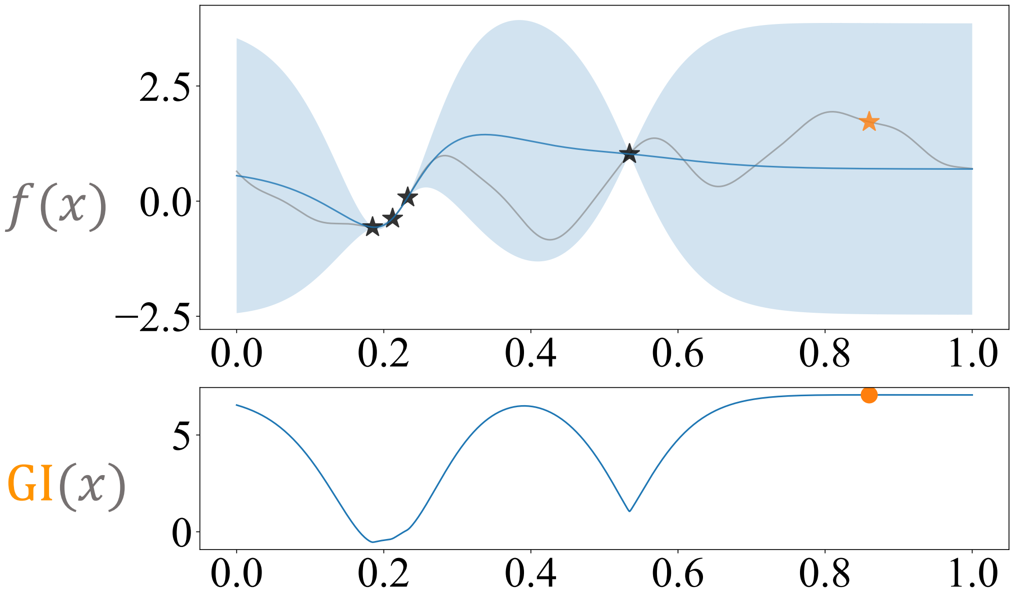

$$\mathrm{EI}_{f|D}(x; y_{\text{best}}) := \mathbb{E}[\max(f(x) - y_{\text{best}}, 0) \mid D]$$

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

Acquisition rule: $\operatorname{argmax}_x \mathrm{EI}_{f|D}(x; y_{\text{best}})$ $\qquad$ $\operatorname{argmax}_x \mathrm{GI}_{f|D,c}(x)$

Stopping rule: $\tau: \mathrm{EI}_{f|D}(x_\tau; y_{\text{best}}) \le c(x_\tau)$ $\qquad$ $\tau: \mathrm{GI}_{f|D,c}(x_\tau) \le y_{\text{best}}$

$$\Leftrightarrow \tau: \mathrm{EI}_{f|D}(x; y_{\text{best}}) \le \mathrm{EI}_{f|D}\left(x; \mathrm{GI}_{f|D,c}(x_\tau)\right) = c(x_\tau)$$

Expected Improvement

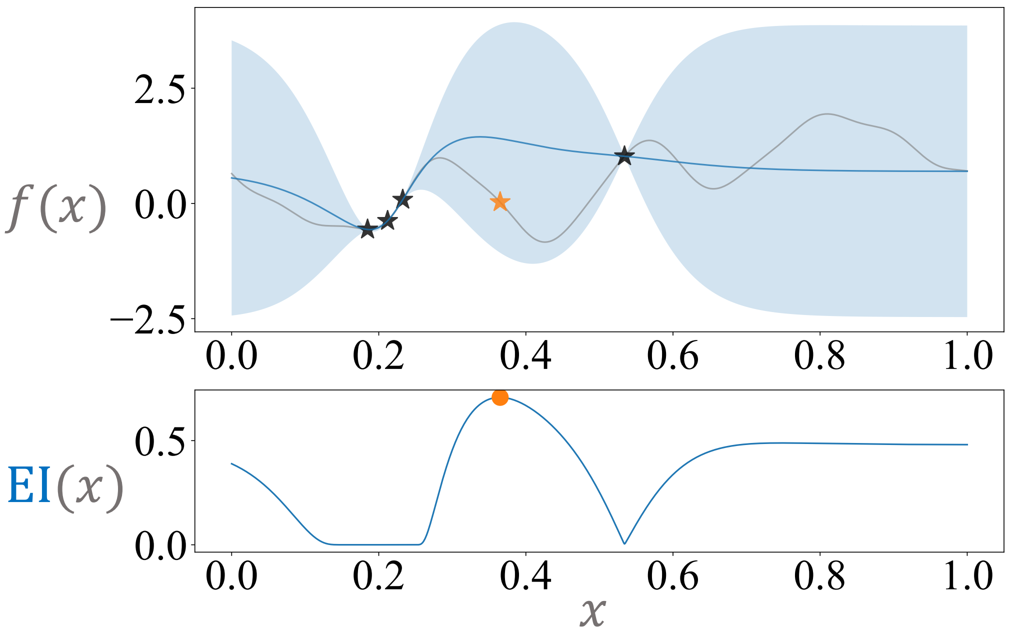

Gittins Index

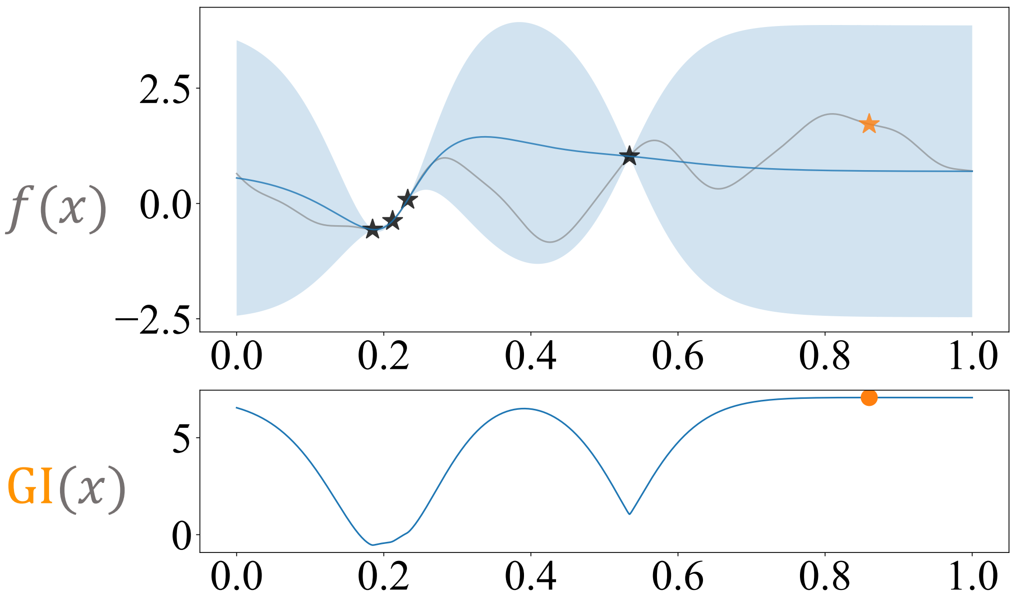

$\mathrm{EI}_{f|D}(x; y_{\mathrm{best}}) := \mathbb{E}[\max(f(x) - y_{\mathrm{best}}, 0) \mid D]$

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

Acquisition rule: $\mathrm{argmax}_x\, \mathrm{EI}_{f|D}(x; y_{\mathrm{best}})$ $\qquad$ $\mathrm{argmax}_x\, \mathrm{GI}_{f|D,c}(x)$

Stopping rule: $\tau: \mathrm{EI}_{f|D}(x_\tau; y_{\mathrm{best}}) \le c(x_\tau)$
$\iff \tau: \mathrm{EIPC}_{f|D,c}(x_\tau; y_{\mathrm{best}}) \le 1$

$\tau: \mathrm{GI}_{f|D,c}(x_\tau) \le y_{\mathrm{best}}$
$\iff \tau: \mathrm{EI}_{f|D}(x; y_{\mathrm{best}}) \le \mathrm{EI}_{f|D}\left(x; \mathrm{GI}_{f|D,c}(x_\tau)\right) = c(x_\tau)$

Expected Improvement vs Gittins Index

Varying evaluation costs

$\text{EIPC}_c(x; y_{\text{best}}) = \text{EI}(x; y_{\text{best}})/c(x)$

Arbitrarily bad

$\text{GI}_c(x) := \text{solution } g \text{ s.t. } \text{EI}(x; g) = c(x)$

Costs built-in

Smart stopping time

$\tau\text{: } \text{EI}(x_\tau; y_{\text{best}}) \le \theta$

Which threshold?

$\tau\text{: } \text{GI}_c(x_\tau) \le y_{\text{best}}$

$\Longleftrightarrow \tau\text{: } \text{EIPC}_c(x_\tau; y_{\text{best}}) \le 1$

Derived shared stopping rule

Temporal (one-step) simplification to MDP

Spatial simplification to MDP

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

"Cost-aware Stopping for Bayesian Optimization." ICML'26

Theoretical Guarantee and Empirical Validation

Theorem (Safeguard Guarantee)

$\mathbb{E}[R(\text{ours}; \text{PBGI})] \leq R[\text{stopping immediately}]$

or LogEIPC

cost-adjusted regret

Implication:

• Matches the best achievable performance in the worst case (evaluations are all very costly).

• Avoids over-spending — a property many cost-unaware stopping rules lack.

New

Proof idea: For all $t < \tau$, $\text{EI}(x_{t+1}) \geq c(x_{t+1})$.

stopping time

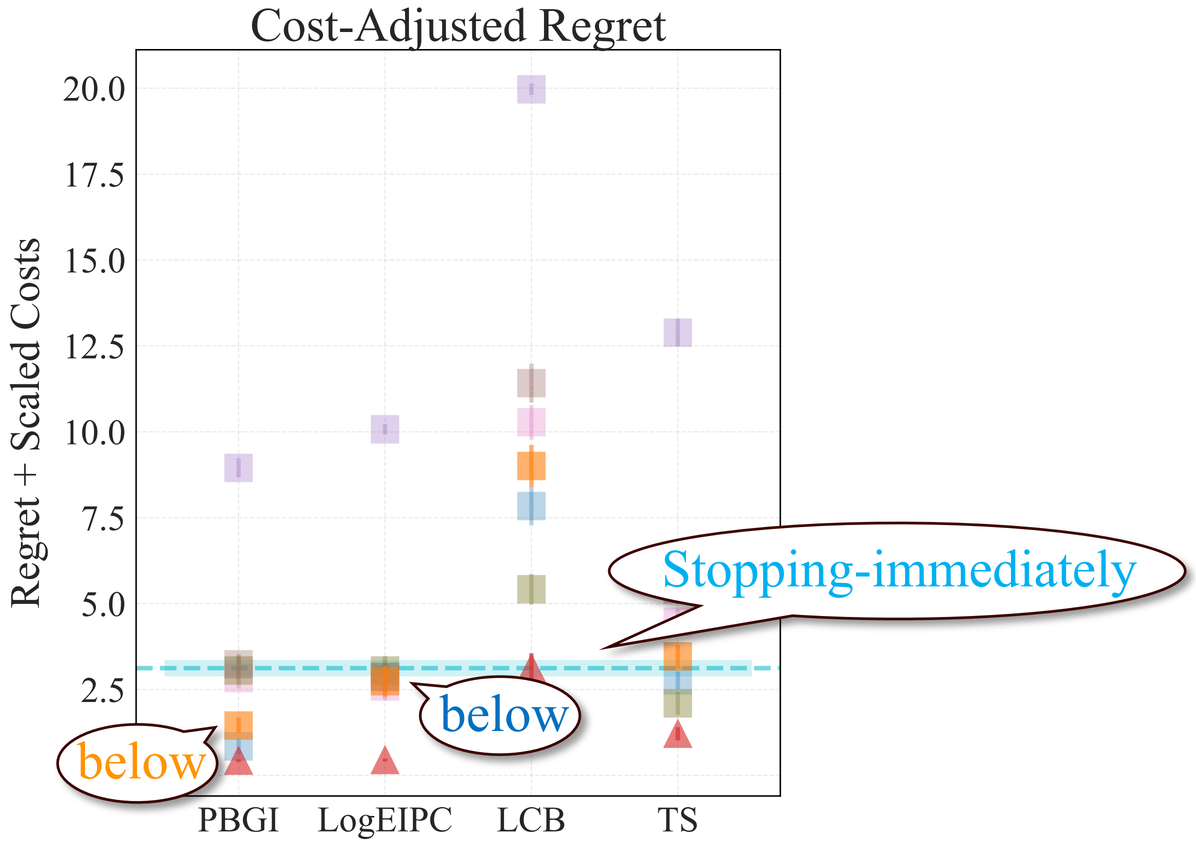

"Cost-aware Stopping for Bayesian Optimization." ICML'26

Gittins Index vs Baselines on LCBench

ML model training:

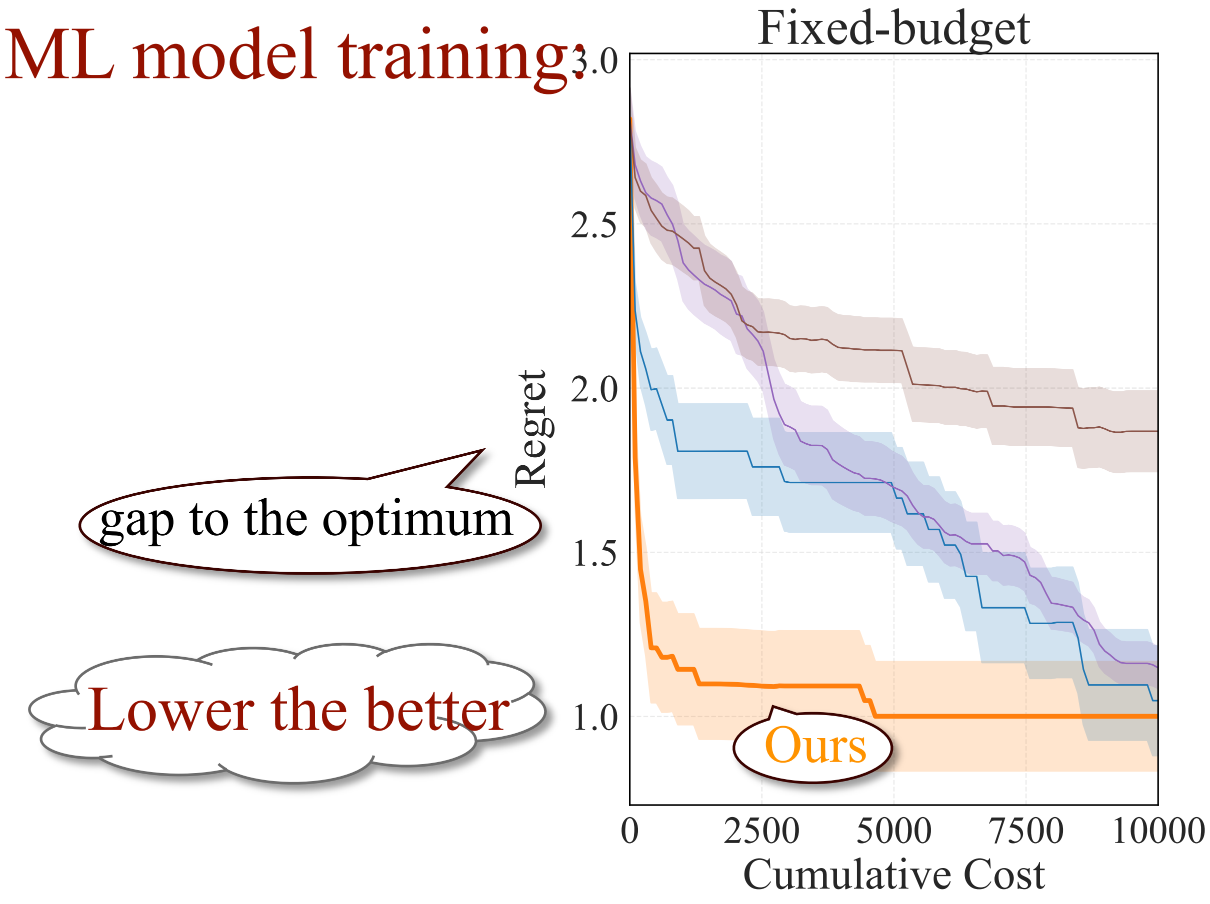

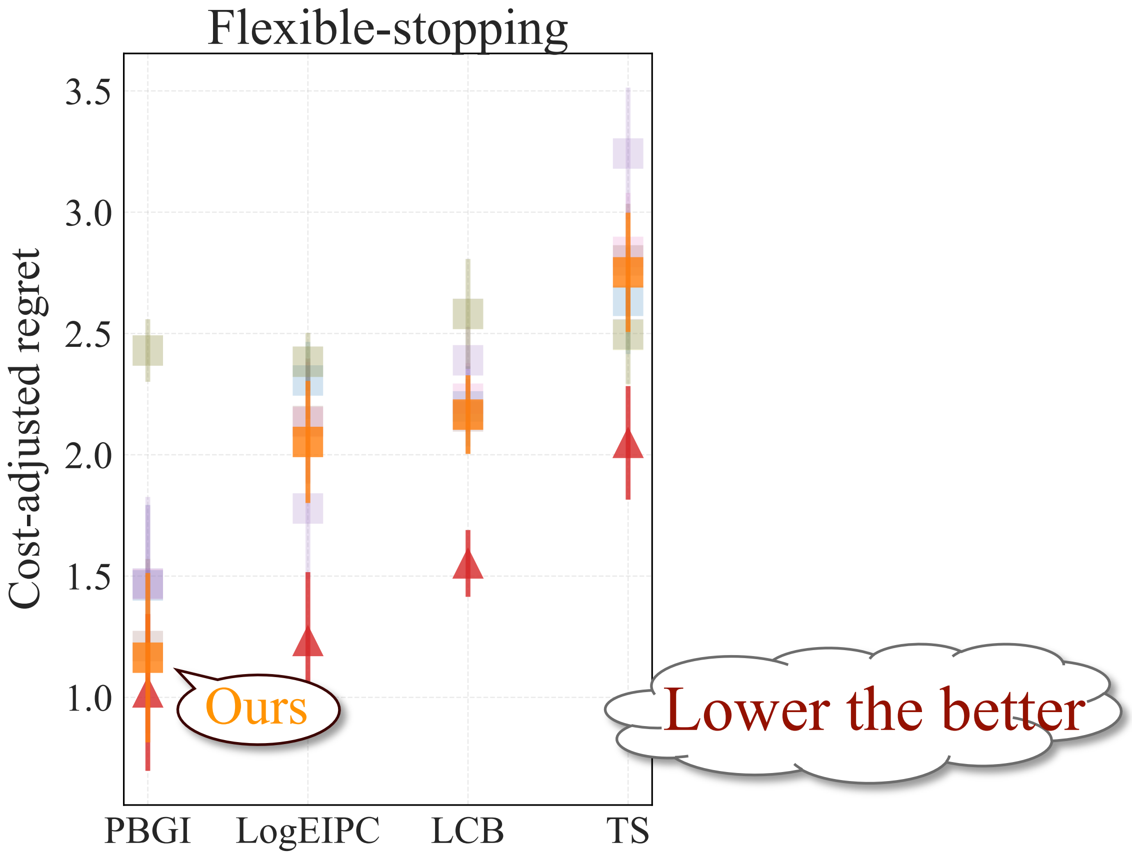

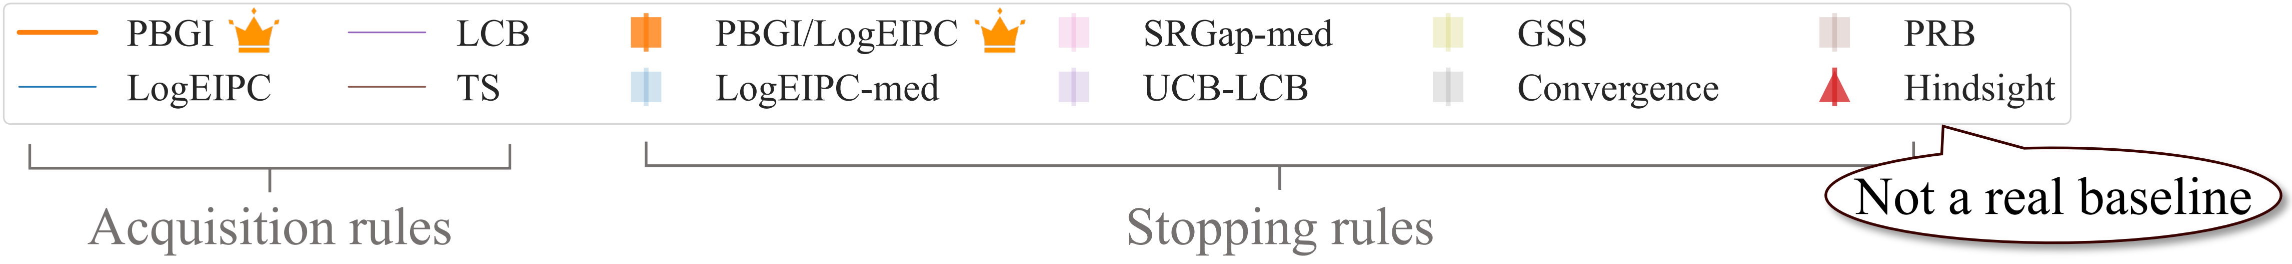

Gittins Index vs Baselines on LCBench

ML model training:

Cost-scaling factor

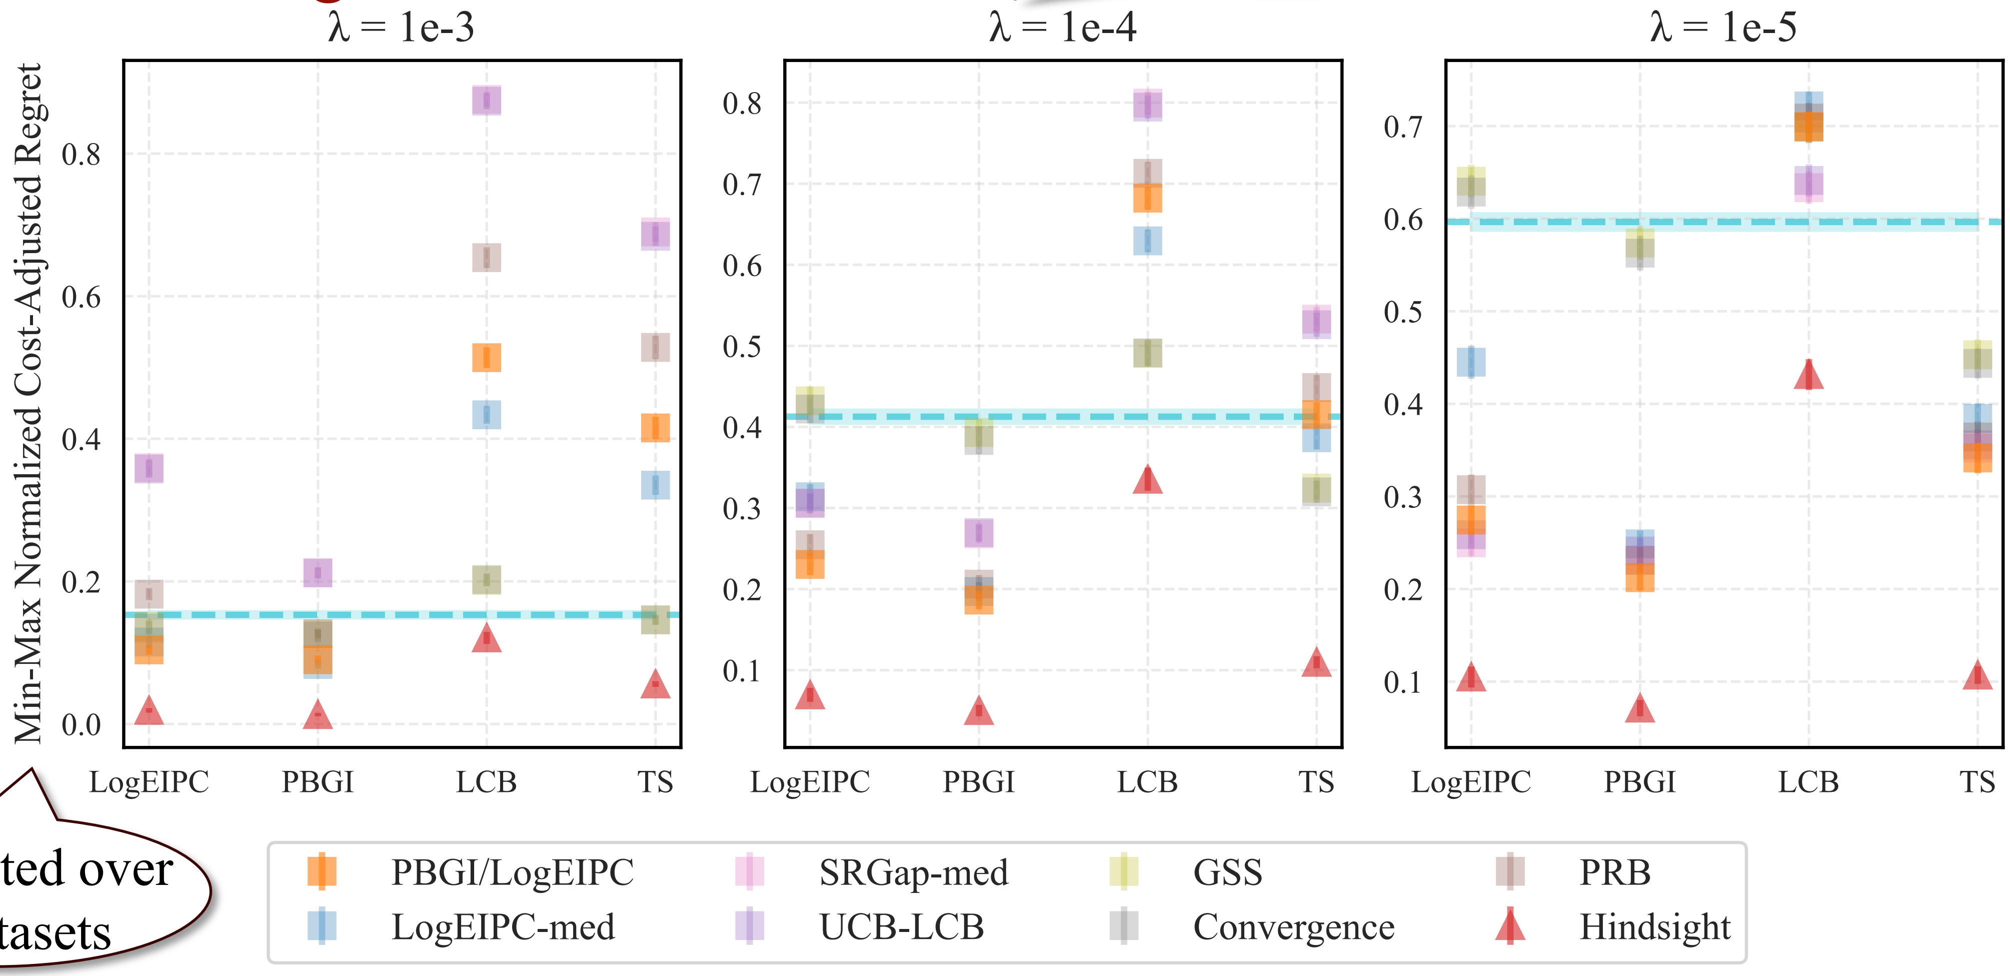

Aggregated over 35 datasets

"Cost-aware Stopping for Bayesian Optimization." ICML'26

Gittins Index vs Baselines on LCBench

ML model training:

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%

(All: 35 datasets, Large: 20 datasets w/ >10,000 instances)

"Cost-aware Stopping for Bayesian Optimization." ICML'26

Part II:

Towards Bayesian Optimization with Multi-Stage Feedback (Bandit-like Bayesian Decision Problems) via Markov Chain Gittins Indices

Model	Q1	Q2	Q3	…	avg performance
ChatGPT GPT-4o	✓	?	?	…	?
Claude 3.5 Sonnet	?	✓	?	…	?
Deepseek	✓	? Y_{ij}	?	…	?
Gemini 1.5 Pro	✖	?	✓	…	?
Llama 3.1-70B	?	✖	?	…	?
Mistral Large	?	?	✖	…	?
…	…	…	…	…	…

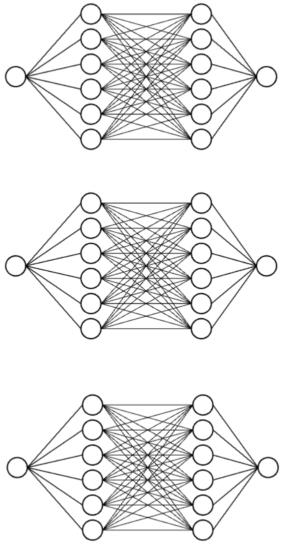

Best-Model Identification in LLM Evaluation

LLM evaluation:

Model →

Model	Q1	Q2	Q3	...	avg performance
ChatGPT GPT-4o	✓	?	?	...	?
Claude 3.5 Sonnet	?	✓	?	...	?
Deepseek	✓	? Y_{ij}	?	...	?
Gemini 1.5 Pro	✖	?	✓	...	?
Llama 3.1-70B	?	✖	?	...	?
Mistral Large	?	?	✖	...	?
...	...	...	...	...	...

→ Average score over questions

Goal: find LLM with best average score

(find row with best latent mean for a matrix)

Best-Model Identification in LLM Evaluation

LLM evaluation:

Model →

Model	Q1	Q2	Q3	...	avg performance
ChatGPT GPT-4o	✓	?	?	...	?
Claude 3.5 Sonnet	?	✓	?	...	?
Deepseek	✓	? Y_{ij}	?	...	?
Gemini 1.5 Pro	✖	?	✓	...	?
Llama 3.1-70B	?	✖	?	...	?
Mistral Large	?	?	✖	...	?
...	...	...	...	...	...

$ → Average score over questions

Goal: find LLM with best average score

(find row with best latent mean for a matrix)

Best-Model Identification in LLM Evaluation

LLM evaluation:

Model →

Model	Q1	Q2	Q3	...	avg performance
ChatGPT GPT-4o	✓	?	?	...	?
Claude 3.5 Sonnet	?	✓	?	...	?
Deepseek	✓	? Y_{ij}	?	...	?
Gemini 1.5 Pro	✖	?	✓	...	?
Llama 3.1-70B	?	✖	?	...	?
Mistral Large	?	?	✖	...	?
...	...	...	...	...	...

$$$

→ Average score over questions

Goal: find LLM with best average score

(find row with best latent mean for a matrix)

Best-Model Identification in LLM Evaluation

LLM evaluation:

Model →

Model	Q1	Q2	Q3	…	avg performance
ChatGPT GPT-4o	✓	?	?	…	?
Claude 3.5 Sonnet	?	✓	?	…	?
Deepseek	✓	? Y_{ij}	?	…	?
Gemini 1.5 Pro	✖	?	✓	…	?
Llama 3.1-70B	?	✖	?	…	?
Mistral Large	?	?	✖	…	?
…	…	…	…	…	…

$$$

→ Average score over questions

Goal: find LLM with best average score

(find row with best latent mean for a matrix)

Cost: full evaluation

Pandora's Box

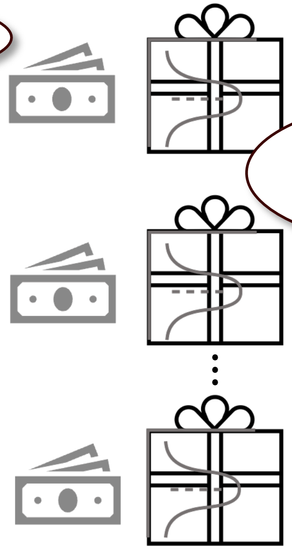

Reward: average score

Best-Model Identification in LLM Evaluation

LLM evaluation:

Model →

Model	Q1	Q2	Q3	...	avg performance
ChatGPT GPT-4o	✓	?	?	...	?
Claude 3.5 Sonnet	?	✓	?	...	?
Deepseek	✓	? Y_{ij}	?	...	?
Gemini 1.5 Pro	✖	?	✓	...	?
Llama 3.1-70B	?	✖	?	...	?
Mistral Large	?	?	✖	...	?
...	...	...	...	...	...

$ → Average score over questions

Goal: find LLM with best average score

(find row with best latent mean for a matrix)

Best-Model Identification in LLM Evaluation

LLM evaluation:

Model →

Model	Q1	Q2	Q3	...	avg performance
ChatGPT GPT-4o	✓	?	?	...	?
Claude 3.5 Sonnet	?	✓	?	...	?
Deepseek	✓	? Y_{ij}	?	...	?
Gemini 1.5 Pro	✖	?	✓	...	?
Llama 3.1-70B	?	✖	?	...	?
Mistral Large	?	?	✖	...	?
...	...	...	...	...	...

$$$

→ Average score over questions

Goal: find LLM with best average score

(find row with best latent mean for a matrix)

Best-Model Identification in LLM Evaluation

LLM evaluation:

Model →

Model	Q1	Q2	Q3	...	avg performance
ChatGPT GPT-4o	✓	?	?	...	?
Claude 3.5 Sonnet	?	✓	?	...	?
Deepseek	✓	? Y_{ij}	?	...	?
Gemini 1.5 Pro	✖	?	✓	...	?
Llama 3.1-70B	?	✖	?	...	?
Mistral Large	?	?	✖	...	?
...	...	...	...	...	...

$$$

→ Average score over questions

Goal: find LLM with best average score

(find row with best latent mean for a matrix)

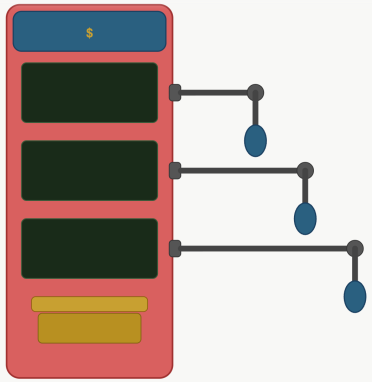

Cost-unaware & fixed-iteration

BanditEval

(Zhou et al. ICLR'25)

Best-Model Identification in LLM Evaluation

LLM evaluation:

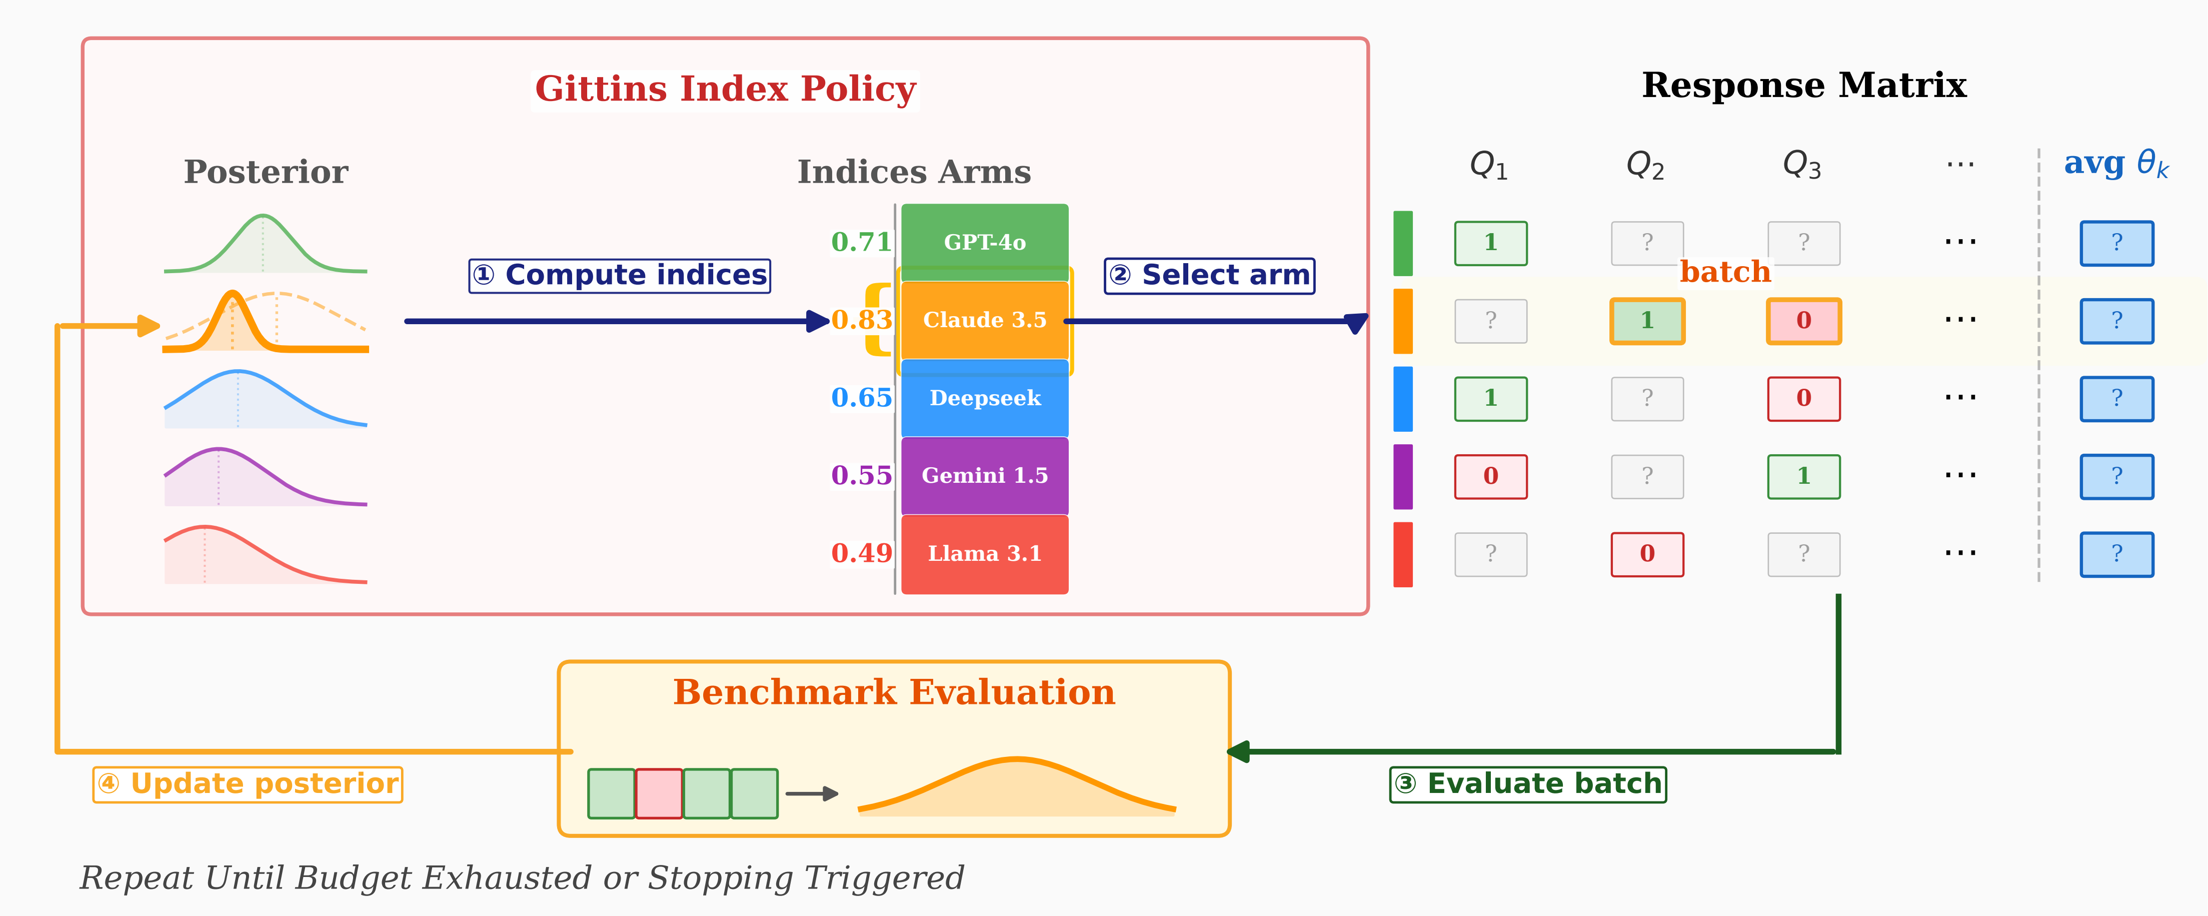

"Efficient Cost-Aware LLM Evaluation via Bayesian Bandit Gittins Indices." ICML'26 Workshop DEMO.

How to update posterior belief of latent mean?

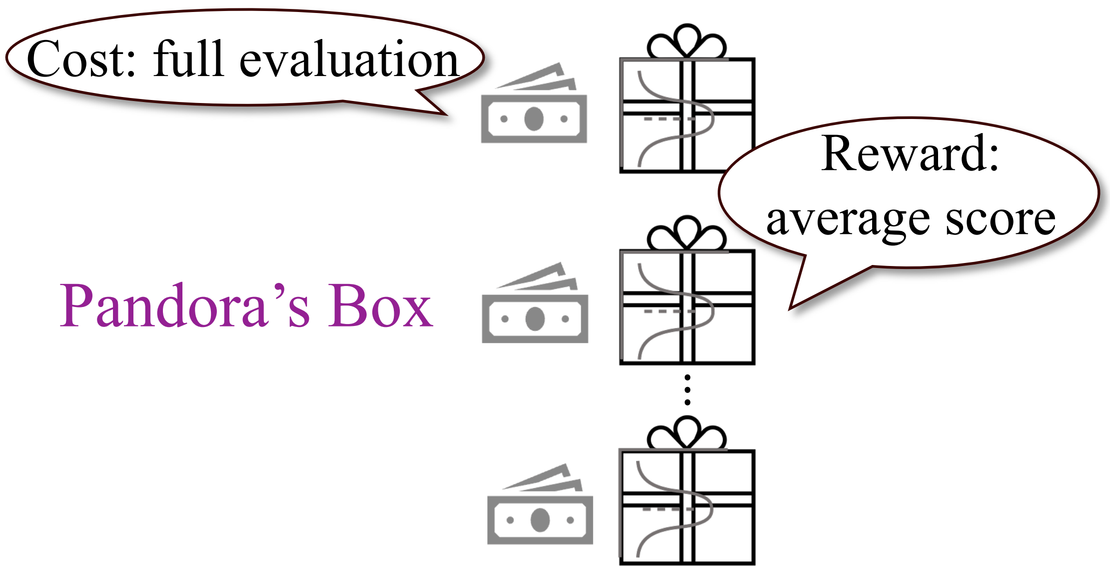

How to update posterior belief of latent mean?

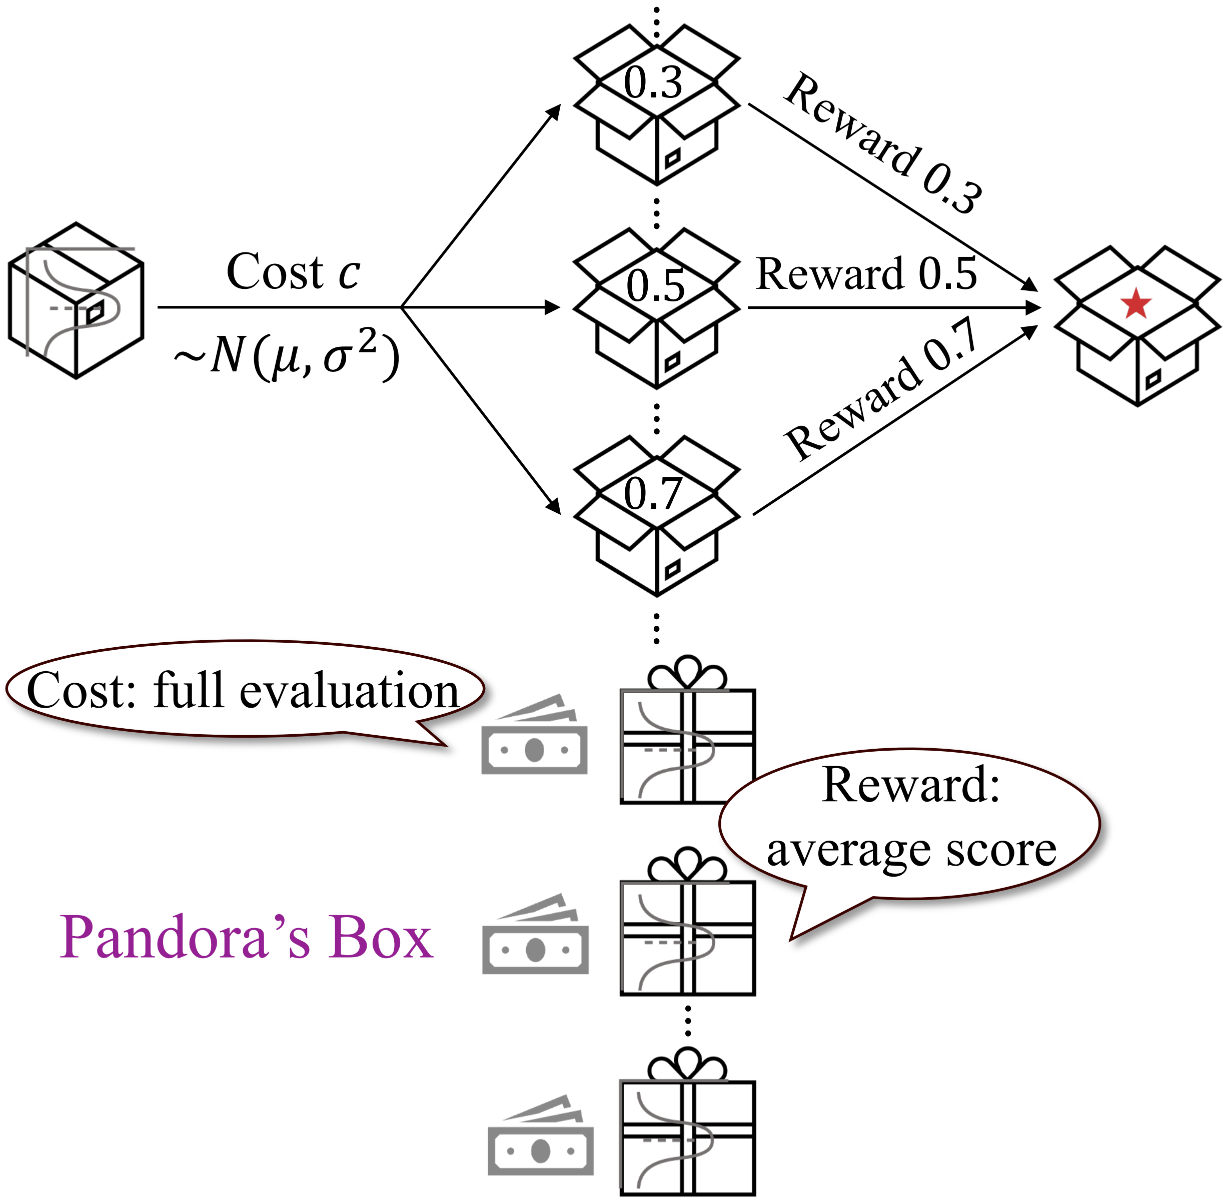

How to update posterior belief of latent mean?

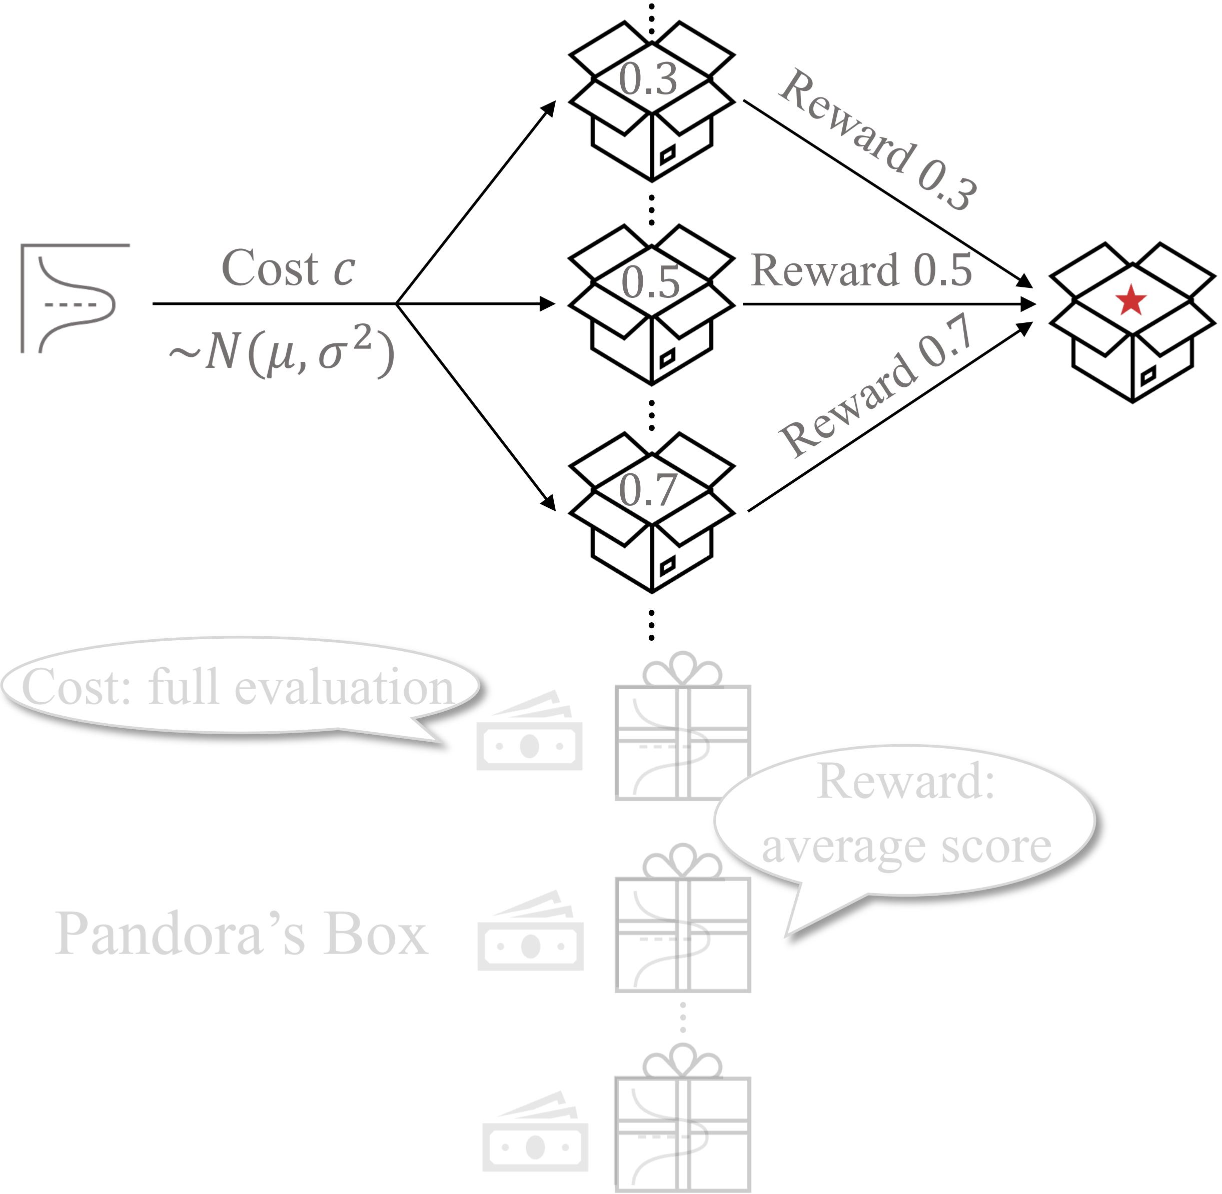

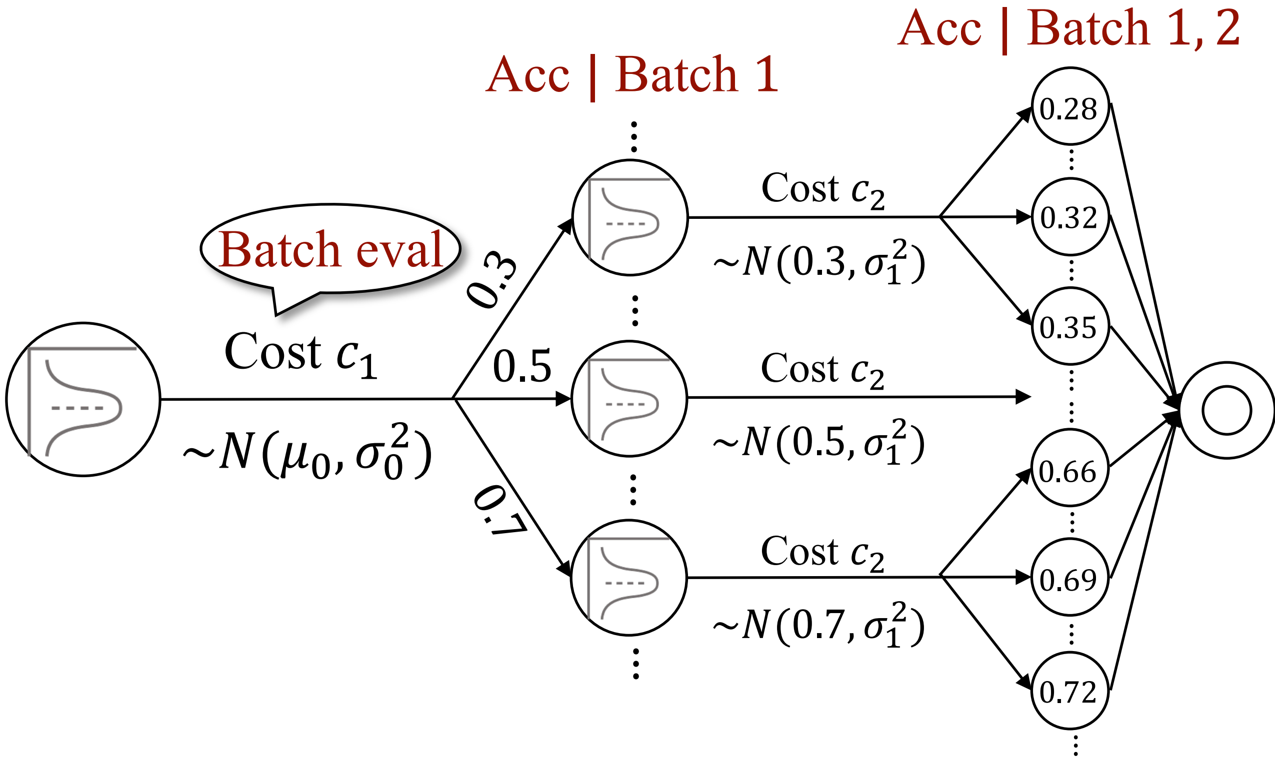

How to update posterior belief of latent mean?

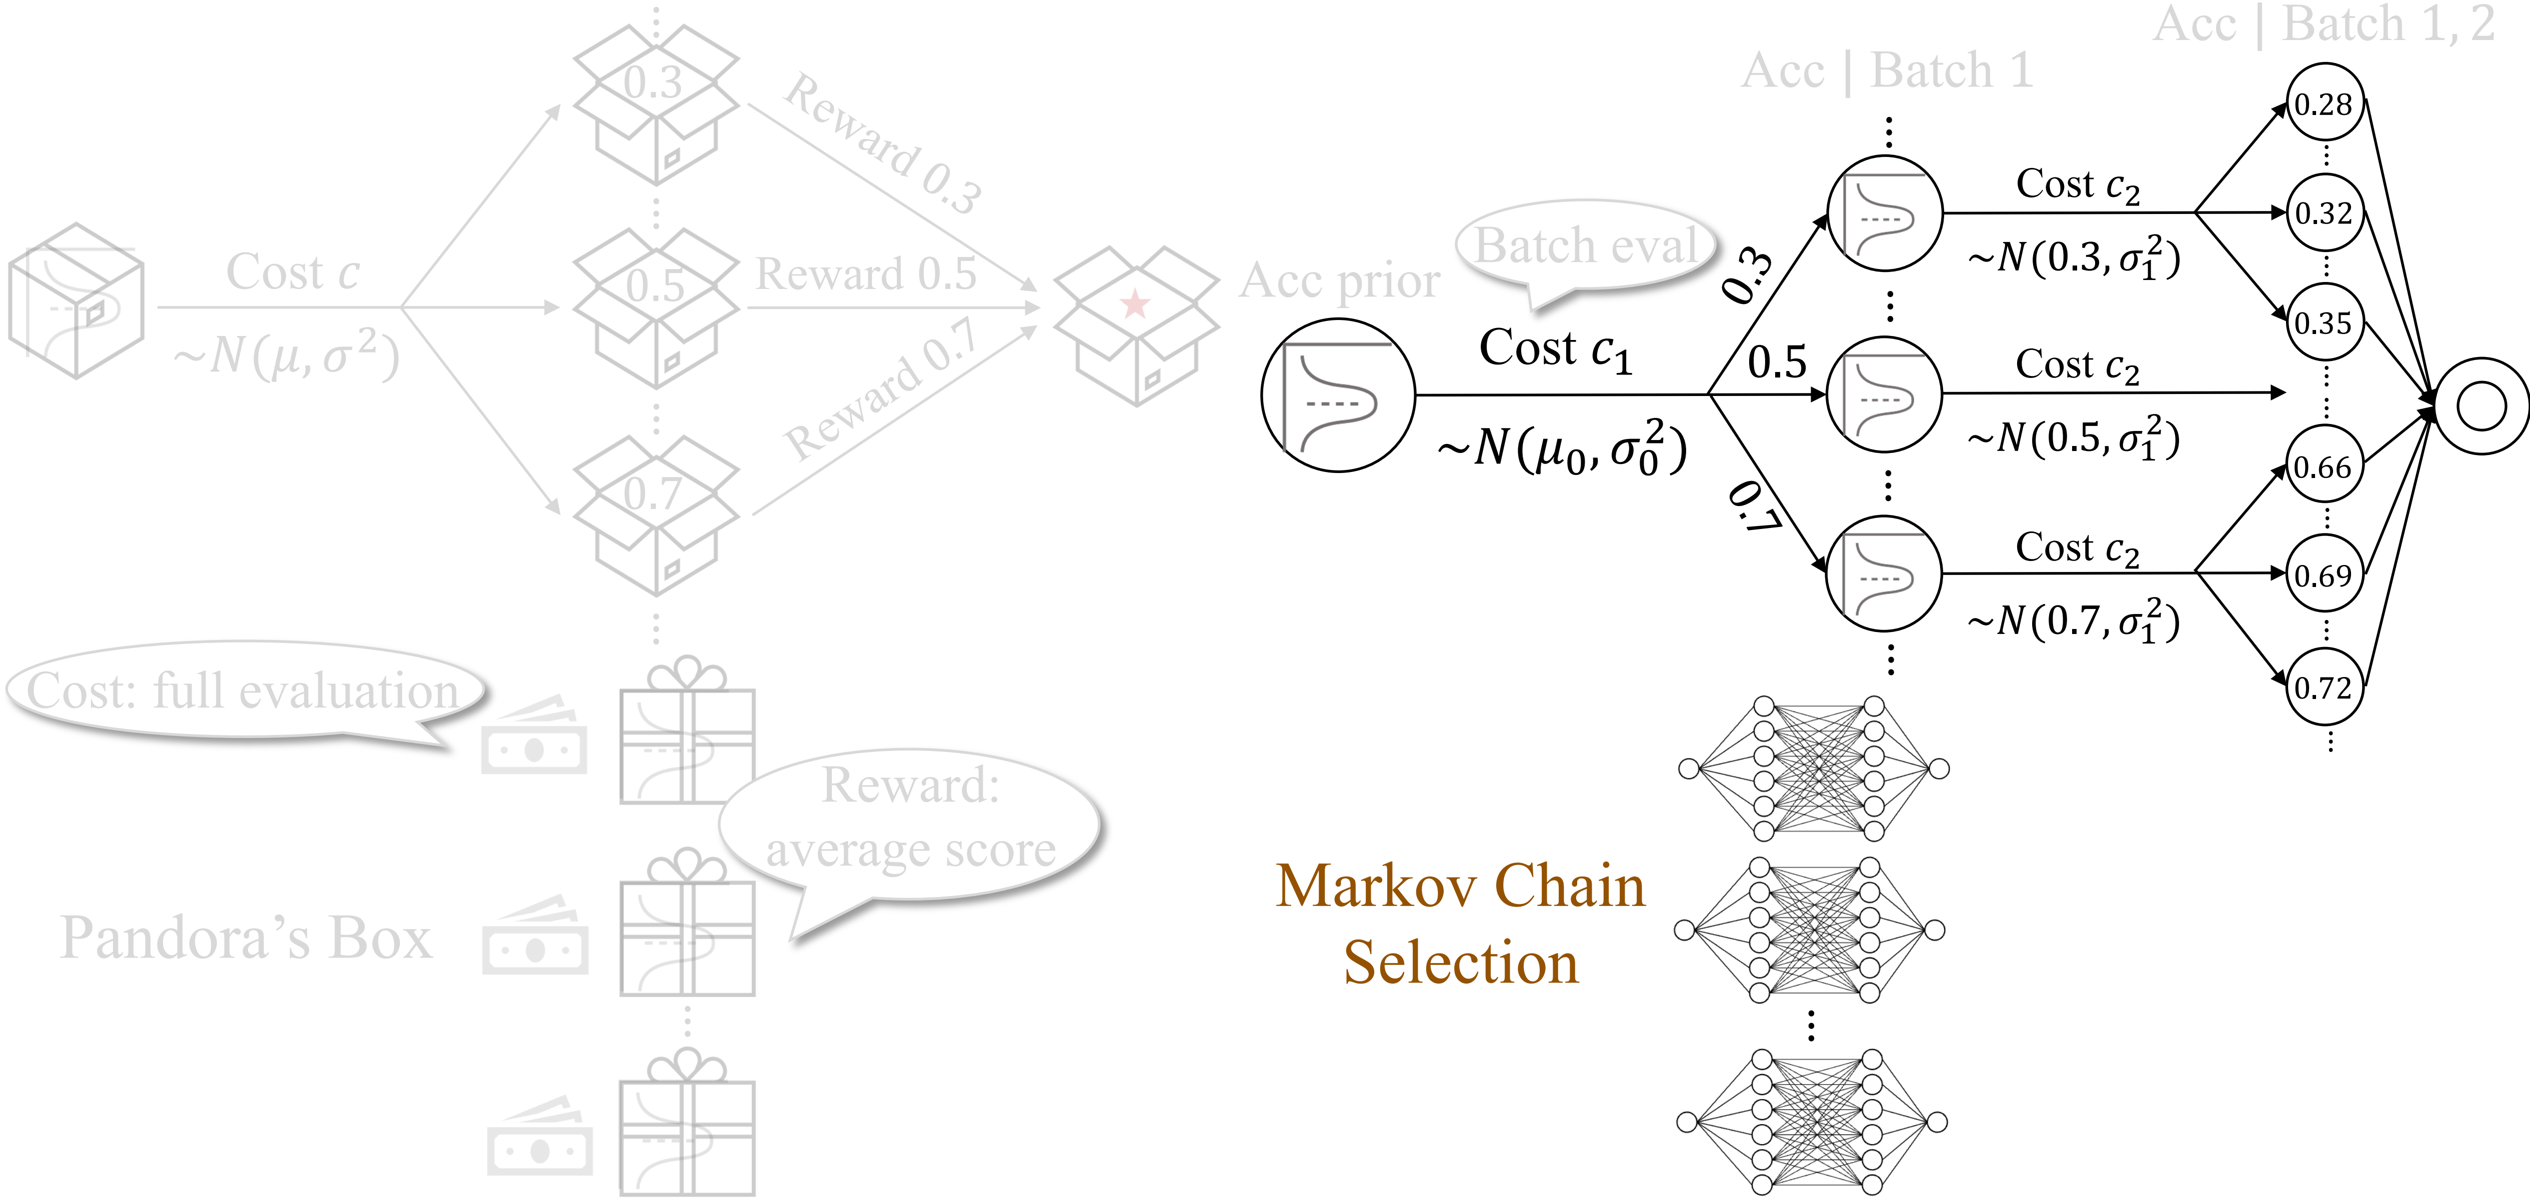

How to compute Gittins indices?

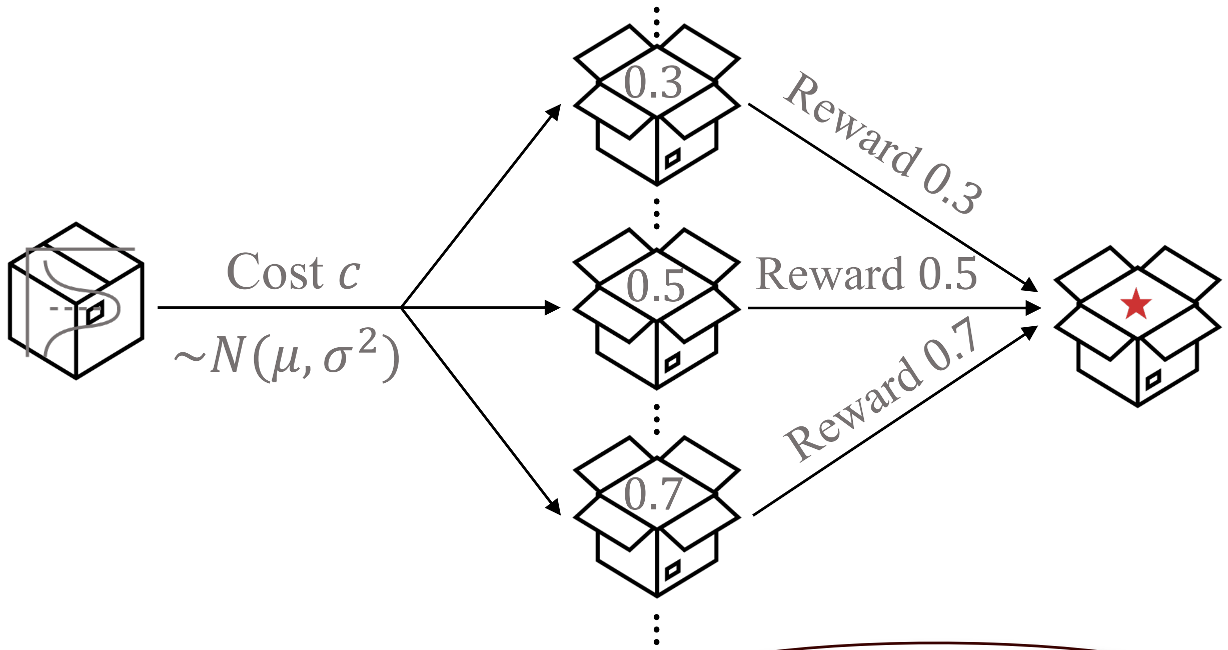

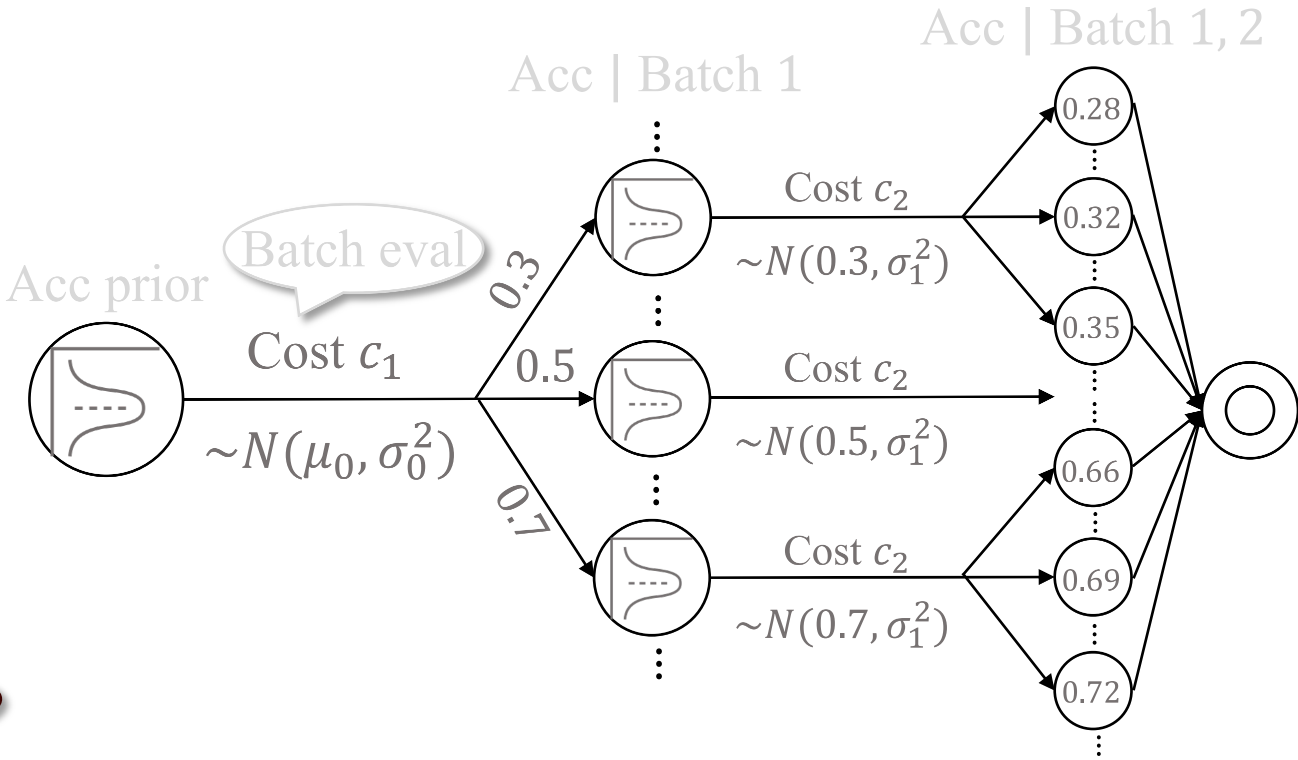

One bisection search

$$\mathrm{GI}_{f|D,c}(x) := \text{solution } g \text{ s.t.}$$

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

analytical expression

& monotonicity in g

How to compute Gittins indices?

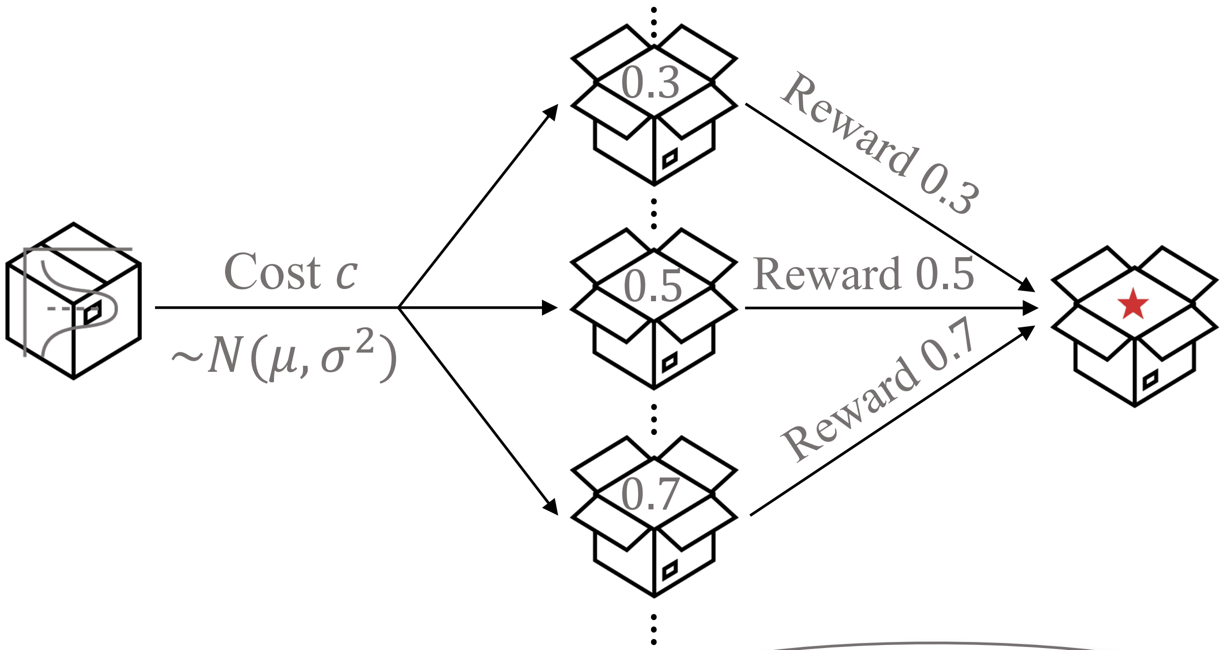

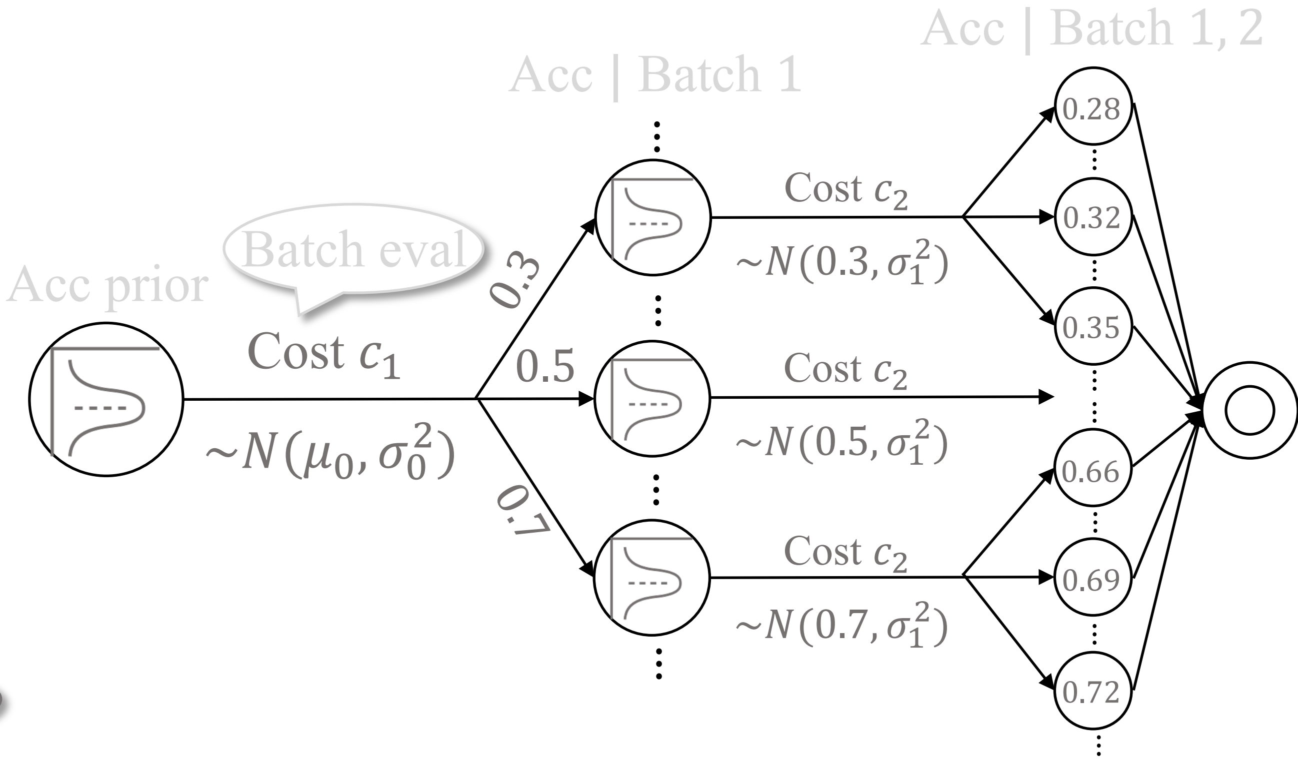

One bisection search

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

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

analytical expression & monotonicity in g

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

$$\mathbb{E}[\max(\mathbb{E}[\max(f(x) - g, 0) \mid D_1] - c_2(x), 0) \mid D] = c_1(x)$$

How to compute Gittins indices?

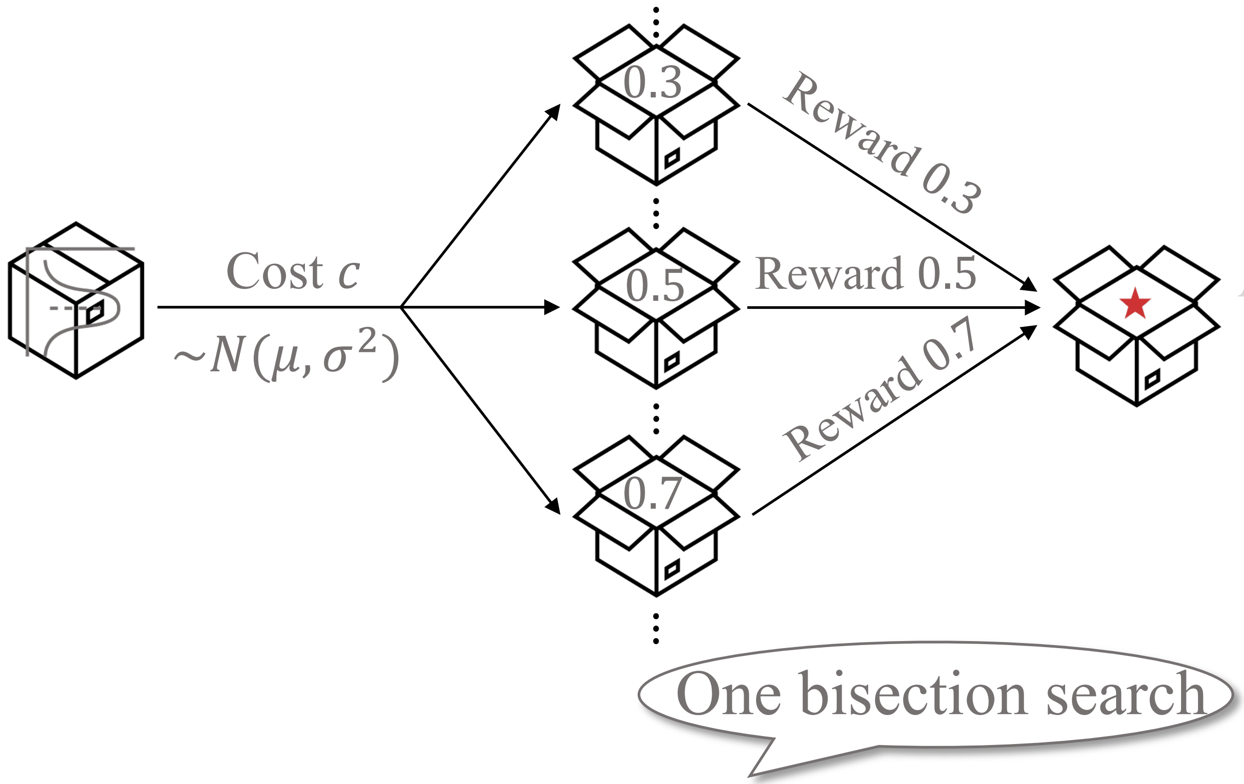

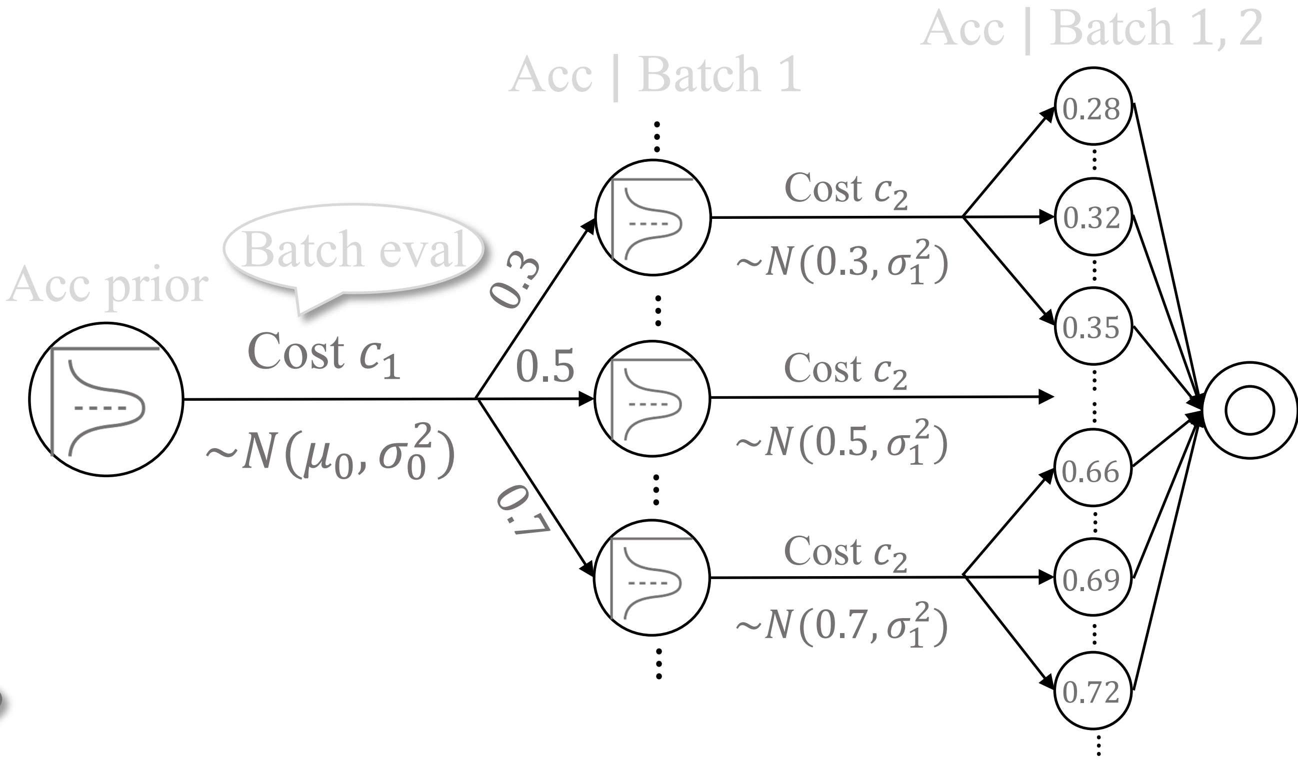

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

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

analytical expression & monotonicity in g

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

$$\mathbb{E}[\max(\mathbb{E}[\max(f(x) - g, 0) \mid D_1] - c_2(x), 0) \mid D] = c_1(x)$$

No analytical expression

How to compute Gittins indices?

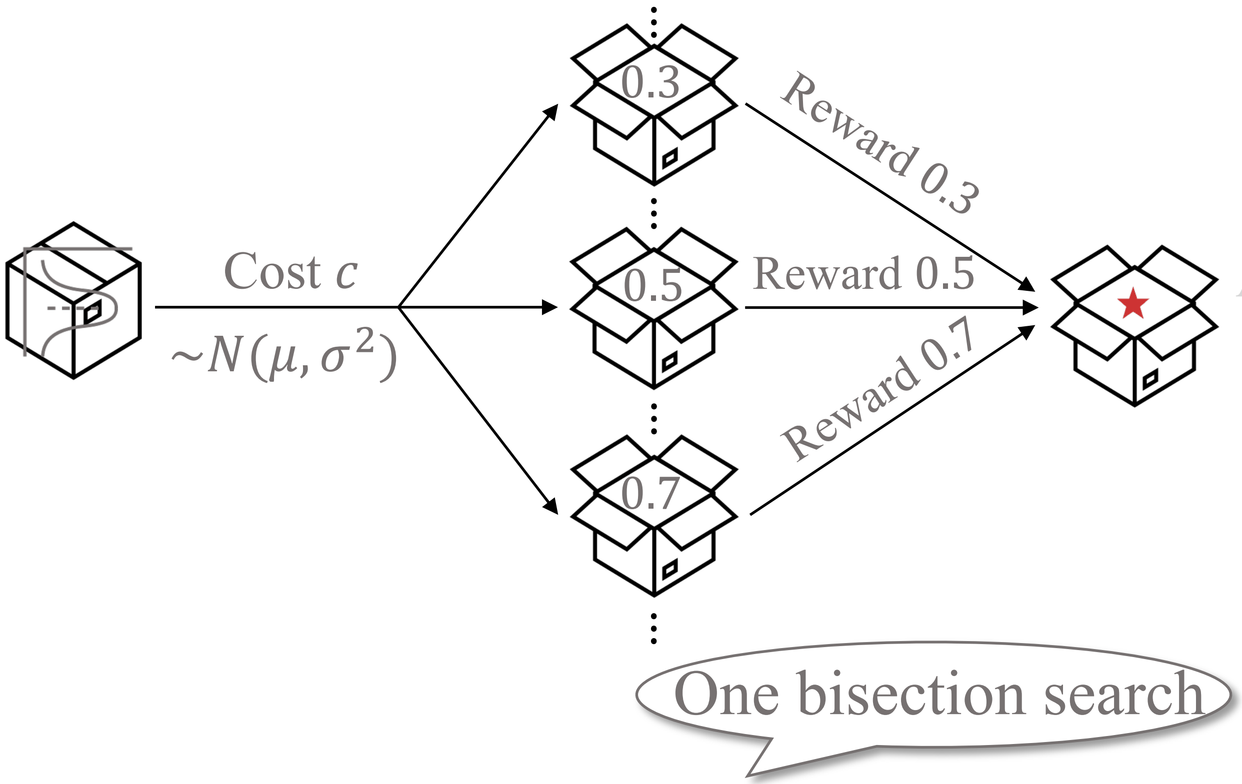

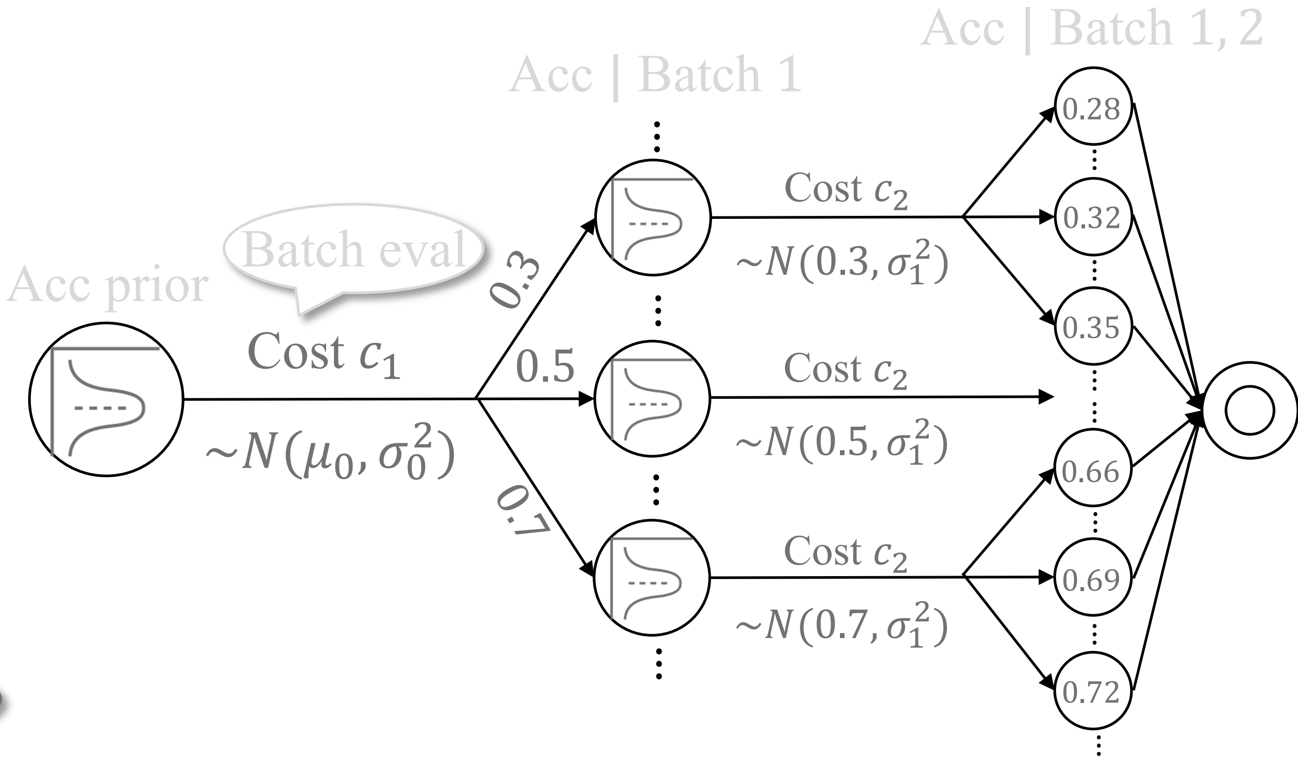

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

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

analytical expression
& monotonicity in g

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

$$\mathbb{E}[\max(\mathbb{E}[\max(f(x) - g, 0) \mid D_1] - c_2(x), 0) | D] = c_1(x)$$

Lookup table over grids for each stage

How to compute Gittins indices?

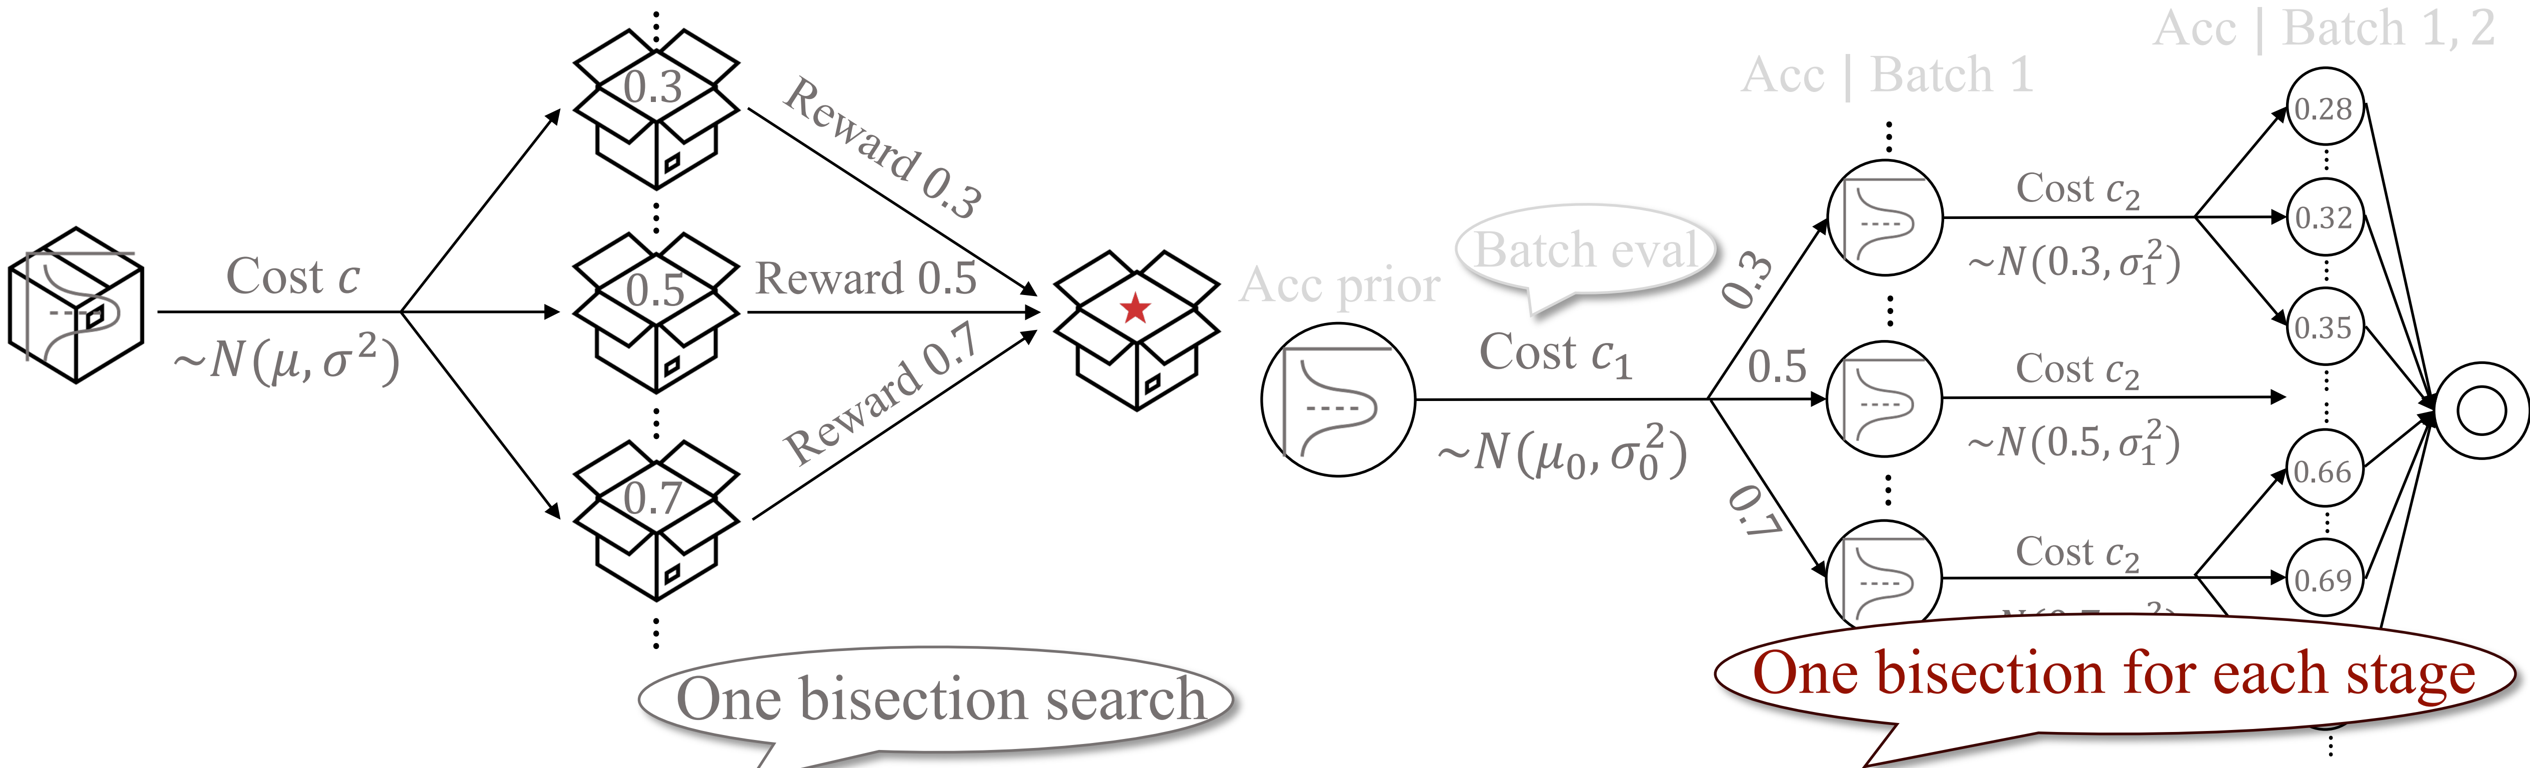

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

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

analytical expression & monotonicity in g

$\mathrm{GI}_{f|D,c}(x) :=$ solution g s.t.

$$\mathbb{E}[\max(\mathbb{E}[\max(f(x) - g, 0) \mid D_1] - c_2(x), 0) | D] = c_1(x)$$

Lookup table over grids for each stage

Choice of Prior

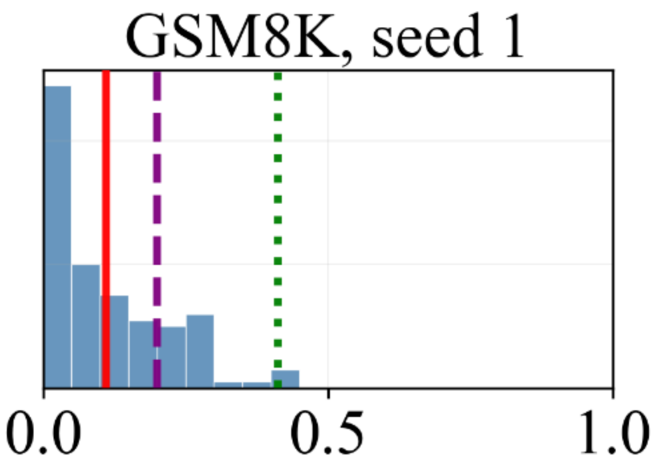

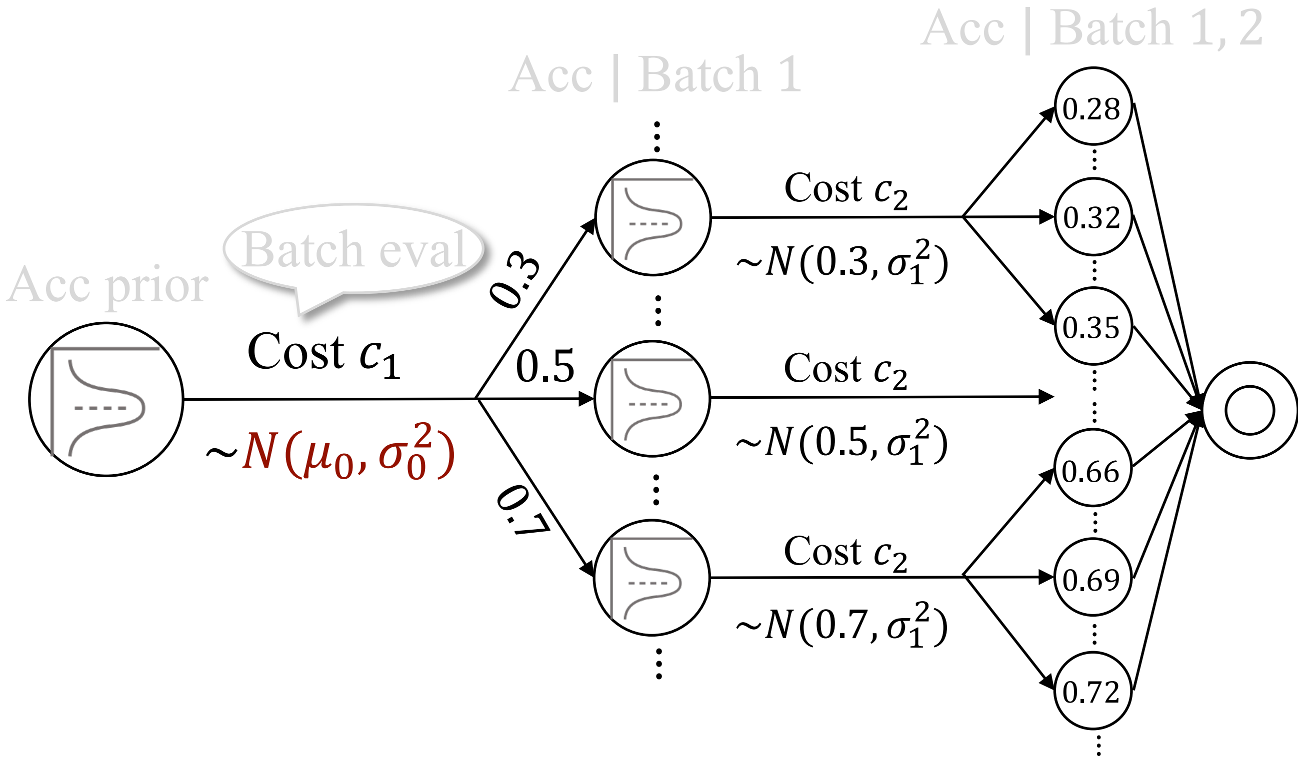

Gittins Variant	Prior	Description
Gittins-G	$N(0.5, 0.04)$	General Default
Gittins-S	e.g., $N(0.2, 0.01)$	Data-Specific

Gittins Index vs Baselines on Fewer-Candidate Tasks

LLM evaluation:

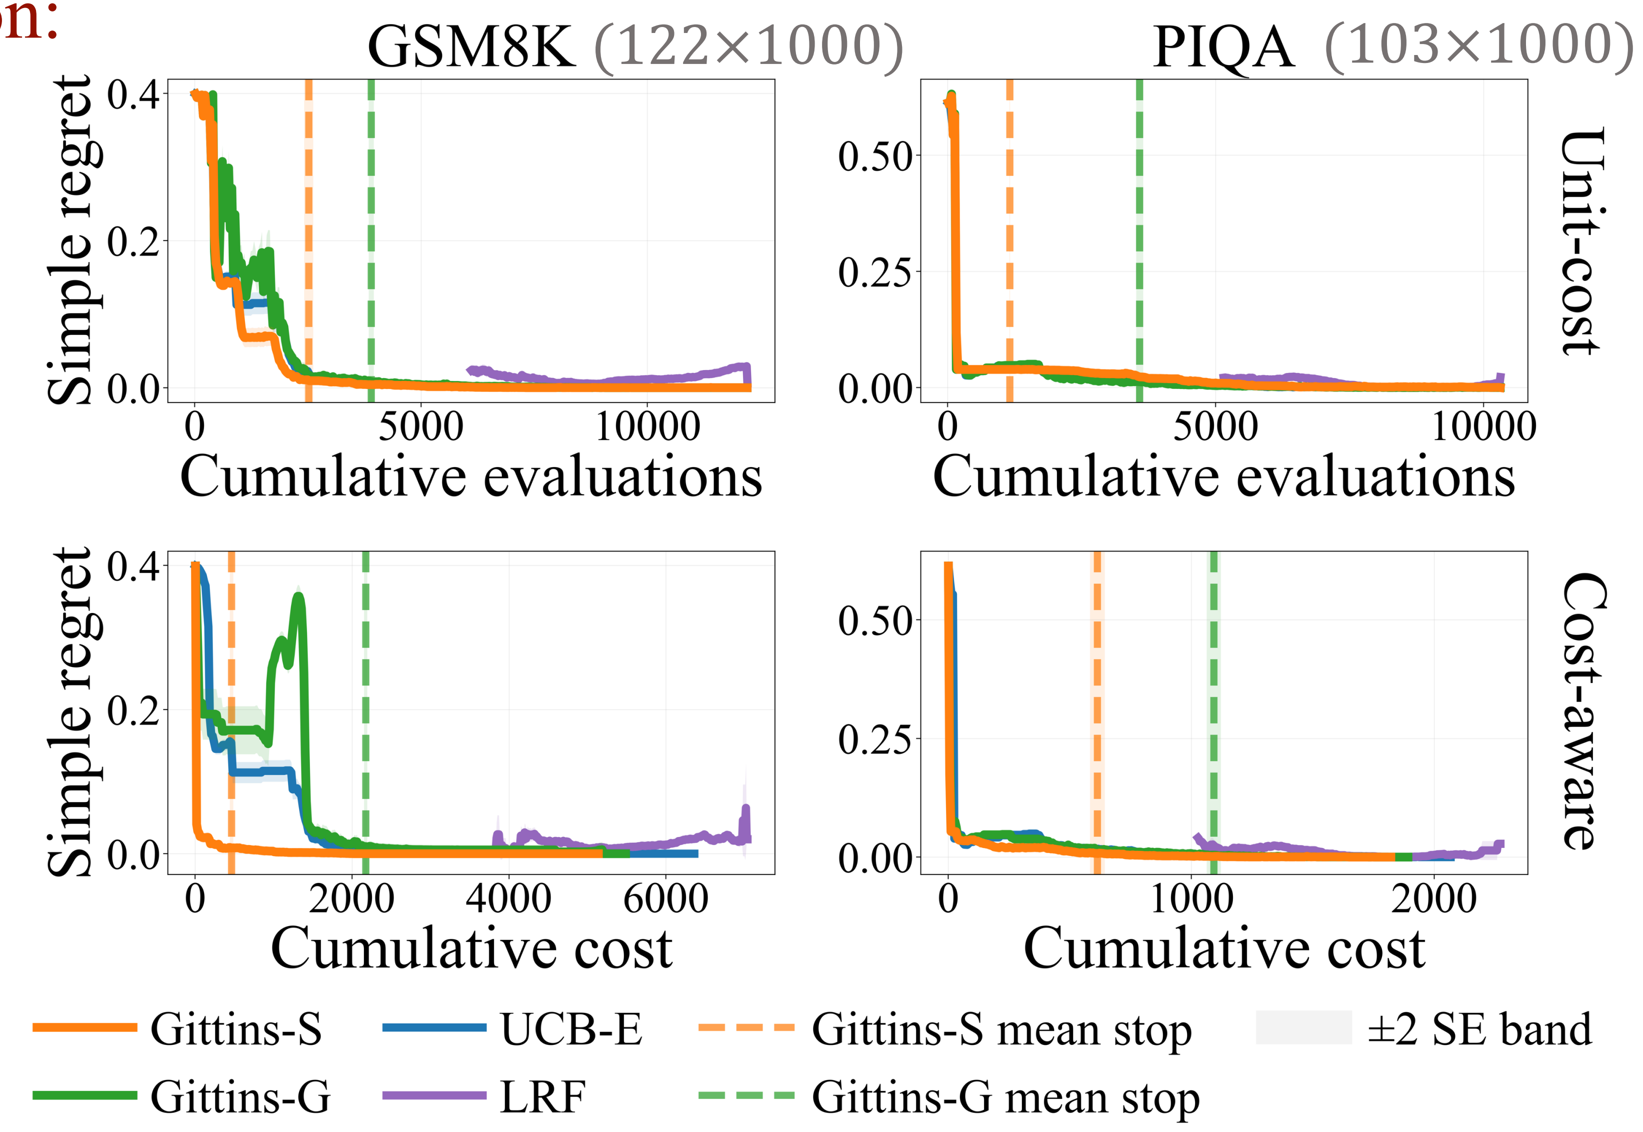

Gittins Index vs Baselines on More-Candidate Tasks

MMLU (1500 arms)

Unit-cost

Cost-aware

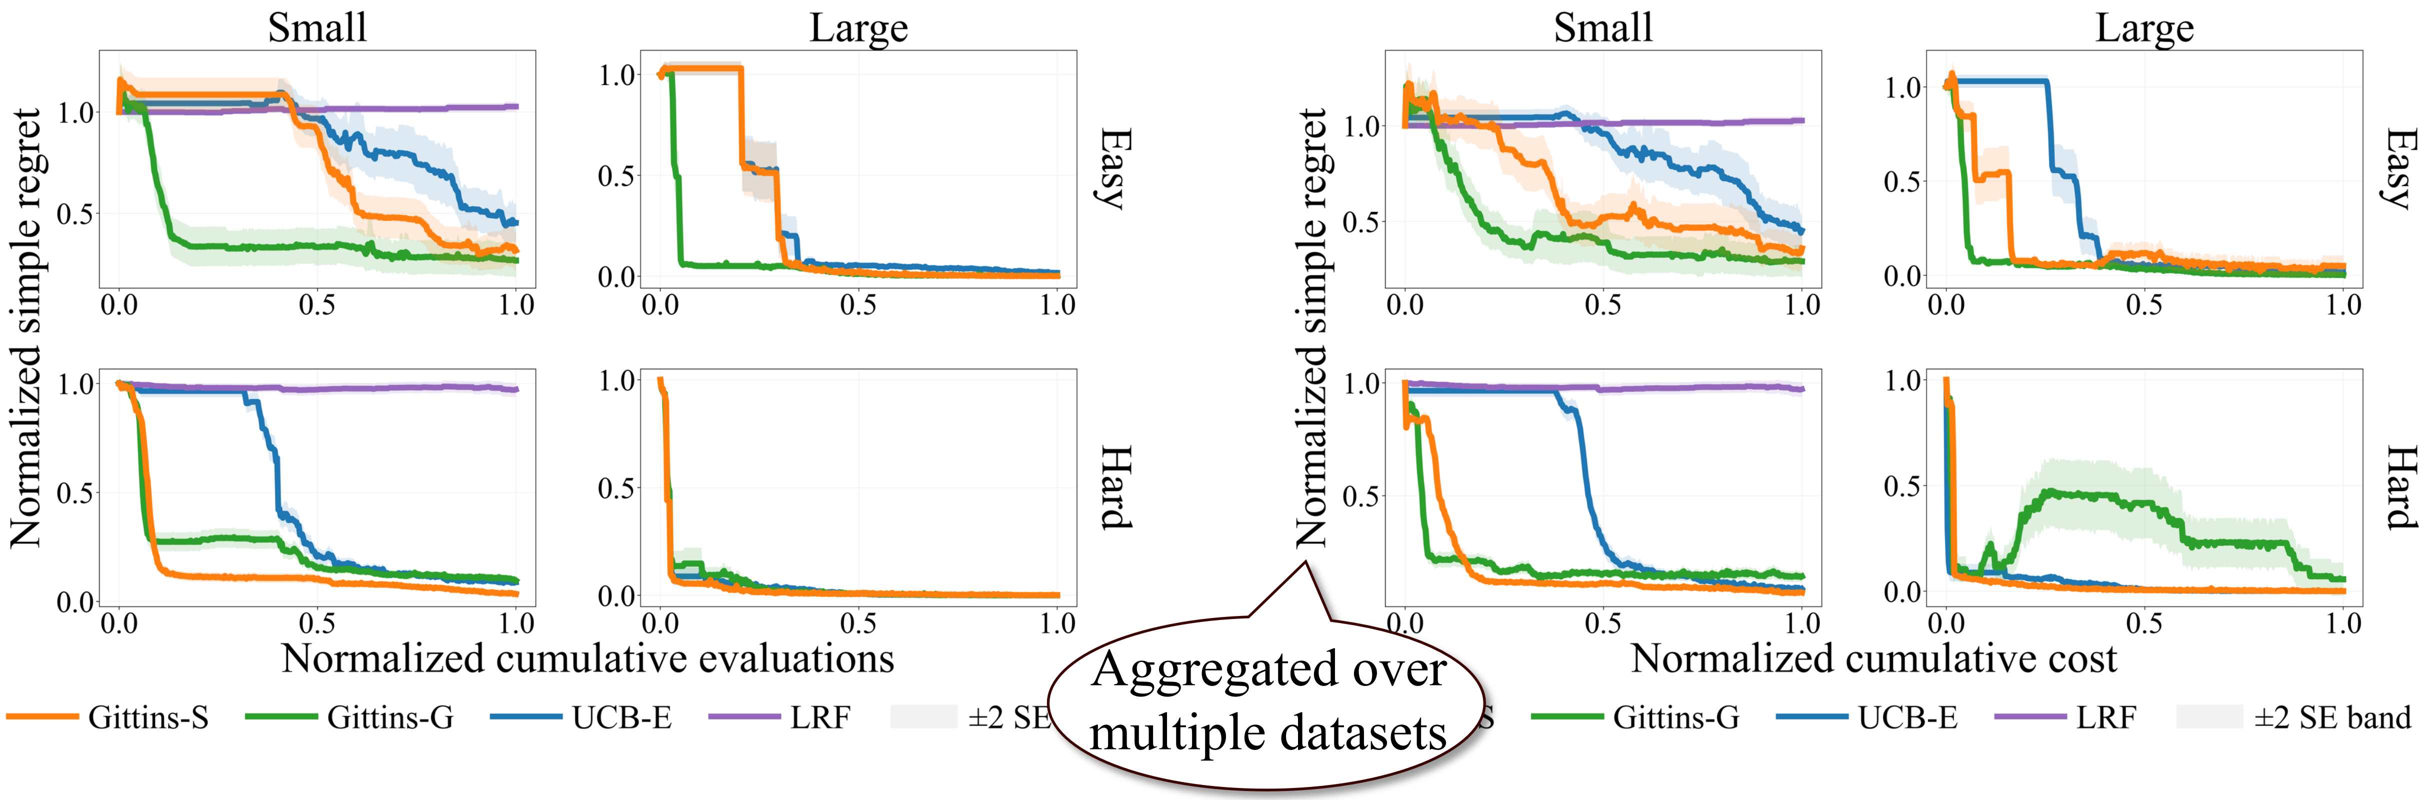

"Efficient Cost-Aware LLM Evaluation via Bayesian Bandit Gittins Indices." ICML'26 Workshop DEMO.

Key Contribution: Gittins Index Design Principle

Novel connection

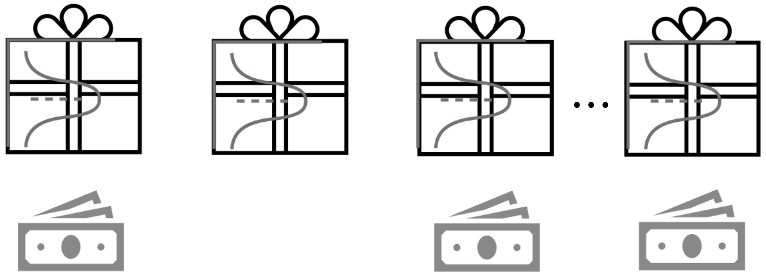

Link to Pandora's Box & Markov chain selection & Gittins index theory

Principled decision rules

- Varying evaluation costs
- Adaptive stopping time
- Multi-stage feedback

Unified framework for acquisition and stopping

Competitive empirical performance

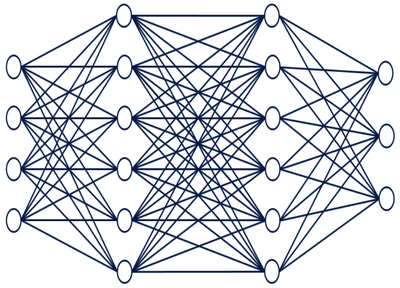

Model	Q1	Q2	Q3	...	avg performance
ChatGPT GPT-4o	✓	?	?	...	?
Claude 3.5 Sonnet	?	✓	?	...	?
Deepseek	✓	? Y_{ij}	?	...	?
Gemini 1.5 Pro	✖	?	✓	...	?
Llama 3.1-70B	?	✖	?	...	?
Mistral Large	?	?	✖	...	?
...	...	...	...	...	...

Applications to AutoML & LLM

Ongoing practical extensions

- Noisy observations
- Batch evaluations
- Arm correlations
- Trajectory-type feedback