IMPERIAL

Causal Discovery In The Age of Weak Supervision

From Preference Data to Language Model Oracles

Sonali Parbhoo

Imperial College London

MCML Workshop on Causal Machine Learning

Aug 13, 2026

In collaboration with **Jean Durand**, **Cesare Bergossi**, **Amartya Roy**, Yashas Annadani, Marcello Negri, Volker Roth, Stefan Bauer and many discussions with the AI4AI lab and others.

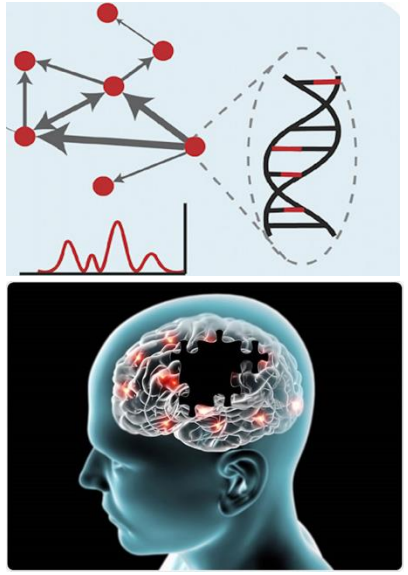
AI4AI AI for Actionable Impact

CIFAR

mcml Munich Center for Machine Learning

Motivation

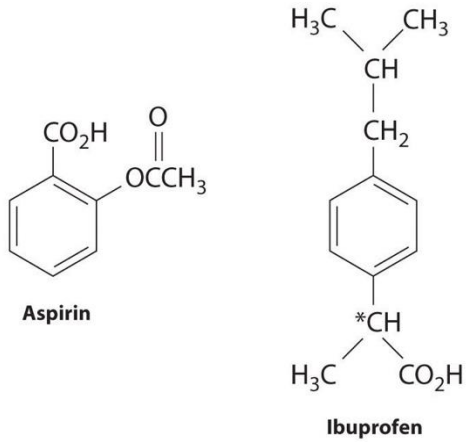
Across science, we must act on complex systems to **improve an outcome we care about**.



Which genes are responsible for Alzheimer's?

Which fertilizers will produce the highest crop yields?





Which compound has the highest efficacy?

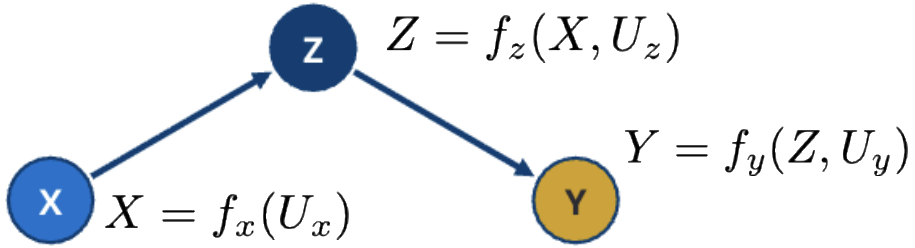
Each question asks **which intervention is best** for an outcome of interest and requires an understanding of the underlying **causal structure**.

The decision we care about

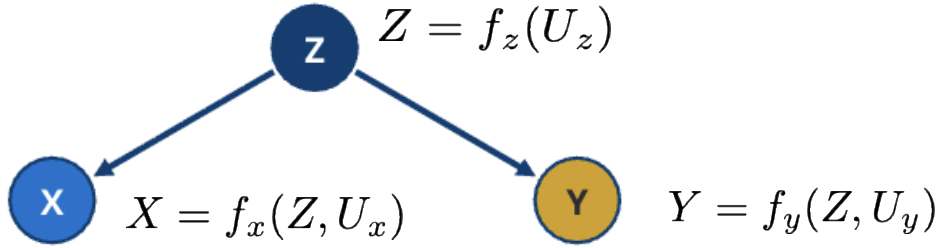
- Each of the earlier examples involves selecting an intervention a to optimize an outcome of interest *Y* based on context *C*:

$$a^* = \arg \min_{a \in P(X)} \mathbb{E}[Y|a, C, \mathcal{G}]$$

- However, the effect of an intervention **depends on the underlying causal graph**. Specifically, the same intervention may either help improve *Y* or worsen *Y* in different scenarios.



Intervening on *X* changes *Y*



Intervening on *X* has no effect on *Y*

Causal Discovery

- Goal: Learn a causal graph from data.
- **Approaches:**
  - Score-based methods
  - Constraint-based methods
  - Restricted SCMs
  - Neural methods

# of Nodes	# of DAGs	# of Graphs
1	1	1
2	3	2^2
3	25	2^6
4	543	2^{12}
5	29281	2^{20}
6	3781503	2^{30}
7	1138779265	2^{42}
8	783702329343	2^{56}
9	1213442454842881	2^{72}
10	4175098976430598143	2^{90}

"I would rather discover one causal law than be King of Persia."
- Democritus, ~400 BC.

Image: https://en.wikipedia.org/wiki/Democritus

Causal Discovery

- Goal: Learn a causal graph from data.
- **Approaches:**
  - Score-based methods
  - Constraint-based methods
  - Restricted SCMs
  - Neural methods

How to search efficiently?

- Cannot be determined from observational data alone.
- Need to know where we should intervene!
- Choose interventions that **optimize our outcomes of interest.**

# of Nodes	# of DAGs	# of Graphs
1	1	1
2	3	2^2
3	25	2^6
4	543	2^{12}
5	29281	2^{20}
6	3781503	2^{30}
7	1138779265	2^{42}
8	783702329343	2^{56}
9	1213442454842881	2^{72}
10	4175098976430598143	2^{90}

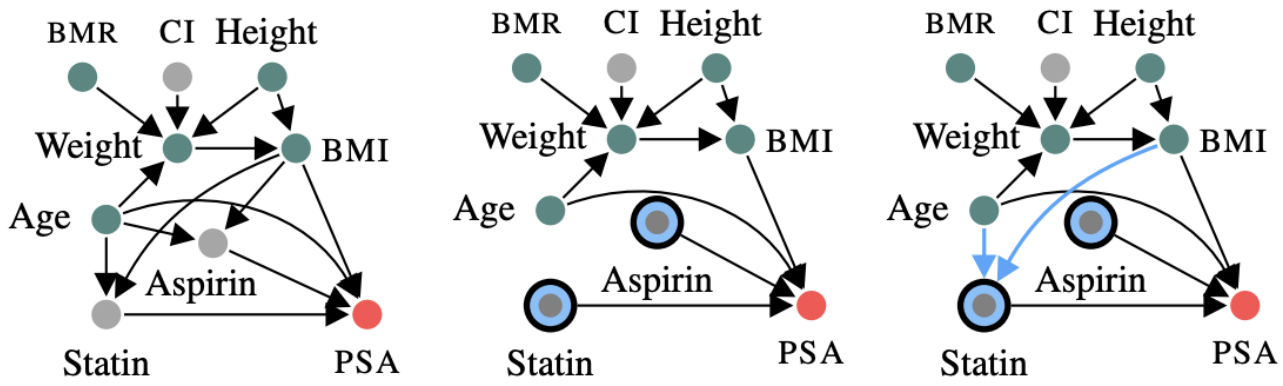
"I would rather discover one causal law than be King of Persia."
- Democritus, ~400 BC.

Image: https://en.wikipedia.org/wiki/Democritus

Causal Bayesian Optimization

- Causal Bayesian Optimization (CBO) is a technique that leverages causal information to choose interventions that optimize an outcome of interest.

E.g. A doctor can prescribe drugs to minimize prostate-specific antigen levels (PSA) in a patient if they know what variables influence PSA levels.



Causal Bayesian Optimization

- Standard CBO iterates through the following stages to solve $a^* = \arg\min_{a \in P(X)} \mathbb{E}[Y|a, C, \mathcal{G}]$
- Specifically, standard CBO uses the known graph $\mathcal{G}(f)$ to keep interventions targeted.

1. Surrogate model

$f_s(v) = \mathcal{GP}(m(v), K(v, v'))$

Model each interventional effect $\mathbb{E}[Y|a = do(X_s = v), C]$ as a GP. Its prior mean and kernel are determined from observational data.

→

2. Acquisition

$a^* = \arg\max \alpha(e), e \in ES$

Maximise expected information over the exploration set ES. Choose both the intervention and its value, trading off exploration and exploitation.

→

3. Intervene

$a = do(X_s = v) \rightarrow y$

Perform hard intervention and observe outcome *Y*, yielding a new interventional sample.

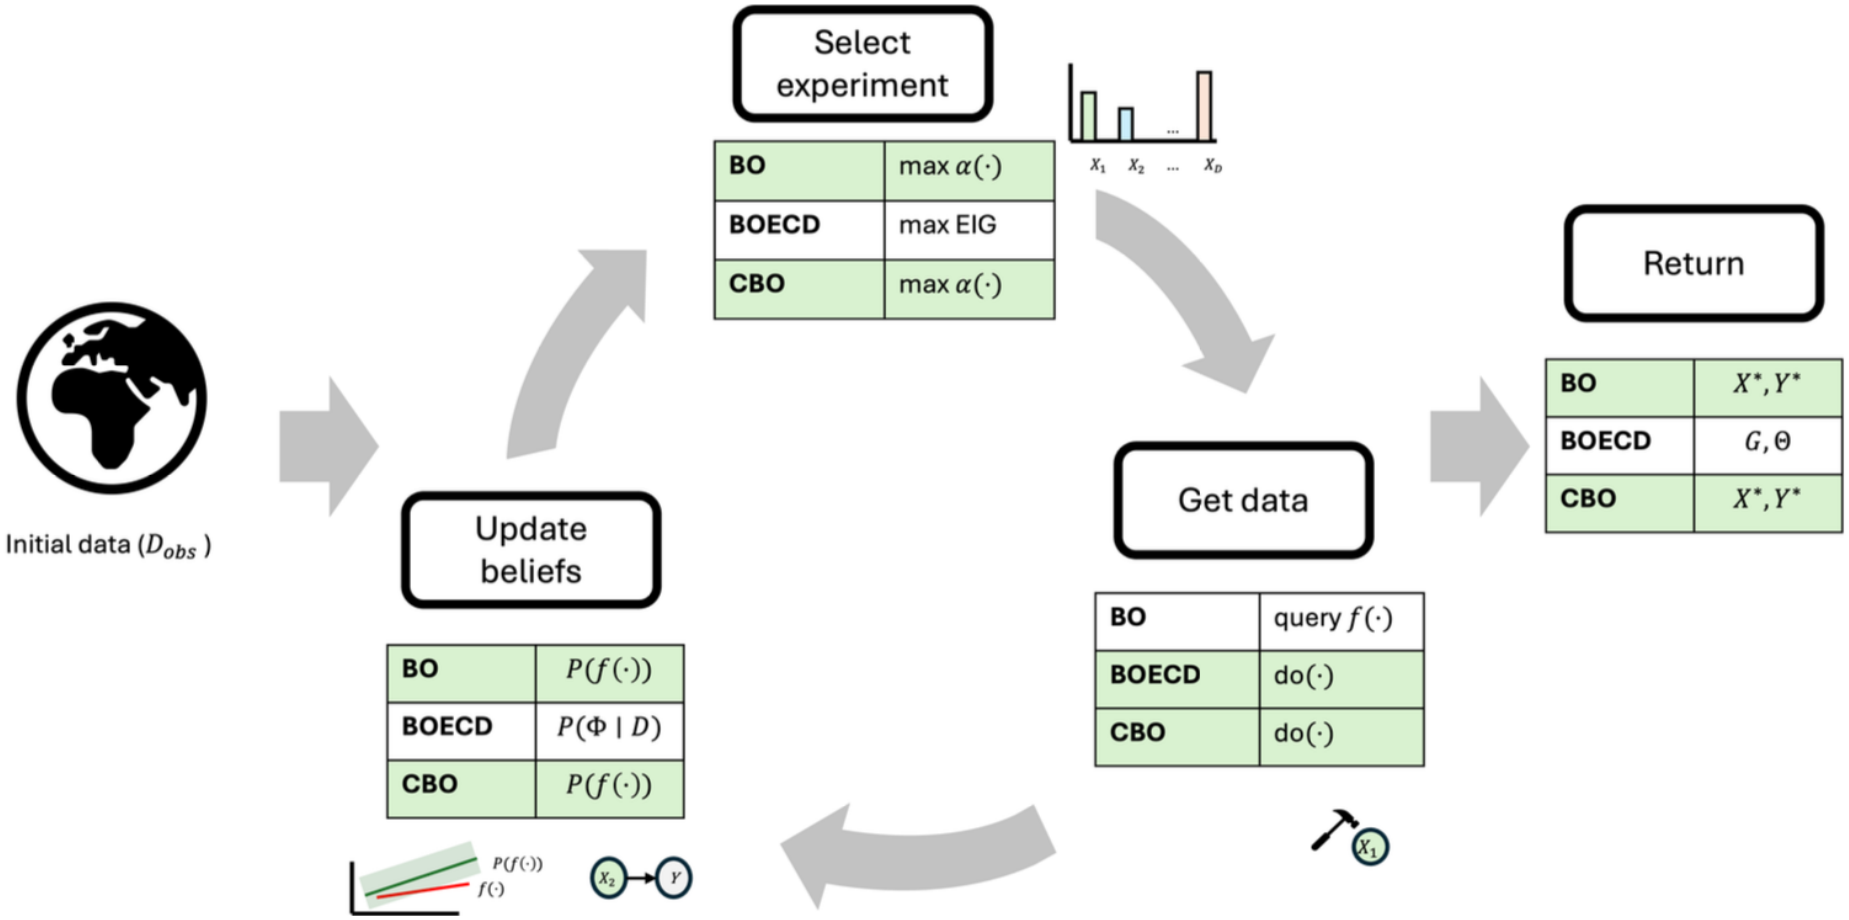
→

4. Update surrogate

$\mathcal{D} \leftarrow \mathcal{D} \cup \{(v, y)\}$

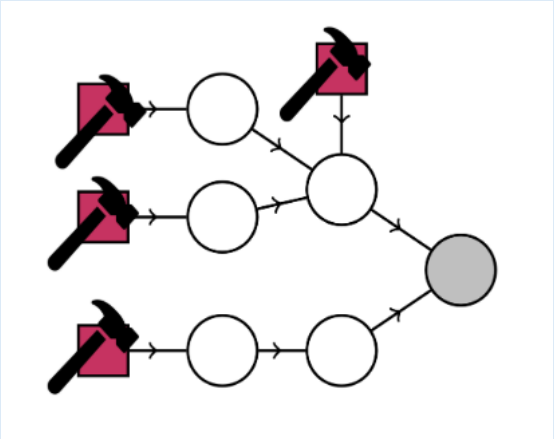
Augment interventional dataset and refit GP surrogate to improve future acquisition based on the interventional data.

Causal Bayesian Optimization



Limitations of Existing Methods for Causal Discovery

Causal Bayesian Optimization
[Aglietti et al 2020]



\+ handles hard interventions

\- assumes we know the true causal graph

In reality, most applications will not have access to all the true causal variables.

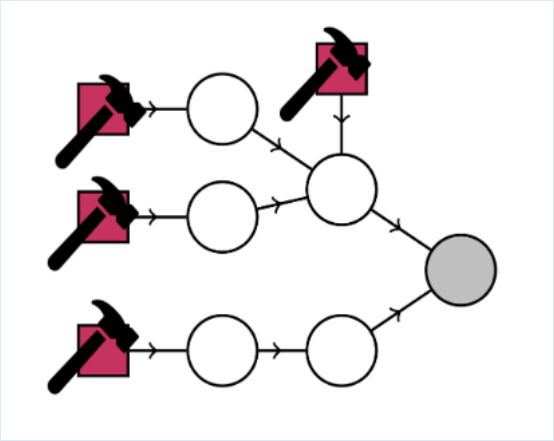
E.g. In agriculture, fertilizers may interact with many other factors such as soil properties, temperature, pests, rainfall etc. to produce a certain crop yield. It is unlikely we know all these factors.



We ask: Can we still perform CBO under <u>weak supervision</u>?

Limitations of Existing Methods for Causal Discovery

Causal Bayesian Optimization
[Aglietti et al 2020]



+ handles hard interventions

- assumes we have access to true numerical rewards

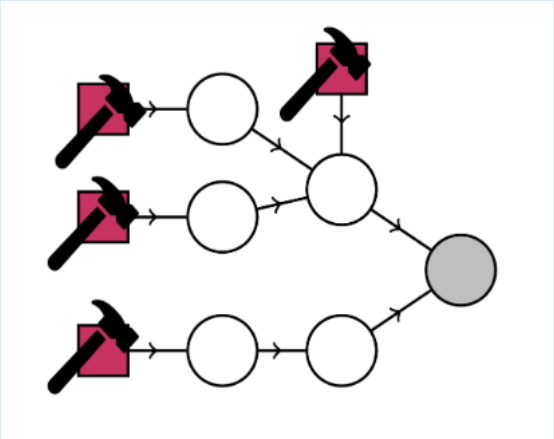
For many applications, we only observe a preference instead of a numerical outcome or reward.

E.g. In medicine, a clinician records biomarkers, physiological indicators, and treatment assignments per patient, however cannot put a scalar on a patient's overall trajectory. Instead, they might express a preference instead of a numerical outcome i.e. 'Regime A is preferrable to Regime B'.

We ask: Can we still perform CBO under <u>weak supervision</u>?

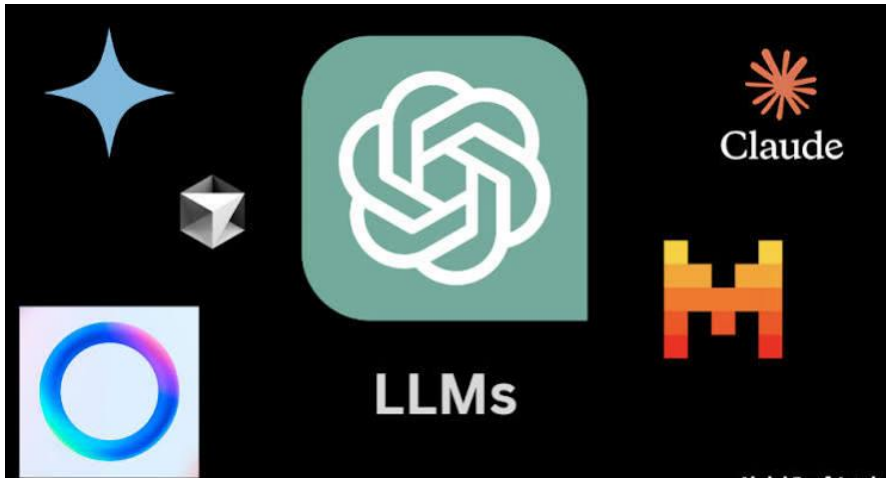
Limitations of Existing Methods for Causal Discovery

Causal Bayesian Optimization
[Aglietti et al 2020]



+ handles hard interventions
- assumes we know the true causal graph
- does not scale to larger graphs

Can we leverage the capabilities of LLMs for causal discovery if we don't know the ground truth causal graph?



We ask: Can we still perform CBO under <u>weak supervision</u>?

Causal Bayesian Optimization Under Weak Supervision

1. Unknown Graphs $\mathcal{G}(f)$

Learn a posterior distribution over the parents of outcomes *Y* and simultaneously optimize this.

Durand et al 2025

2. Unknown Outcomes

We observe preferences between outcomes instead of numerical outcomes. Optimize from preferences between outcomes instead of reward

Bergossi et al 2026

3. Language Oracles

Use an LLM as a noisy oracle over causal relations and refine these beliefs via CBO

Roy et al 2026

Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025

Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.

Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

Causal Bayesian Optimization Under Weak Supervision

1. Unknown Graphs $\mathcal{G}(f)$

Learn a posterior distribution over the parents of outcomes *Y* and simultaneously optimize this.

Durand et al 2025

2. Unknown Outcomes

We observe preferences between outcomes instead of numerical outcomes. Optimize from preferences between outcomes instead of reward

Bergossi et al 2026

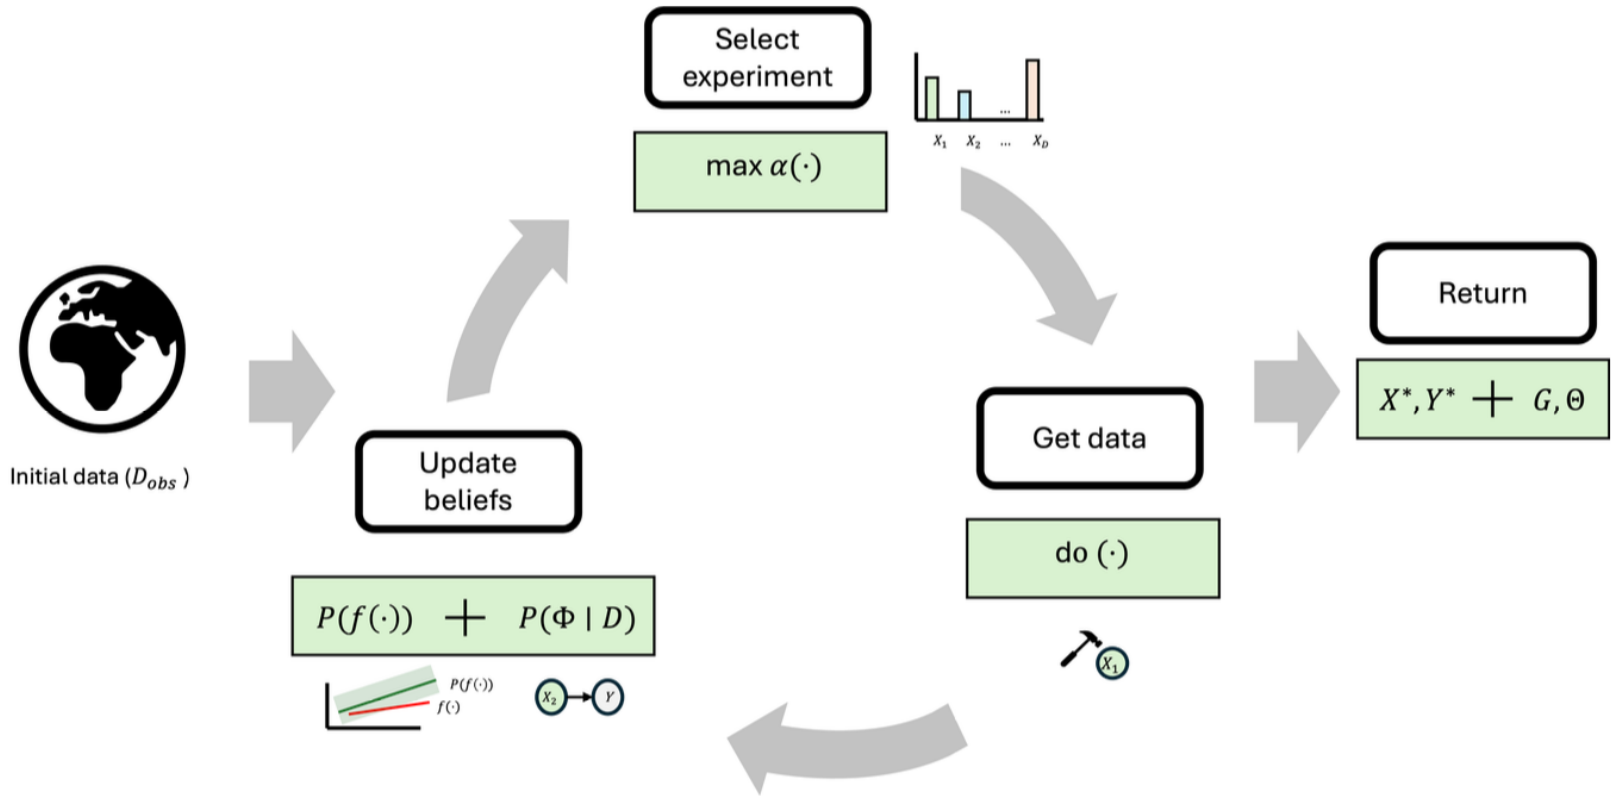
3. Language Oracles

Use an LLM as a noisy oracle over causal relations and refine these beliefs via CBO

Roy et al 2026

Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.
Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.
Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

1. Causal Bayesian Optimization with Unknown Graphs



- Idea: Maintain a posterior distribution over causal effect **as well as over causal parents of Y.**

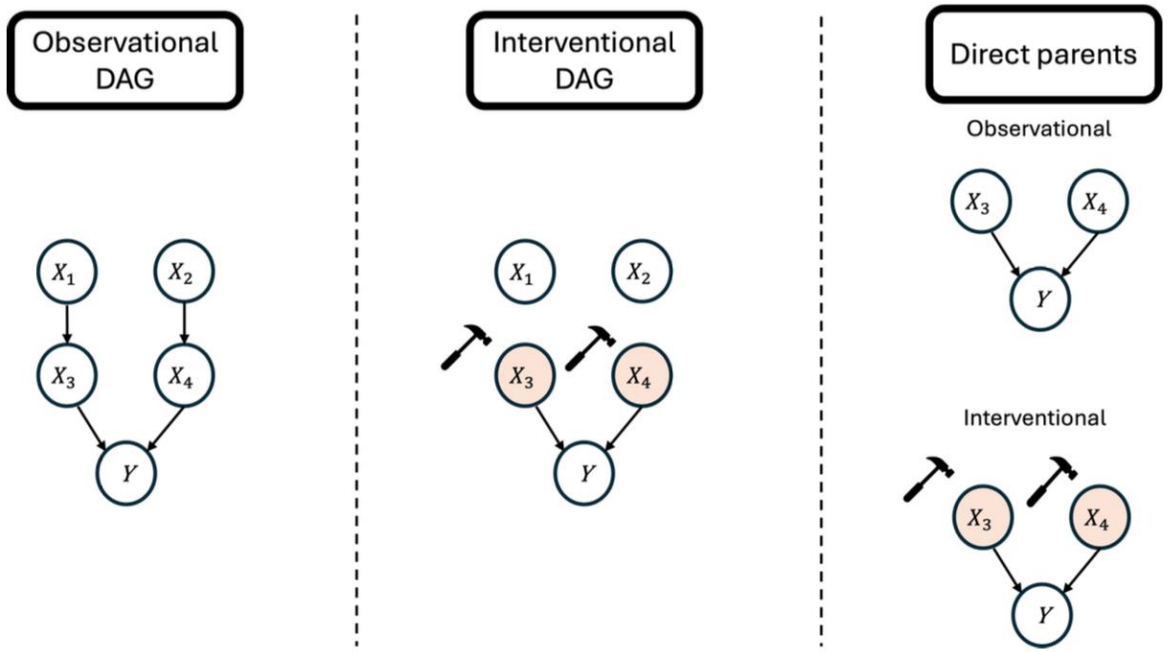
Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

Key Insight: Learning the causal parents of *Y* alone suffices to optimize *Y*.

- We do not need to learn a posterior distribution over the entire causal graph $\mathcal{G}(f)$, but maintain a posterior distribution over **only the parents of *Y*.**

Assuming a Gaussian Additive Noise Model, we are interested in the causal mechanism of target *Y:*

$$Y := f(\mathbf{pa}_Y) + \varepsilon_Y \text{ where } \varepsilon_Y \sim \mathcal{N}(0, \sigma_Y^2)$$



Both interventional DAGs are **equivalent**

Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

Learning a posterior distribution over causal parents of *Y*

$$Y := f(\mathbf{pa}_Y) + \varepsilon_Y \text{ where } \varepsilon_Y \sim \mathcal{N}(0, \sigma_Y^2)$$

The likelihood is given by $P(\boldsymbol{x}, y \mid \boldsymbol{g}, \boldsymbol{\theta}_Y) = P_{\varepsilon_Y}(y - f(\boldsymbol{x}_{\mathrm{pa}_Y})) \equiv \mathcal{N}(y - f(\boldsymbol{x}_{\mathrm{pa}_Y}); 0, \sigma_Y^2)$

Linear case

$$Y := \boldsymbol{\theta}_s^\top \boldsymbol{X}_s + \varepsilon_Y \text{ where } \varepsilon_Y \sim \mathcal{N}(0, \sigma_Y^2)$$

The probability that $\boldsymbol{g}_s$ is the parent set of Y is calculated as follow

$$P(\boldsymbol{g}_s \mid \boldsymbol{x}, y) = \int_{\boldsymbol{\theta}_s} \frac{\overbrace{P_{\varepsilon_y}(y - \boldsymbol{\theta}_s^\top \phi(\boldsymbol{x}_s))}^{\text{likelihood}} \overbrace{P(\boldsymbol{\theta}_s \mid \boldsymbol{g}_s) P(\boldsymbol{g}_s)}^{\text{priors}}}{\underbrace{P(\boldsymbol{x}, y)}_{\text{marginal}}} \mathrm{d}\boldsymbol{\theta}_s$$

Non-linear case

$$Y := \boldsymbol{\theta}_s^\top \phi(\boldsymbol{X}_s) + \varepsilon_Y \text{ where } \varepsilon_Y \sim \mathcal{N}(0, \sigma_Y^2)$$

Suppose that

$$f(\boldsymbol{x}_s) \sim \mathcal{GP}(\mathbf{0}, K_{\mathrm{RBF}}(\boldsymbol{x}_s, \boldsymbol{x}_s'))$$

Approximate $\phi(\boldsymbol{X}_s)$ using **Fourier Transforms**

Additional theoretical results show that this can be done at scale using a **double machine learning** approach.

Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

Causal Bayesian Optimization with Unknown Graphs (CBO-U): Methodology

Problem Setup

$$a^* = \arg \min_{a \in P(X)} \mathbb{E}[Y|a, C, \mathcal{G}]$$

Optimize target Y and learn the causal graph

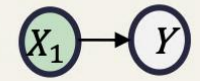
Surrogate Model

$$\mathbb{E}[Y|a, C] = \mathbb{E}_{P(g)}[\mathbb{E}[Y|a, C, g]]$$
$$\mathbb{V}[Y|a, C] = \mathbb{V}_{P(g)}[\mathbb{E}[Y|a, C, g]] + \mathbb{E}_{P(g)}[\mathbb{V}[Y|a, C, g]]$$

Introduce a latent variable g to capture uncertainty over the graph structure as well

Algorithm

1. Get initial probabilities for $P(\boldsymbol{g})$ using the **observational data** and **double machine learning**



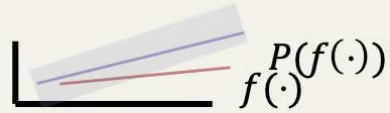
2. Use acquisition function to **select intervention** and **intervene** on the system



3. Update $P(\boldsymbol{g})$ using derived formula and $\mathrm{P}(\boldsymbol{\theta})$ using **Bayesian Linear Regression**



4. Update **GP** $P(f(\cdot))$ using the new interventional data point

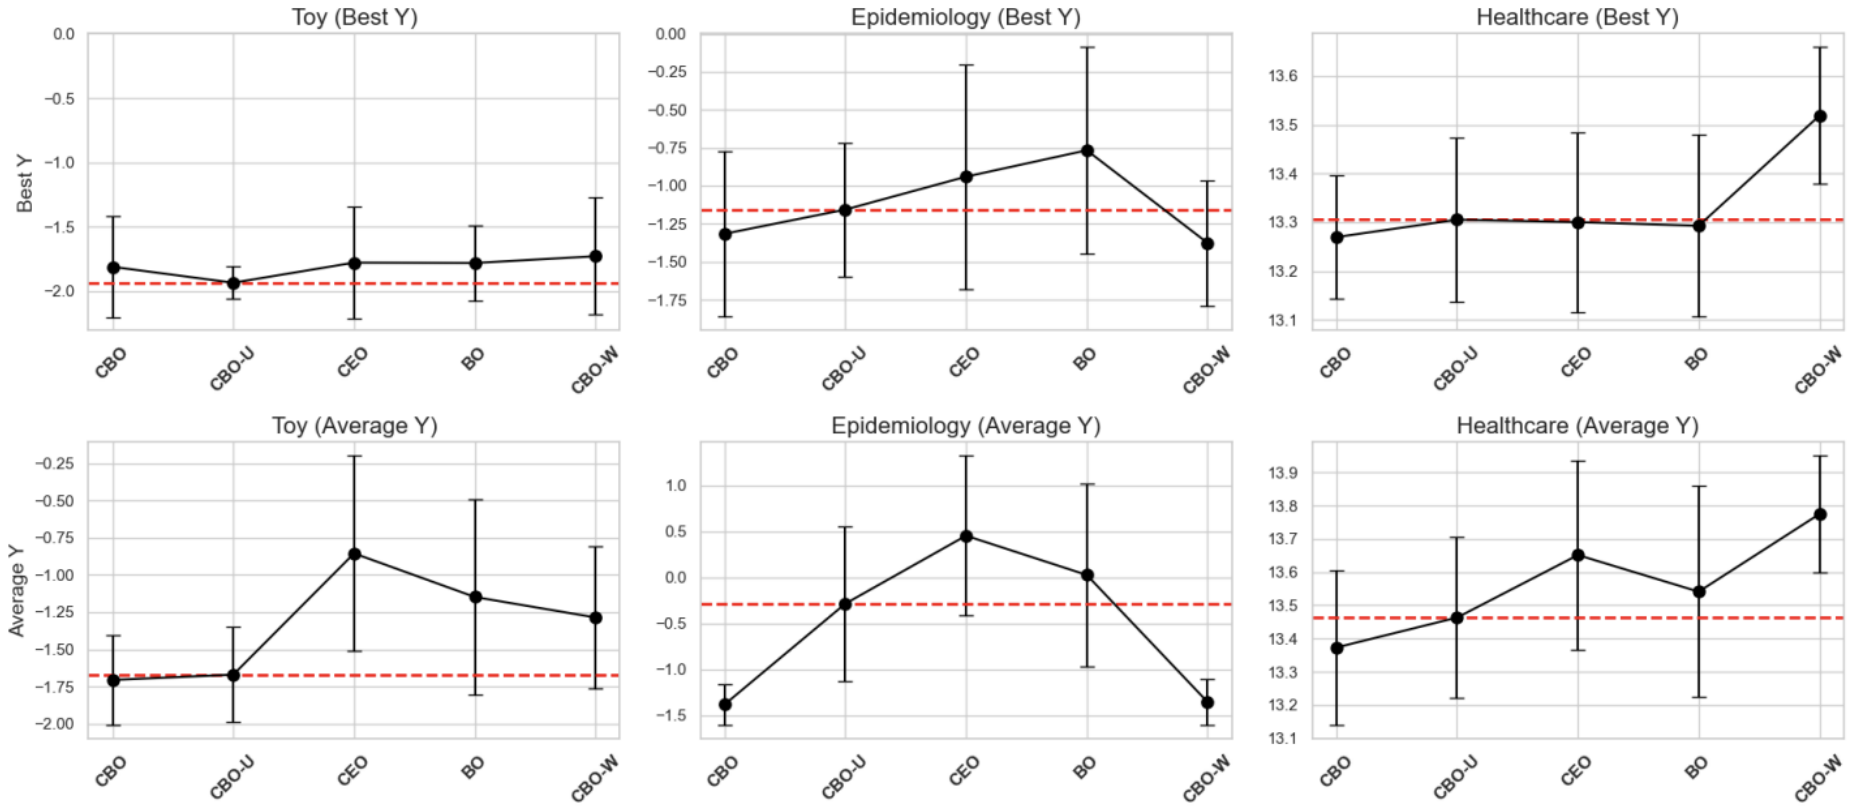


5. Return **best input** in the data set



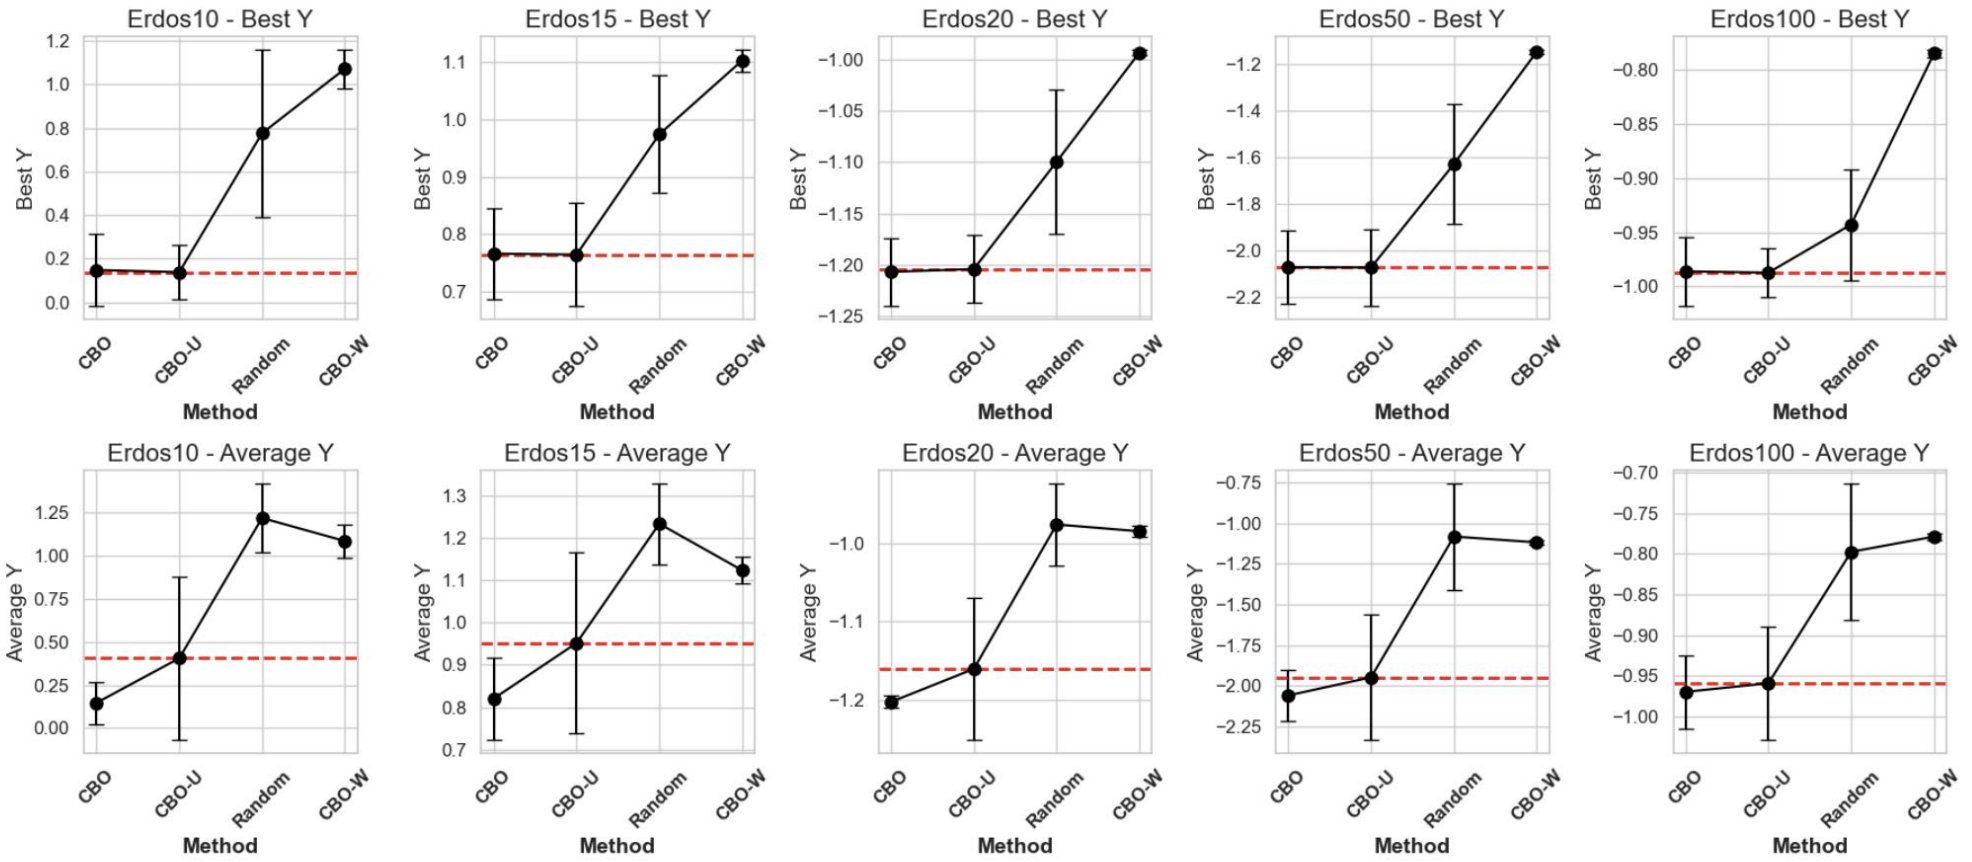
Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

CBO-U converges to the optimal target



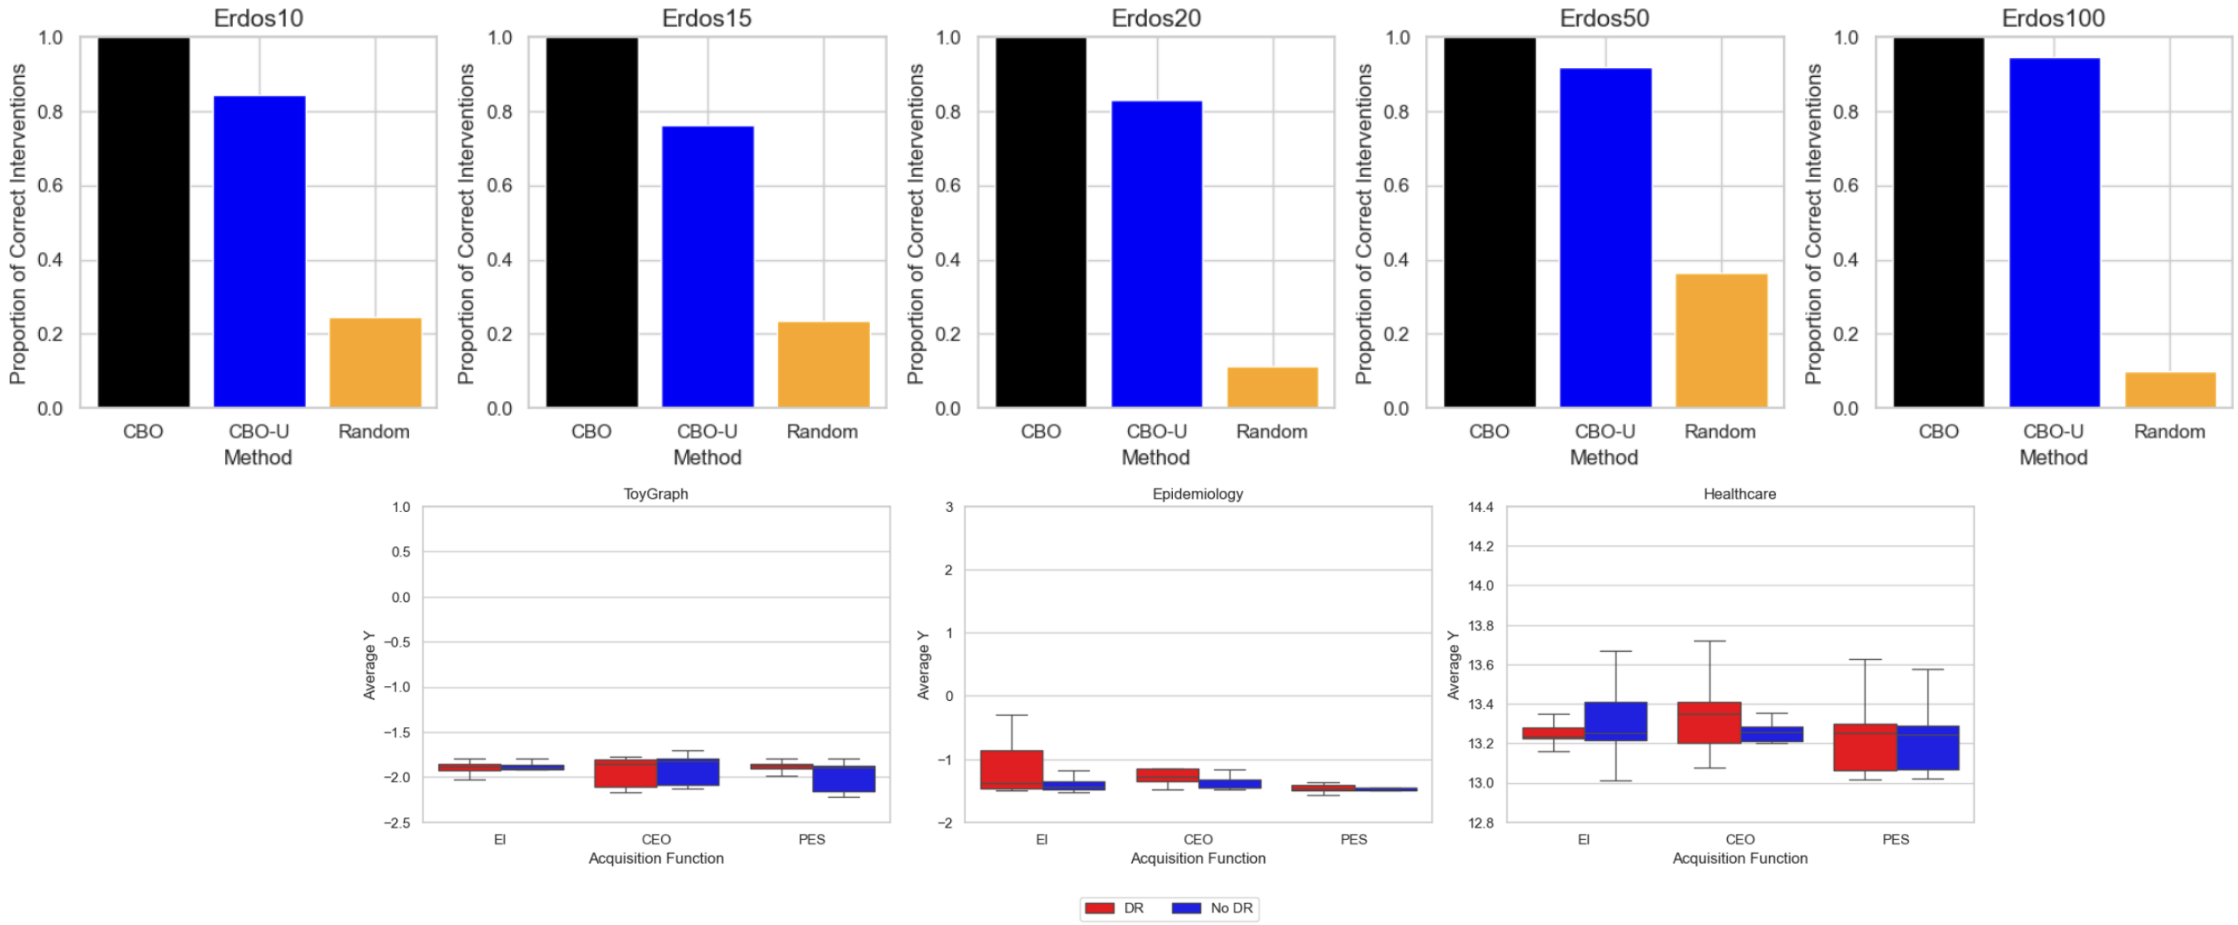
Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

CBO-U scales to larger graphs



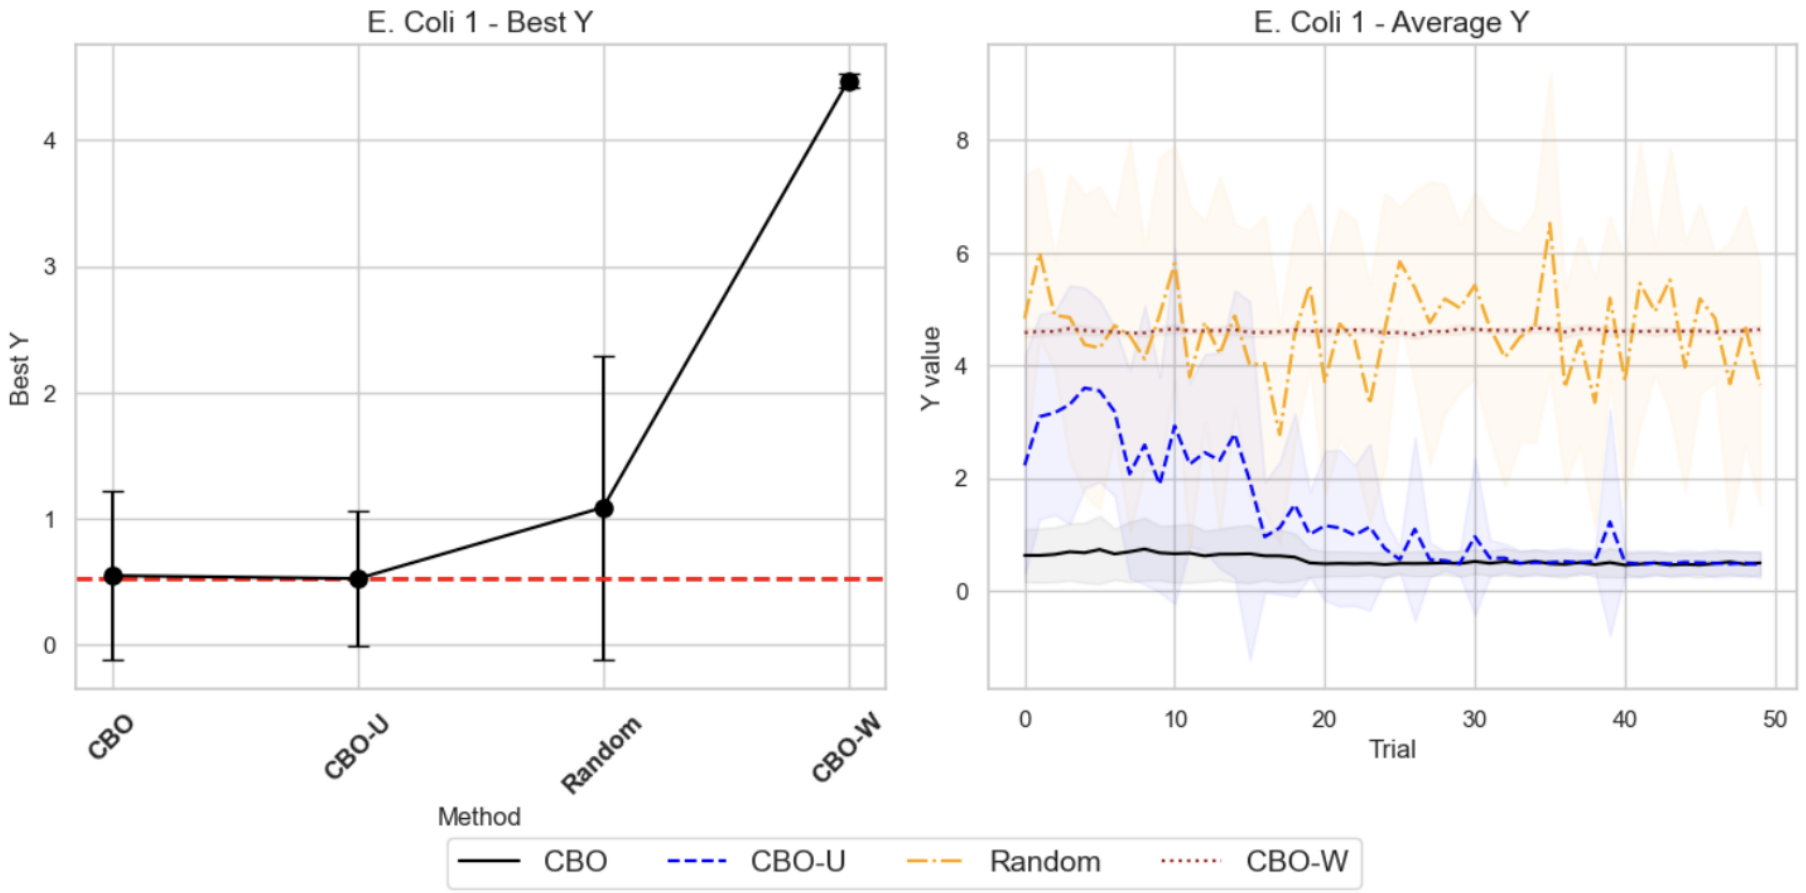
Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

CBO-U chooses the optimal intervention, in spite of weaker supervision



Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

CBO-U translates to real graphs



Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.

Causal Bayesian Optimization Under Weak Supervision

1. Unknown Graphs $\mathcal{G}(f)$

Learn a posterior distribution over the parents of outcomes *Y* and simultaneously optimize this.

Durand et al 2025

2. Unknown Outcomes

We observe preferences between outcomes instead of numerical outcomes. Optimize from preferences between outcomes instead of reward

Bergossi et al 2026

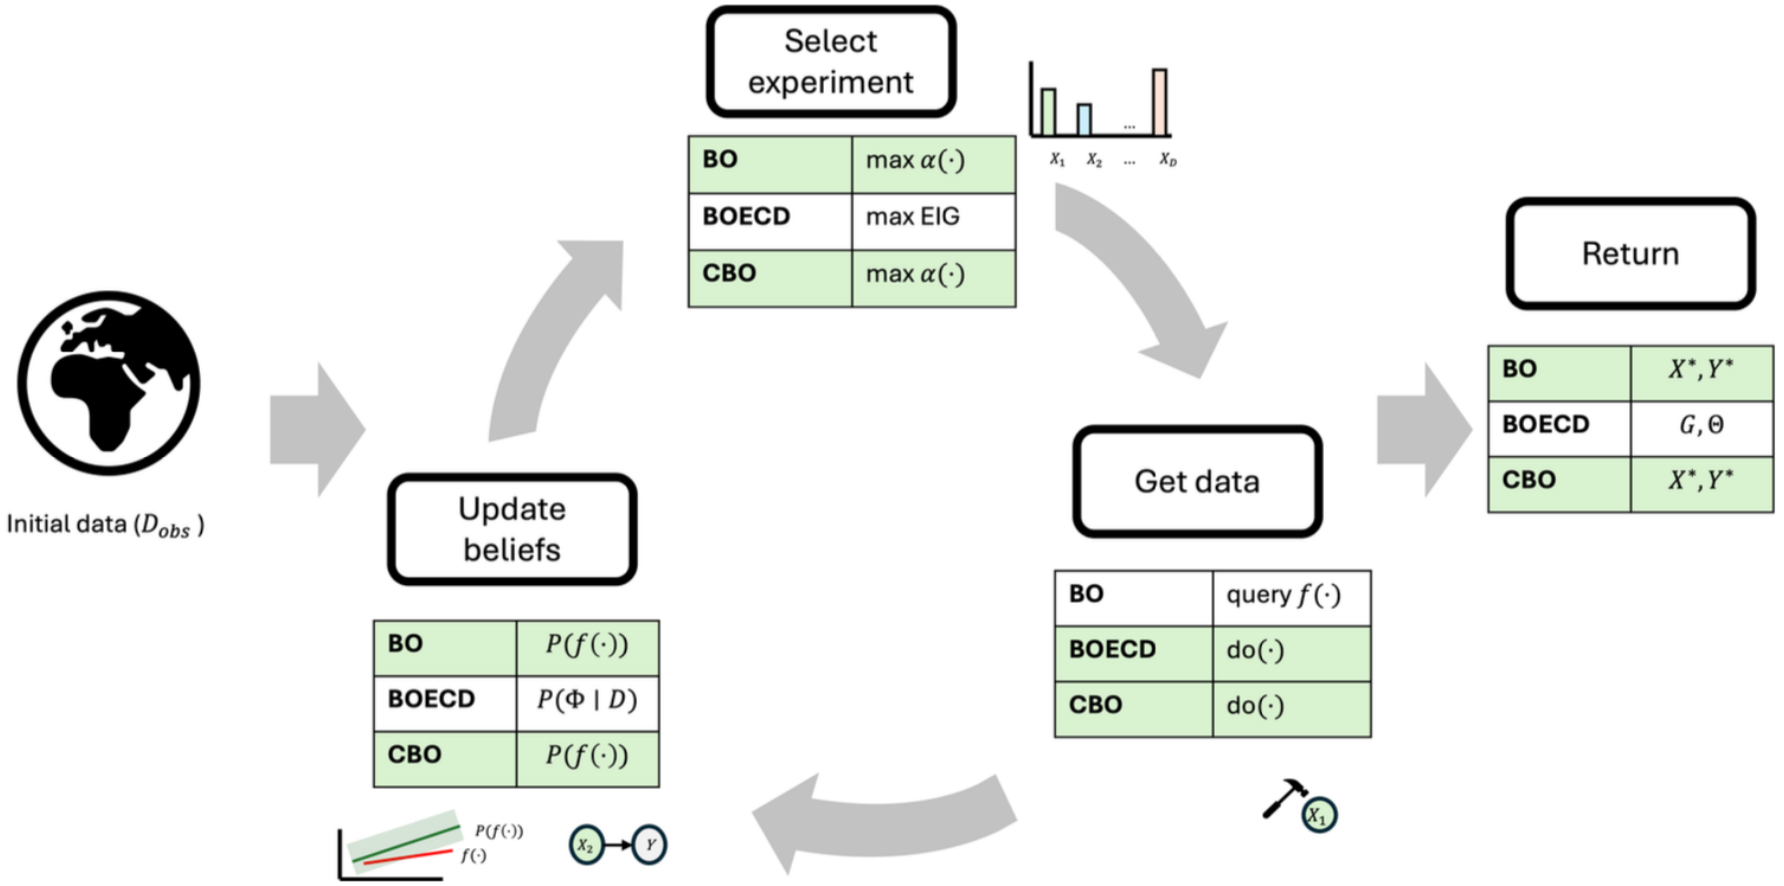
3. Language Oracles

Use an LLM as a noisy oracle over causal relations and refine these beliefs via CBO

Roy et al 2026

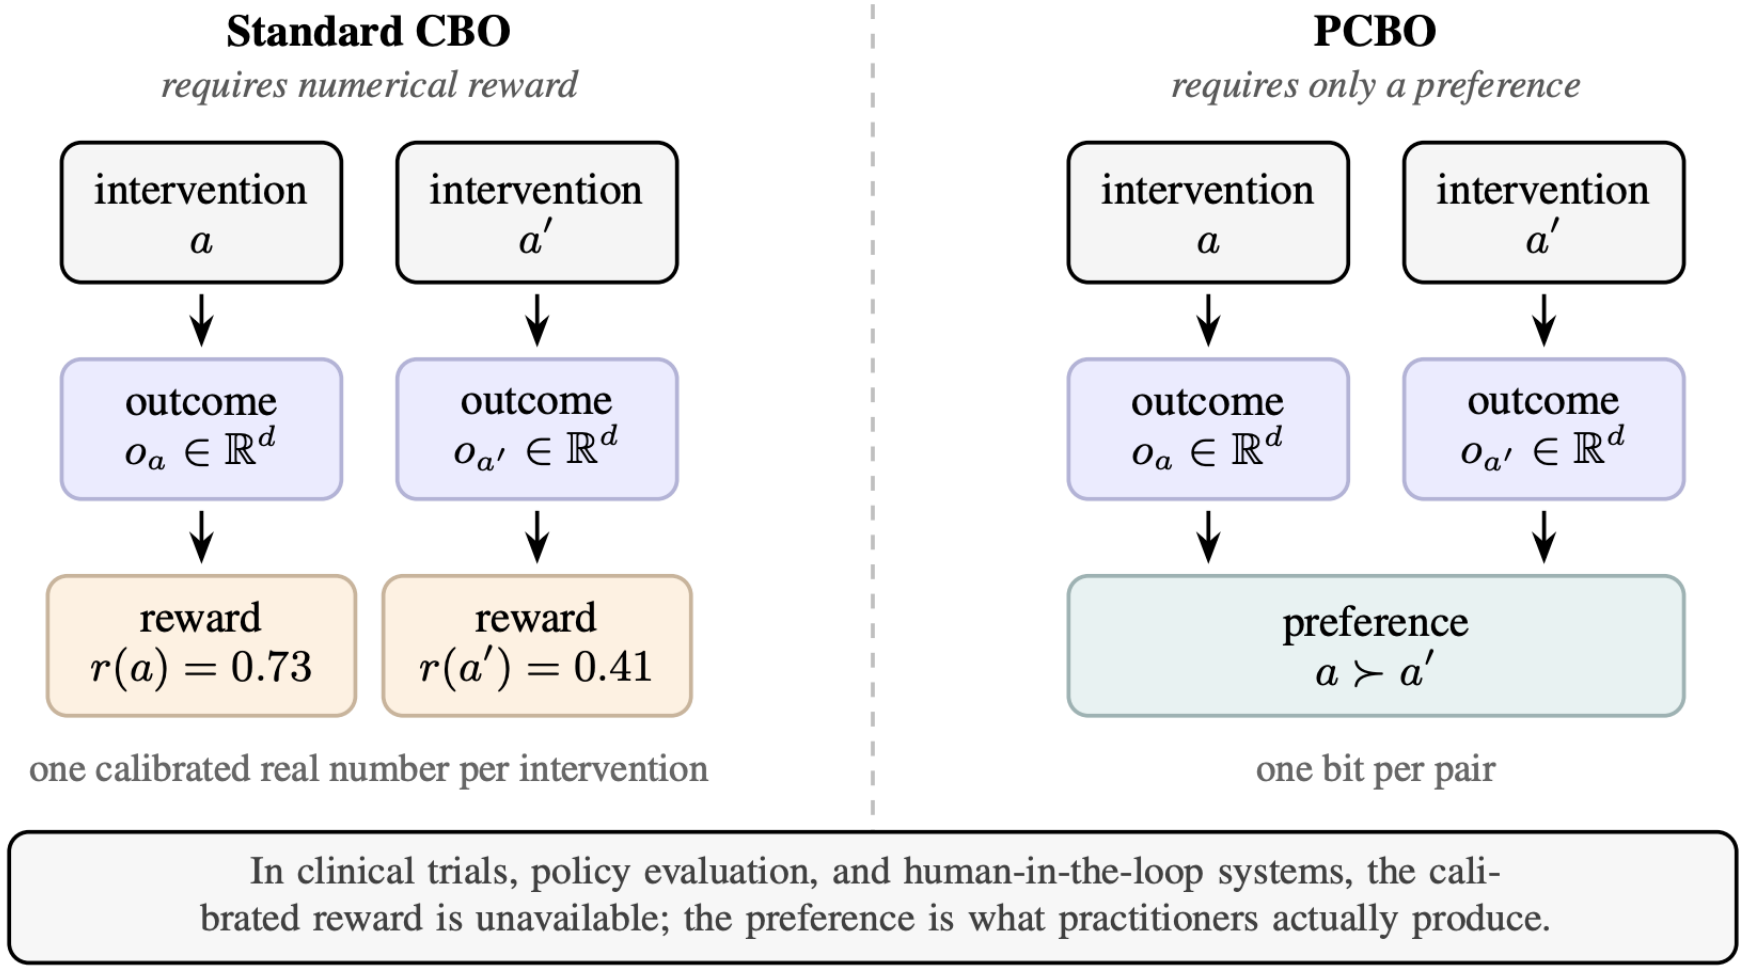
Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.
Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.
Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

2. Unknown Outcomes: Active Causal Discovery Under Ordinal Feedback



Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.

2. Unknown Outcomes: Active Causal Discovery Under Ordinal Feedback



In clinical trials, policy evaluation, and human-in-the-loop systems, the calibrated reward is unavailable; the preference is what practitioners actually produce.

Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.

Key Insight: Intervention Selection Only Requires Ordinal Feedback

- CBO-U performs two tasks: learns the causal parents of a node and selects optimal interventions a^*.
  - Learning the causal parents of a node requires the full outcome vector.
  - Choosing the next intervention only requires knowing **which intervention is better via a preference.**

$$p(a \succ a' \mid o_a, o_{a'}) = \sigma\left(\frac{\psi(o_a) - \psi(o_{a'})}{s}\right),$$

  - Importantly, utility ψ is never observed directly. We only observe the rankings of a and a'.

Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.

How do we model the utility based on the preference-data?

Given

We observe pairwise rankings $a > a'$ but not ψ. Assume the pairwise rankings come from a random utility model

$$p(a \succ a' \mid o_a, o_{a'}) = \sigma\left(\frac{\psi(o_a) - \psi(o_{a'})}{s}\right),$$

Modelling ψ via preferential normalizing flows

Recover ψ indirectly using normalizing flows. That is, fit a normalizing flow $u \approx \psi$ whose preferences match the observed pairwise rankings

Why use normalizing flows instead of a GP?

- GPs implicitly encode a smoothness assumption i.e if two outcomes o_a and $o_{a'}$ are similar, their utilities are similar too.
- Yet two interventions can **produce numerically similar outcomes with very different utilities**.
- E.g. Two patients with similar blood pressure and oxygen saturation. One may be on heavy vasopressor support, while the other may be stable.
- Normalizing flows do not assume such smoothness.

Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.

Preference-Based CBO (PCBO): Methodology

Problem Setup

Given a budget of *T* interventions, jointly recover $p(\mathcal{G}(f)|\mathcal{D})$ and optimal interventions *a**

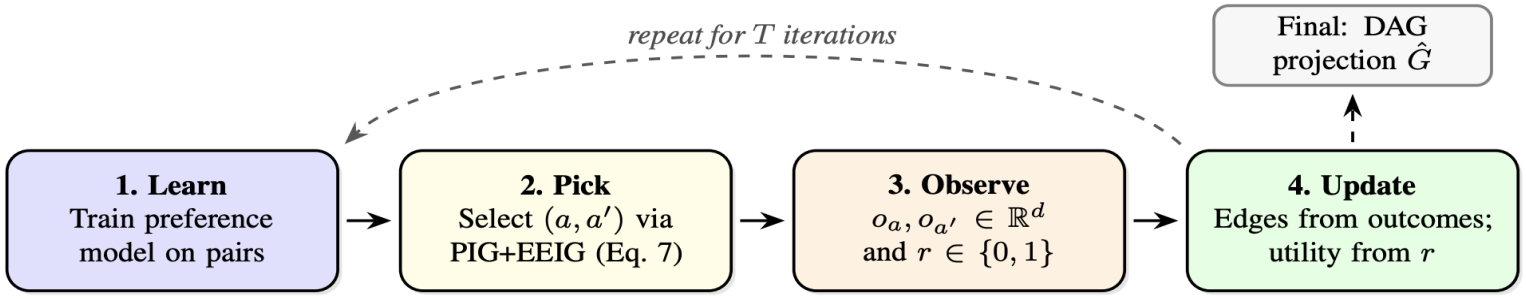
Structure Modelling

From outcome vectors o_a and $o_{a'}$, compute edge posteriors via Bayes'.

Intervention Selection via Utility

Based on preferences *a*> *a'*. Use a preferential normalizing flow to compute $u(o) = \log p_\phi(h(o))$

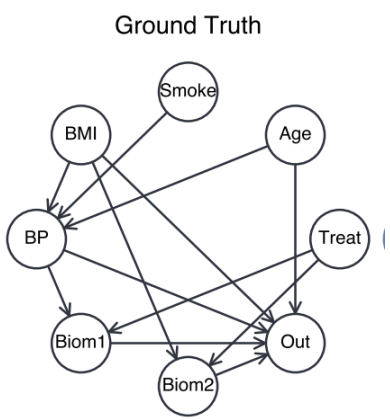
Algorithm

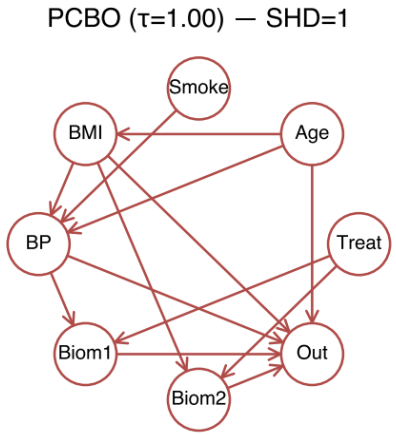


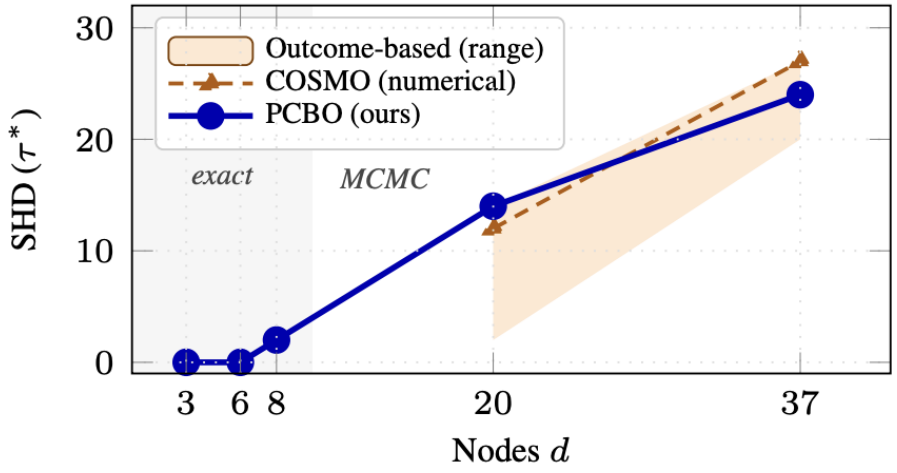
Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.

PCBO recovers graphs accurately in spite of weaker supervision

Graph	Nodes	Edges	Prec.	Recall	F1	SHD(τ^*)	SHD($\tau \approx 1$)
Synthetic-3	3	2	1.00	1.00	1.00	0	0
Synthetic-6	6	6	1.00	1.00	1.00	0	0
Medical	8	12	0.92	1.00	0.96	2	1
CHILD	20	25	0.78	0.84	0.81	14	16
ALARM	37	46	0.71	0.76	0.73	24	28







Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.

Causal Bayesian Optimization Under Weak Supervision

1. Unknown Graphs $\mathcal{G}(f)$

Learn a posterior distribution over the parents of outcomes *Y* and simultaneously optimize this.

Durand et al 2025

2. Unknown Outcomes

We observe preferences between outcomes instead of numerical outcomes. Optimize from preferences between outcomes instead of reward

Bergossi et al 2026

3. Language Oracles

Use an LLM as a noisy oracle over causal relations and refine these beliefs via CBO

Roy et al 2026

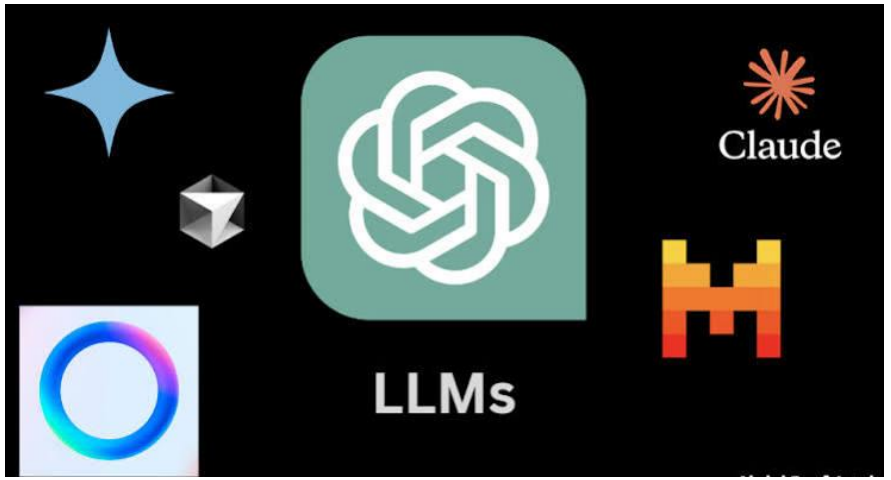
Causal Bayesian Optimization with Unknown Graphs. Durand J., Annadani Y., Bauer S., Parbhoo S. https://arxiv.org/abs/2503.19554, 2025.
Preference-Based Causal Bayesian Optimization: Active Causal Discovery from Ordinal Feedback. Bergossi C, Negri M., Roth V., Parbhoo S., arxiv preprint, 2026.
Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

3. Language Model Oracles for Unknown Graphs

Given a set of premises of independences and correlations, and a causal hypothesis, an LLM must deduce if the hypothesis is true or false.

Assume 4 variables A, B, C, D. A correlates with B, C and D. B correlates with D. C correlates with D. B and D are independent given A and C.

Hypothesis: There is at least one collider between A and B.

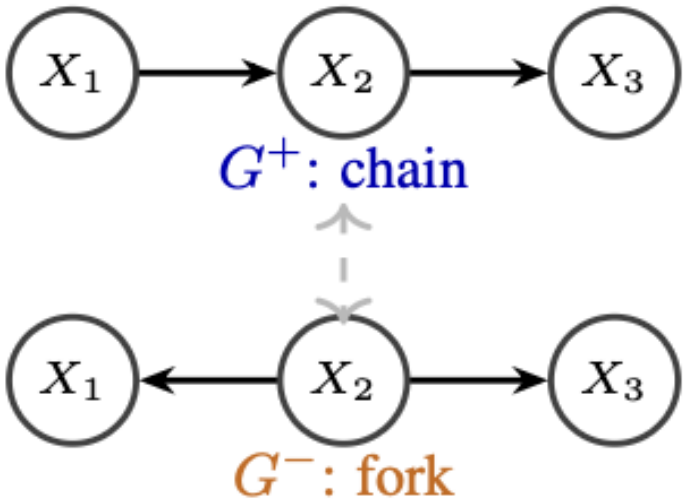


'False'

Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

Key Insight: Language models do not suffice for causal discovery

- Pretraining methods such as SFT, ICL cannot distinguish between seemingly similar pairs.



Why?

Kernel obstruction theorem

The discrimination margin between seemingly similar pairs is bounded by $B\kappa\sqrt{2\delta}$. *Distinguishing between seemingly similar graphs requires* $B > \gamma/\kappa\sqrt{2\delta}$ *which diverges as* $\delta \to 0$.

Implication:

Pre-training methods degrade with graph complexity due to geometric impossibility.

- LLMs suffer from similar non-identifiability if **relying purely on observational data**.

Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

Agentic-CBO with Language Model Oracles

Generate Potential Causal Hypotheses

A frozen LLM proposes candidate causal graphs $\{\mathcal{G}_1, \ldots, \mathcal{G}_n\}$ based on a uniform prior i.e each candidate is equally likely $\pi_0 = 1/n$

Choose a pair of variables to query → **Query the LLM Oracle** → **Perform A Bayesian Belief Update**

For each pair of variables (V_i, V_j), intervene on V_i and observe what happens to V_j

Ask the LLM: Does intervening on variable V_i change V_j? Record the output $r \in \{0, 1\}$

Reweight the candidate graphs by whether LLM prediction matches r.

Converges after m rounds

Return the MAP graph

$$\mathcal{G}^* = \arg\max_i \pi_i$$

Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

Agentic CBO identifies all types of causal relationships including colliders, without retraining the LLM

Tier	Model	Par.	Chi.	Anc.	Des.	Col.	Con.	F1	Acc.
LOW	Gemma-3-12B	81.4	79.8	76.2	74.6	70.3	83.1	77.6	89.4
	LLaMA-7B†	76.2	74.8	71.3	69.5	65.1	78.4	72.6	86.1
MID	Qwen3.5-122B	89.8	88.7	86.1	85.3	81.6	91.2	87.1	93.9
	Llama-3.3-70B	88.1	87.4	84.9	83.5	79.8	90.3	85.7	93.2
HIGH	GLM-5.1*	**95.3**	**94.7**	**92.6**	**91.8**	**88.4**	**96.2**	**93.2**	**97.1**
	Qwen3-30B*	94.6	93.9	91.4	90.5	87.1	95.4	92.1	96.5
Baselines from Jin et al. (2023) (no A-CBO):									
	GPT-4 zero-shot	—	—	—	—	—	—	29.1	64.6
	LLaMA-7B finetuned	—	—	—	—	—	—	92.0	97.5

Zhijing Jin, Jiarui Liu, Zhiheng Lyu, Spencer Poff, Mrinmaya Sachan, Rada Mihalcea, Mona Diab, and Bernhard Scholkopf. Can large language models infer causation from correlation? arXiv preprint arXiv:2306.05836, 2023.

Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

Agentic CBO scales to higher dimensions

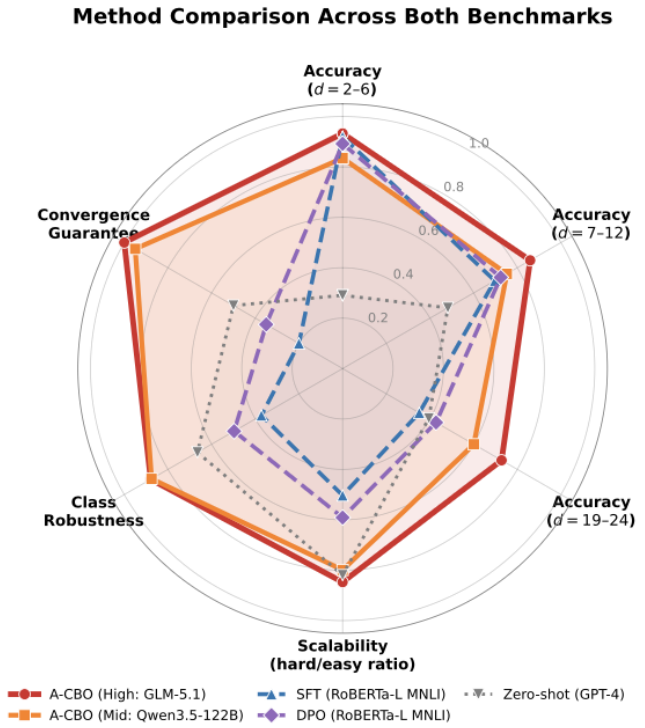
Method	Avg.	Accuracy by Graph Size (%)				Per-Class Accuracy (%, avg. all d)					
		$d=7$–10	$d=11$–15	$d=16$–20	$d=21$–24	Par.	Chi.	Anc.	Des.	Col.	Con.
RoBERTa-L SFT, orig. $d \leq 6$	98.2*	*— not evaluated on extended —*				96.2	95.7	93.9	96.6	92.2	98.7
RoBERTa-L SFT, extended	52.2	69.7	58.2	45.8	35.1	54.7	53.2	50.8	50.3	46.5	57.8
RoBERTa-L DPO, extended	58.0	73.0	63.4	52.7	42.9	60.6	58.8	56.8	56.0	51.6	64.1
A-CBO (GLM-5.1*)	**86.0**	**91.4**	**88.7**	**84.2**	**79.8**	**84.8**	**84.1**	**81.2**	**80.2**	**75.2**	**86.0**
Δ (A-CBO − SFT)	+33.8	+21.7	+30.5	+38.4	+44.7	+30.1	+30.9	+30.4	+29.9	+28.7	+28.2
Δ (A-CBO − DPO)	+28.0	+18.4	+25.3	+31.5	+36.9	+24.2	+25.3	+24.4	+24.2	+23.6	+21.9

∗From Jin et al. (2023); per-class values are F1 on original benchmark. ⋆ GLM-5.1 evaluated in thinking mode.

Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

Agentic CBO's inference procedure drives empirical gains

	Corr2Cause						Ext. C2C ($d = 7$–24)					
	Zero-shot		+A-CBO		Δ		Zero-shot		+A-CBO		Δ	
Model	F1	Acc	F1	Acc	ΔF1	ΔAcc	F1	Acc	F1	Acc	ΔF1	ΔAcc
LLaMA-7B†	26.8	17.36	72.6	74.2	+45.8	+56.84	37.8	39.2	52.9	54.6	+15.1	+15.4
Gemma-3-12B	27.4	30.2	77.6	79.1	+50.2	+48.9	42.9	44.7	56.2	57.9	+13.3	+13.2
Qwen3.5-122B	30.1	33.4	87.1	88.3	+57.0	+54.9	52.6	54.8	73.7	75.1	+21.1	+20.3
Qwen3-30B⋆	32.8	36.1	92.1	93.0	+59.3	+56.9	57.8	60.1	81.0	82.3	+23.2	+22.2
GLM-5.1⋆	34.2	37.5	**93.2**	**94.1**	+59.0	+56.6	59.9	62.4	**81.9**	**83.2**	+22.0	+20.8



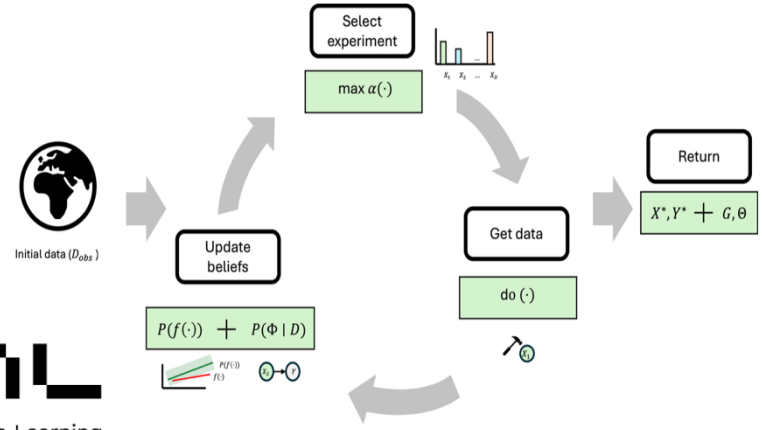
Why LLMs fail at Causal Discovery and How Interventional Agents Escape. Roy A, Parbhoo S, https://arxiv.org/pdf/2605.27567, 2026

Summary

- Determining how to act in complex systems requires a **causal understanding**. Such an understanding cannot be obtained from purely observational data.
- Causal Bayesian Optimization is a powerful technique for learning how to act based on available causal knowledge.
- In practice, we **rarely have complete knowledge** of the causal graph.
- **Weaker supervision** in the form of unknown graphs, ordinal feedback and LLM oracles **can help. Are there other types of weaker supervision that can help?**

Thank you!

- Jean Durand, Cesare Bergossi, Amartya Roy, Yashas Annadani, Marcello Negri, Volker Roth, Stefan Bauer, AI4AI lab and many others!



AI4AI AI for Actionable Impact

IMPERIAL

CIFAR

mcml Munich Center for Machine Learning