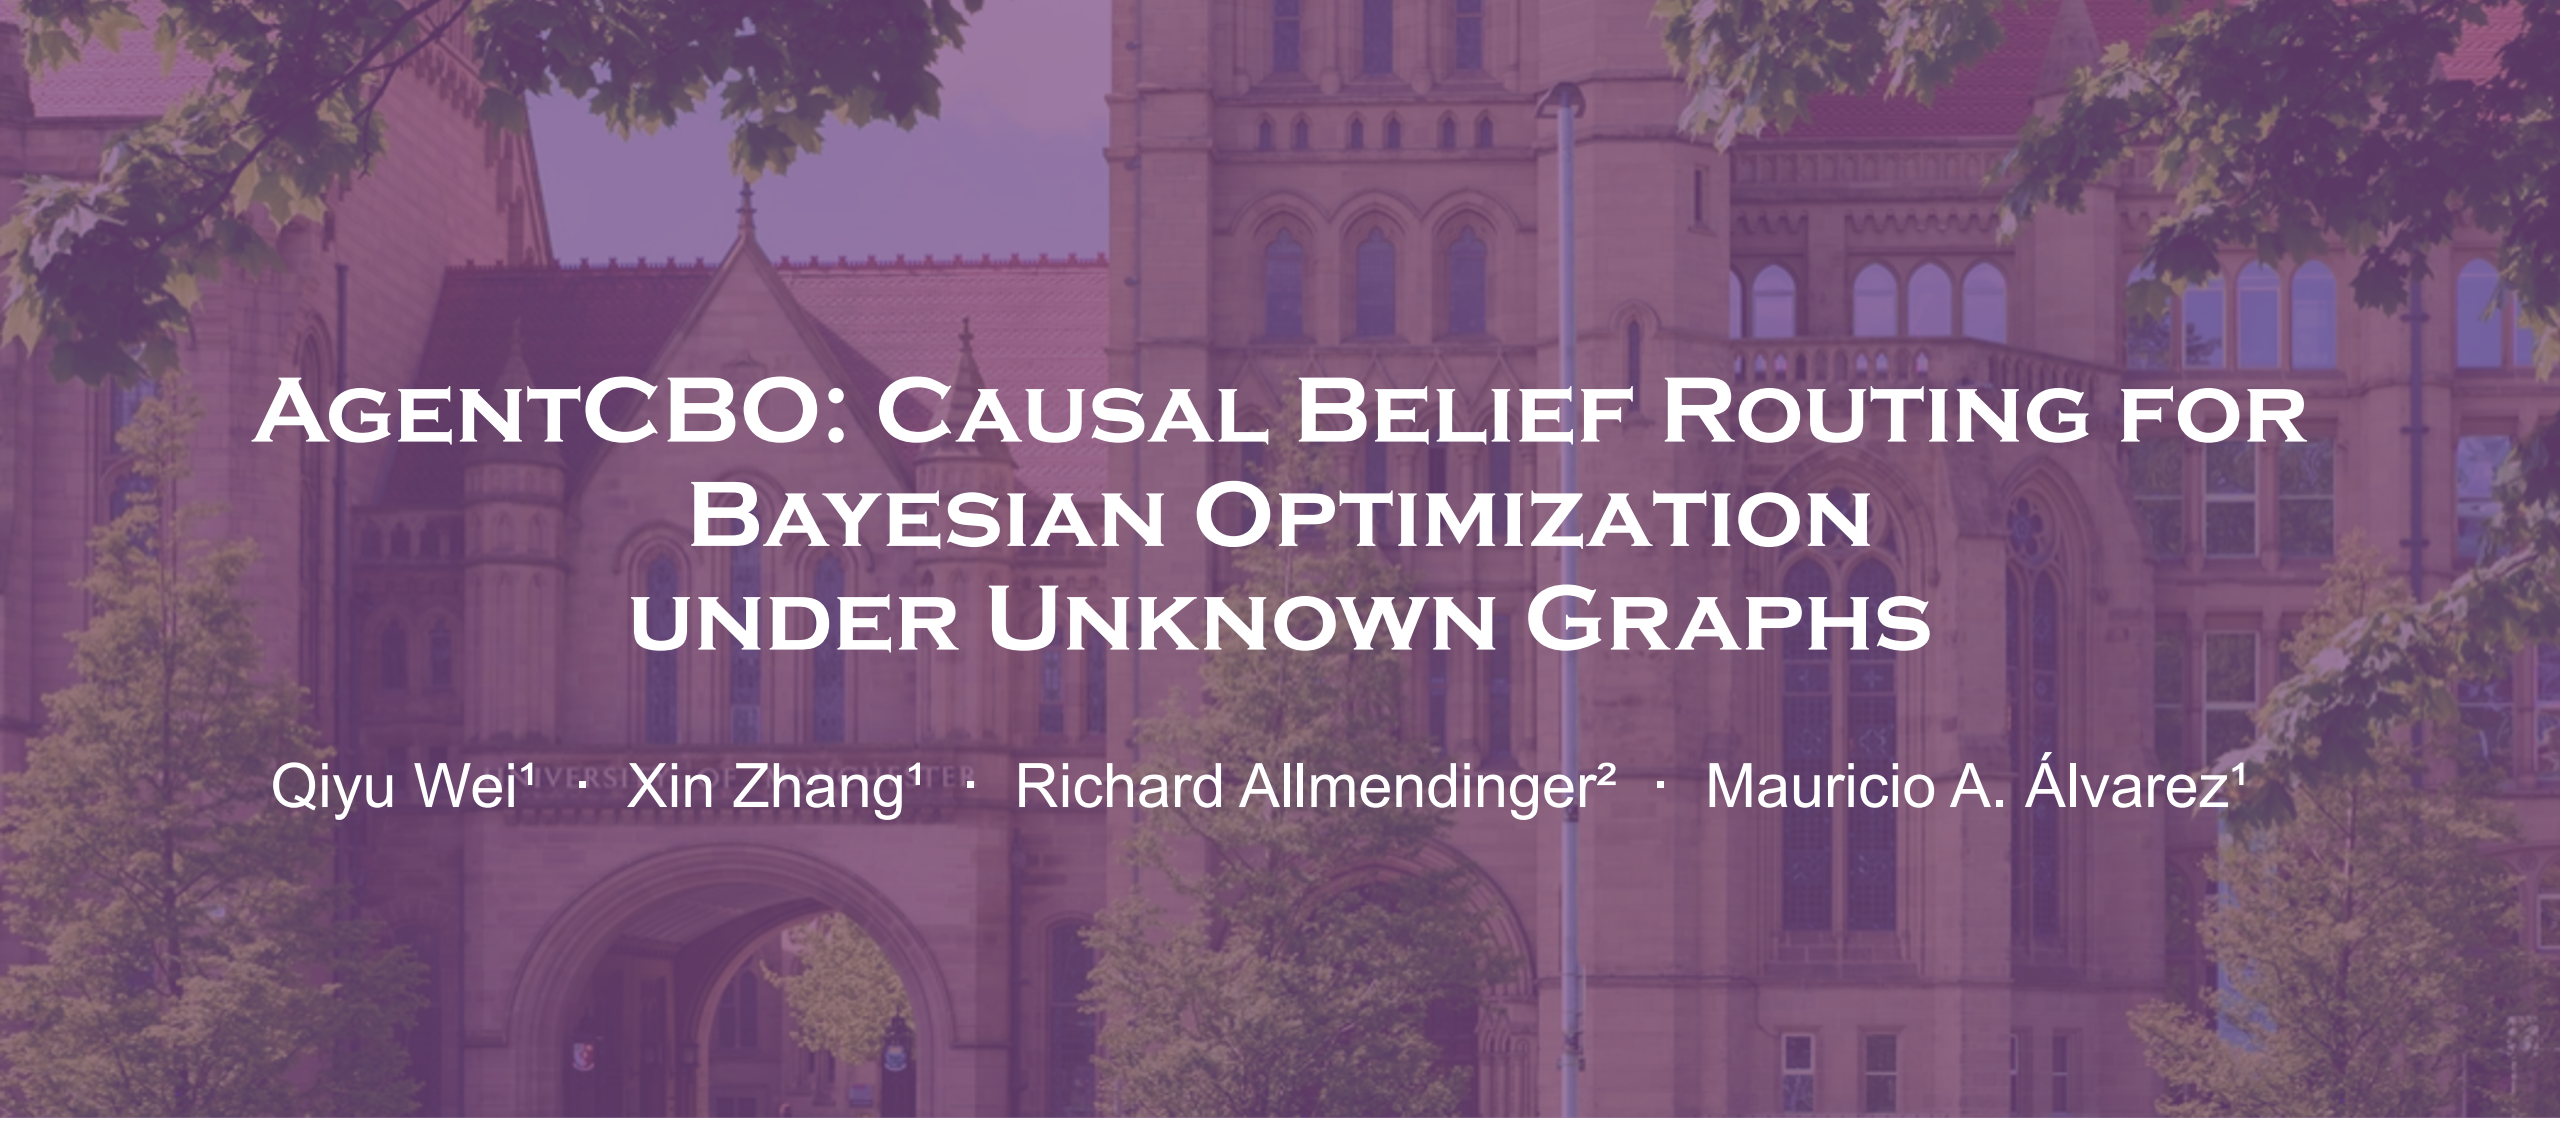

AgentCBO: Causal Belief Routing for Bayesian Optimization under Unknown Graphs

Qiyu Wei[1] · Xin Zhang[1] · Richard Allmendinger[2] · Mauricio A. Álvarez[1]

Presenter: Qiyu Wei

The University of Manchester, United Kingdom

Global Optimization

Consider a function $f : \mathcal{X} \to \mathbb{R}$ where $\mathcal{X} \subseteq \mathbb{R}^D$ is a bounded domain.

$$x_M = \arg\min_{x \in \mathcal{X}} f(x).$$



Option 1: Use previous knowledge? Option 2: Grid search? Option 3: Random search?

Bayesian Optimization is a better solution.

Bayesian Optimization

BO = Surrogate model + Acquisition Function

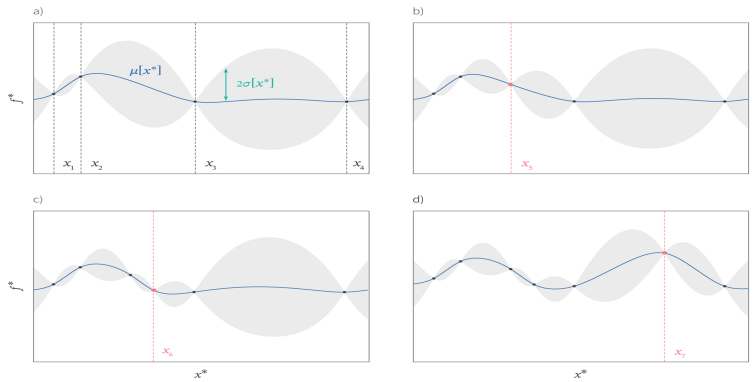

Fig. 1 Gaussian Process model

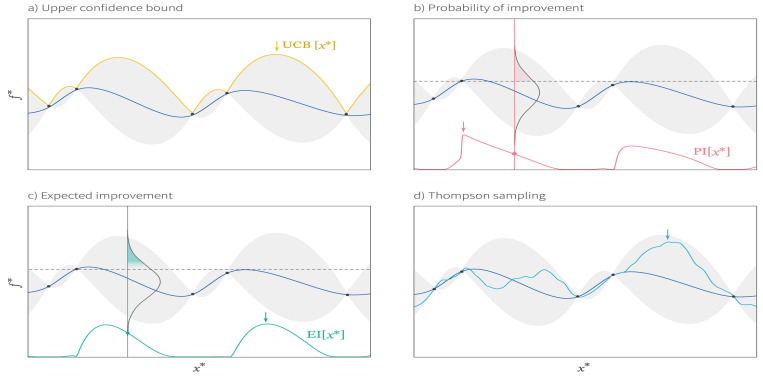

Fig. 2 Acquisition Function

AutoML ⊃ BO ⊃ GP, TPE (Tree-structured Parzen Estimator)

UCB, EI, PI, TS, KG ...

Causal Bayesian Optimization

- **BO:** Independently and cost-ineffective → suboptimal
- **CBO:** **Use Causal structure** existing among input parameters
  - Observe on all parameters (in BO) → **Intervention** (in CBO).

$$P(Y \mid X = x) \rightarrow P(Y \mid \mathrm{do}(X = x))$$

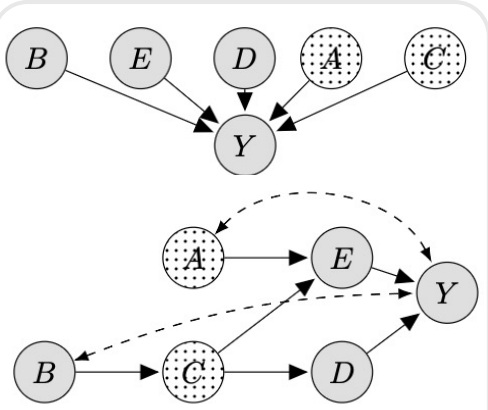

Fig. 3 BO vs CBO

- **Intervention**: do effects estimated via do-calculus.
  - do: A operator.
  - do-calculus: Estimations of do(·). Object: Find the highest informative subset.

[1] *Aglietti V, Lu X, Paleyes A, et al. Causal bayesian optimization[C] International Conference on Artificial Intelligence and Statistics. PMLR, 2020: 3155-3164.*

What a known graph gives CBO

1 · Choice

exploration sets

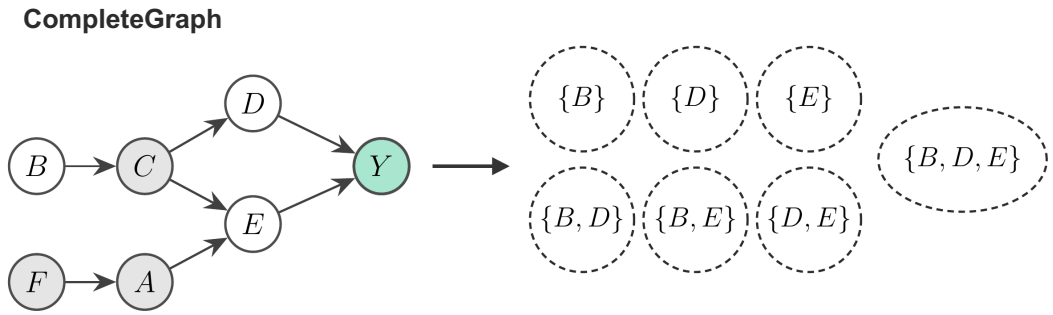

which intervention sets are worth trying

white B, D, E are manipulable: all 7 subsets; gray F, A, C are not

2 · Starting point

do-prior

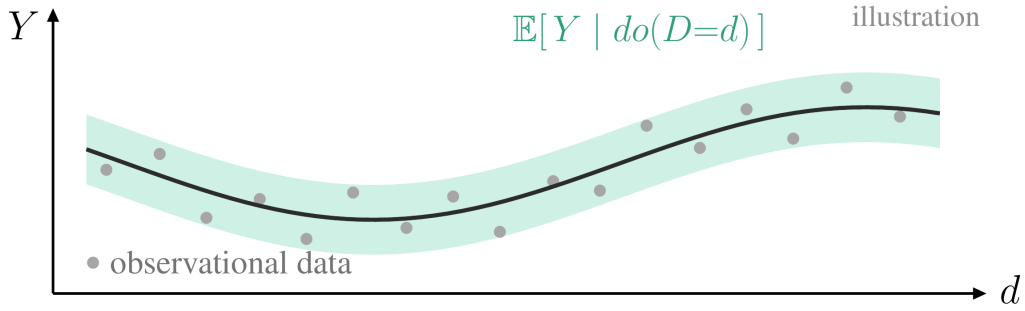

GP prior mean of E[Y | do(D=d)] for each set

estimated from observational data + do-calculus

Gray dots	cheap passive records (patient charts), not experiments (treatments); n_obs = 100
CBO decides	causal EI picks the next costly experiment from the menu (budget T = 50)
Unknown graph	no choice, no starting point

Choose interventions

CompleteGraph · which experiment should we run next?

1 Choose variables

2 Set their values

3 Observe

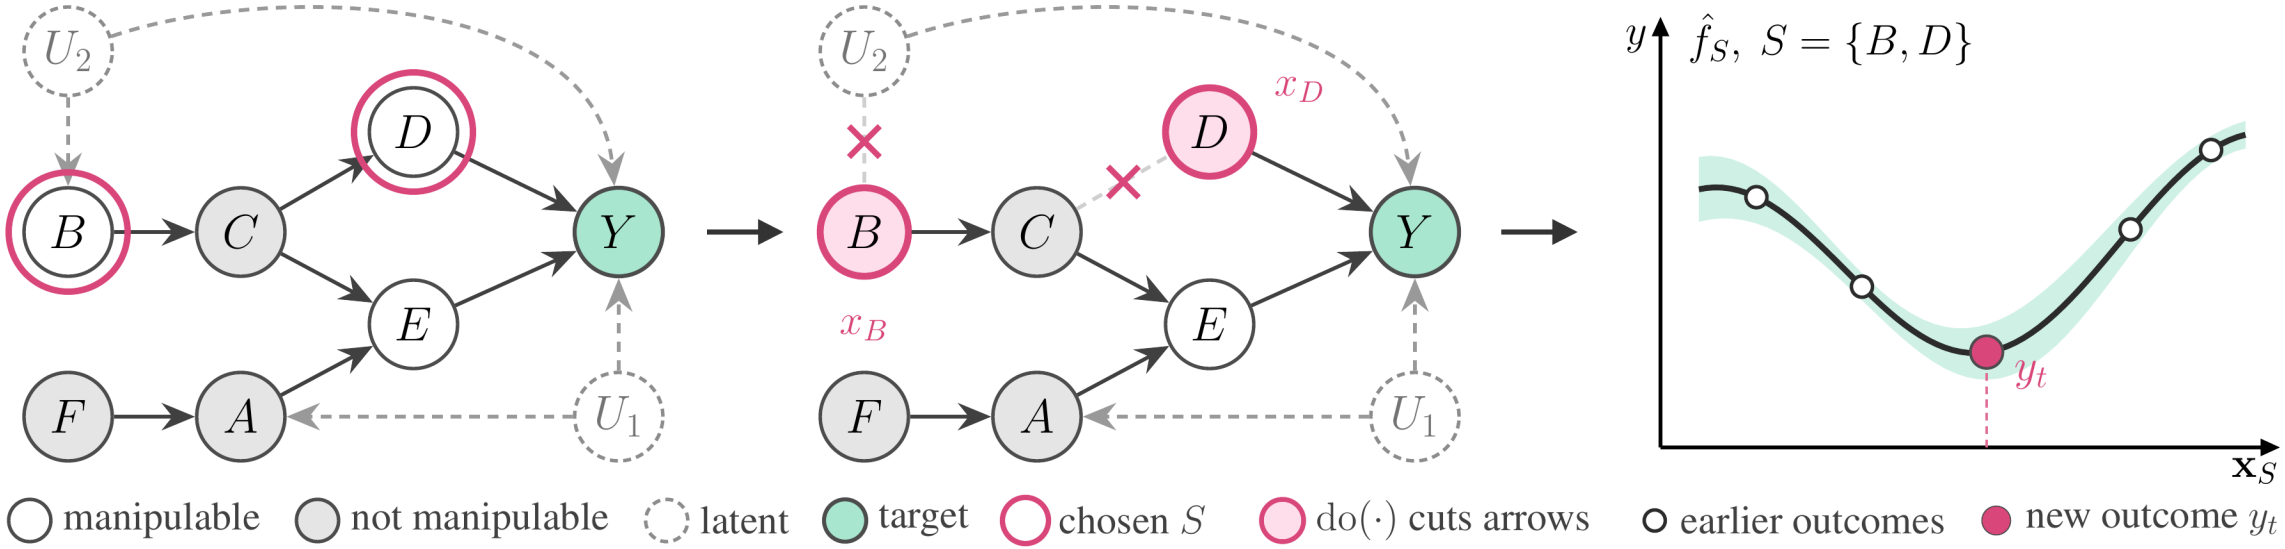

Choose (S, x) to minimize E[Y | do(S = x)]

Causal graphs guide two decisions

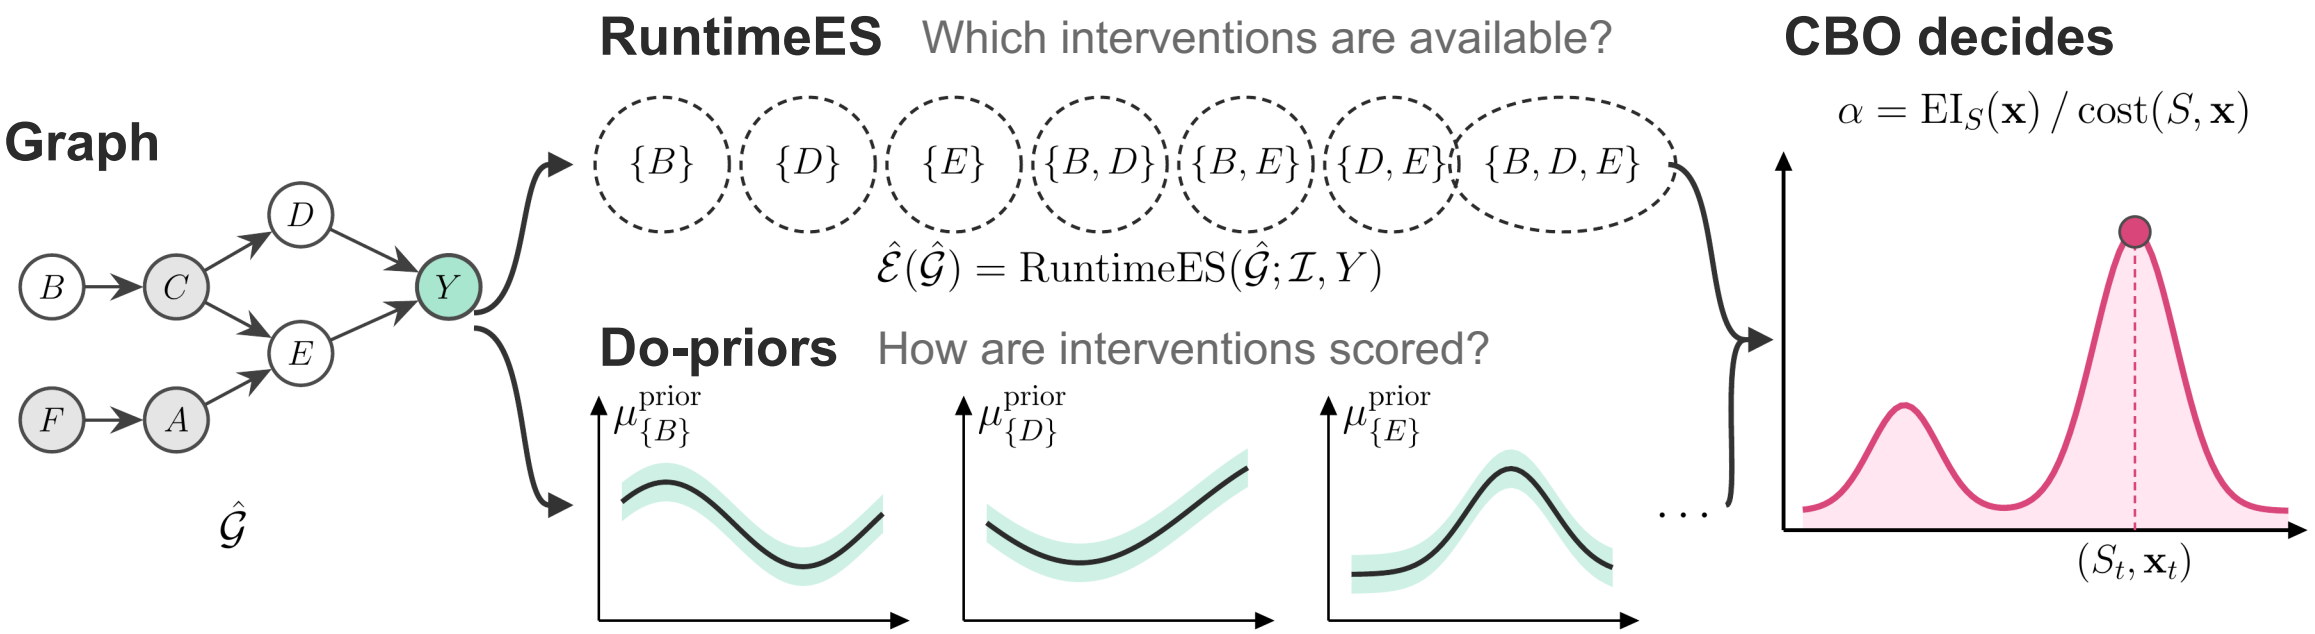

CEI = expected improvement / intervention cost

Existing methods

LIMIT: enumerates DAGs → small graphs only

CEO

Graph posterior

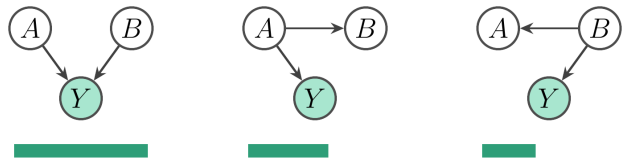

MAINTAINS

posterior over DAGs

EVIDENCE

obs + int likelihood

With BO

entropy-search acquisition

trades off structure learning and optimization

LIMIT: enumerates DAGs → small graphs only

GACBO

Graph hypotheses

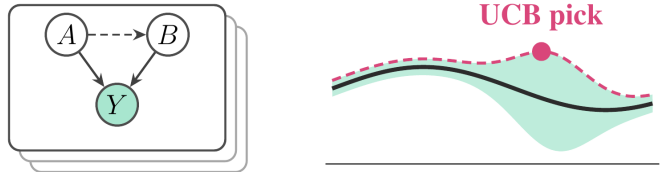

MAINTAINS

plausible (G, f, a) triples

EVIDENCE

Bayes scores + GP confidence sets

With BO

UCB over (G, f, a)

discovery only when it serves reward

CBO-U

Target-local summary

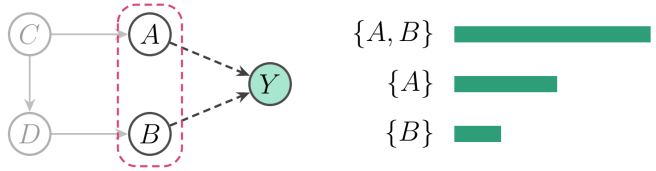

MAINTAINS

posterior over Y's direct parents

EVIDENCE

parent-set conditional-independence tests

With BO

parent-aware surrogate

scales by focusing inference on Y's parents

Shared limitation Graph coupled tightly with discovery; budget can go to structure irrelevant to the optimizer

[1] Branchini N, Aglietti V, Dhir N, et al. Causal entropy optimization, AISTATS2023

[2] Mukherjee S, Zhang M, Flaxman S, et al. Graph agnostic causal Bayesian optimisation, NeurIPS BDU Workshop 2024

[3] Durand J, Annadani Y, Bauer S, et al. Causal Bayesian optimization with unknown graphs, ICLR 2025 submission

Existing methods vs. Ours

	CEO	GACBO	CBO-U	**AgentCBO (ours)**
Graph	candidate DAGs,	plausible DAGs + GP confidence bands	posterior over Y's parent sets	**one committed DAG (commit gate)**
Learned from	obs + int data	obs + int data	obs + int data	**obs data + LLM prior;**
Experiments	info on (y*, G) per cost	UCB over graphs and functions	causal EI on a parent-averaged surrogate	**only to improve Y (CBO's causal EI)**
What matters	every graph difference	structure that serves finding the best action	only Y's direct parents	**only differences that change the menu or starting point**
Scale	enumerates candidates; ≤ 10 nodes	enumerates all DAGs; small graphs only	ER up to 100 nodes; assumes Y unconfounded	**ordered up to 50**

The graph is a routing interface

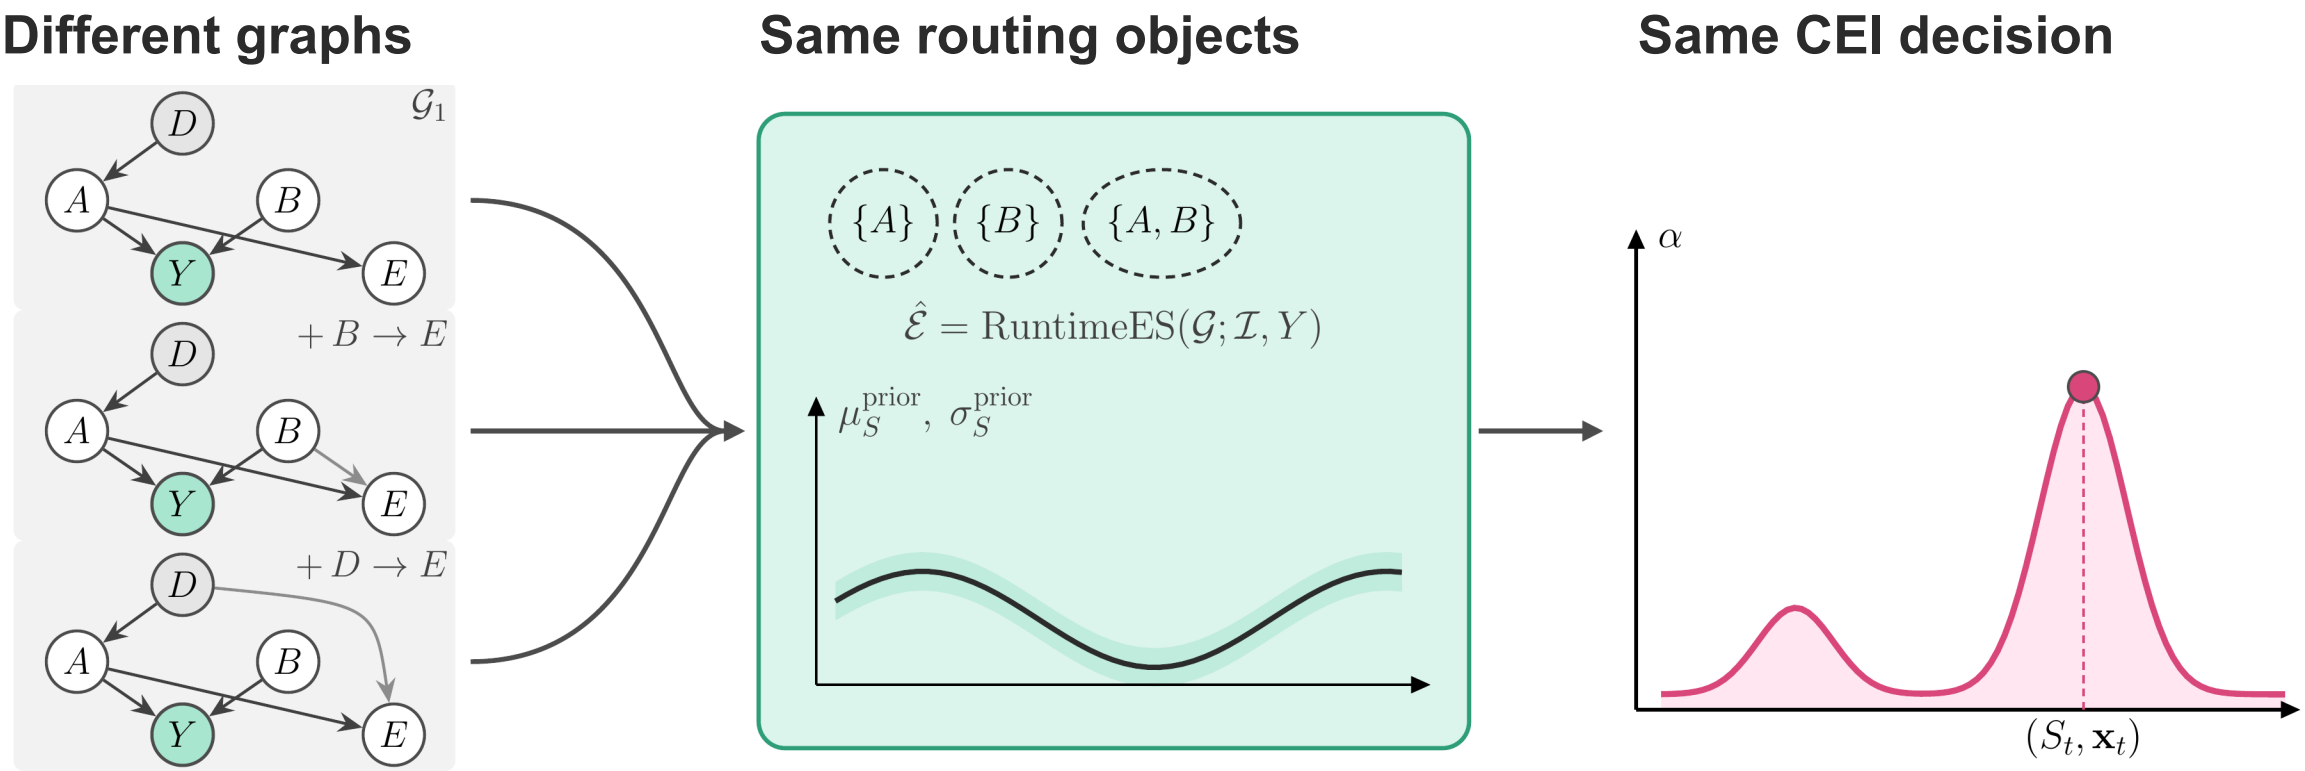

Different structural hamming distance, same CEI.

Equal routes give the same next decision

At one matched update

Different graphs

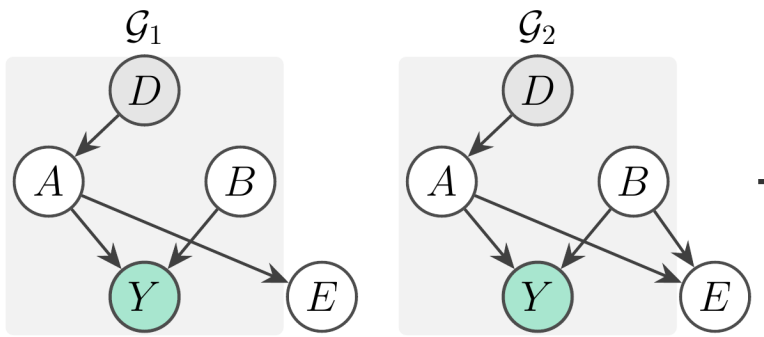

$\mathcal{G}_1 \neq \mathcal{G}_2$: extra edge $B \to E$

Equal routing inputs

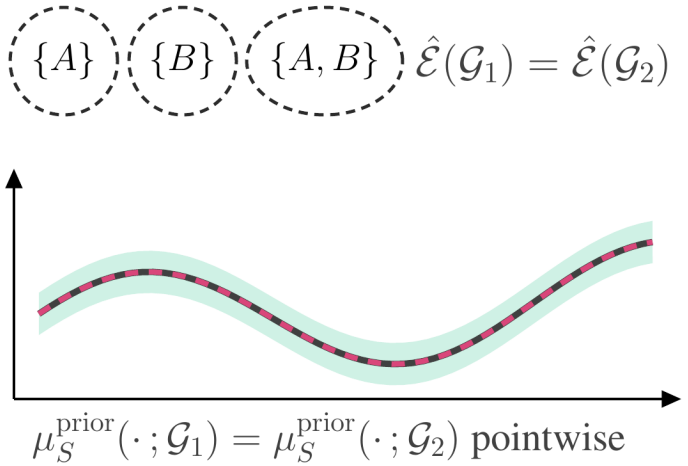

Same next-intervention distribution

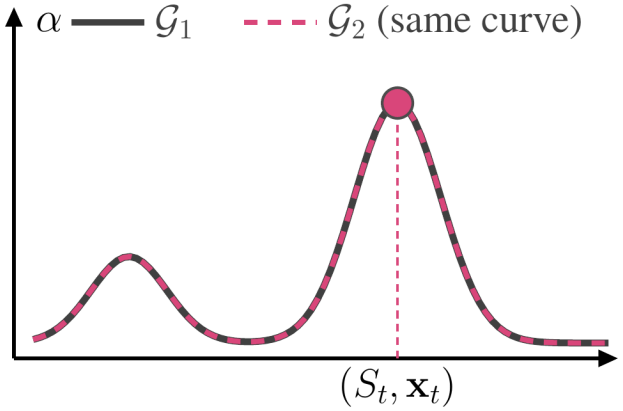

Same data, archive, domains, costs and rules; coupled randomness.
All graph dependence passes through the matched interface.

Edge accuracy is not the right target

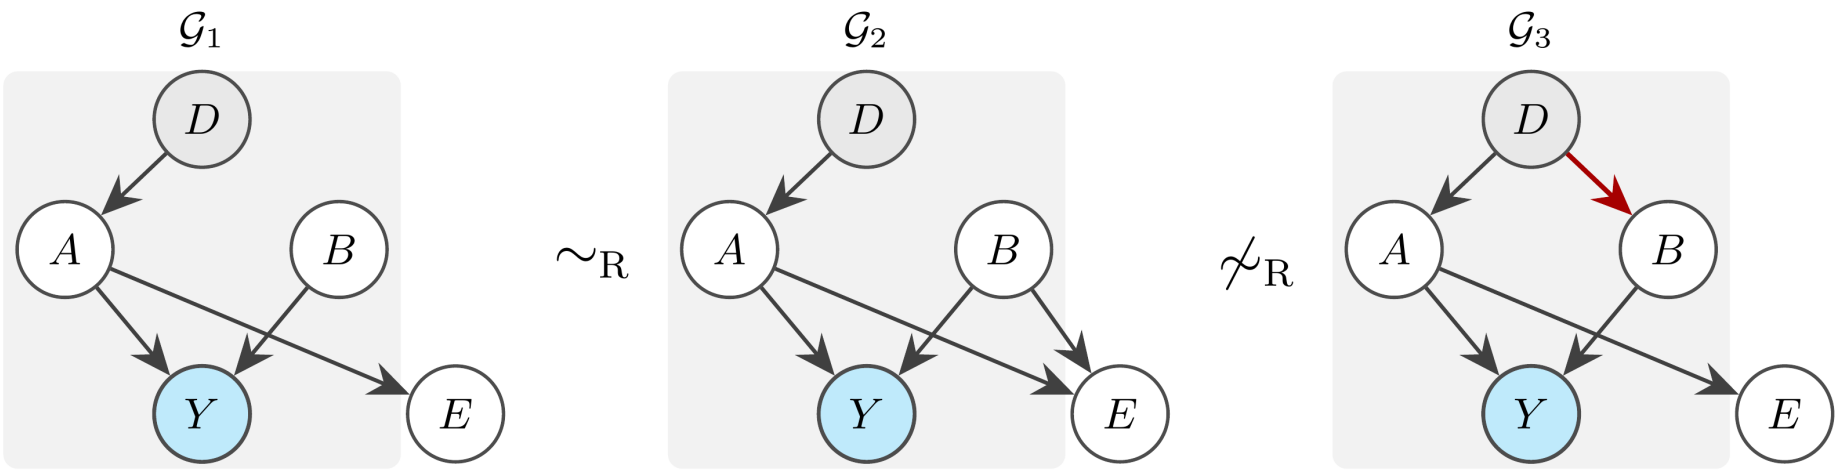

Routing-equivalent $G_1 \sim_R G_2$

Same labeled RuntimeES + pointwise-equal do-prior objects

G_3 is routing-critical

1 · Same choice

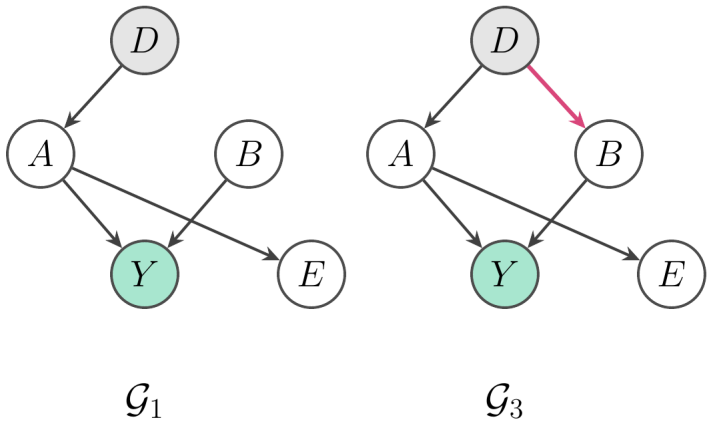

Ancestors {A, B} unchanged → same RuntimeES

2 · B is generated differently

$do(A{=}a)$: fix A, generate the rest

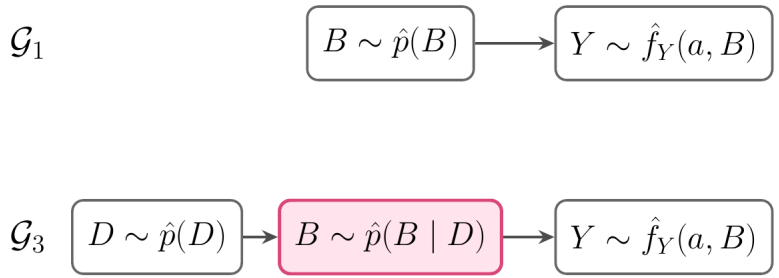

D→B changes how B is generated

3 · The prior of {A} may shift

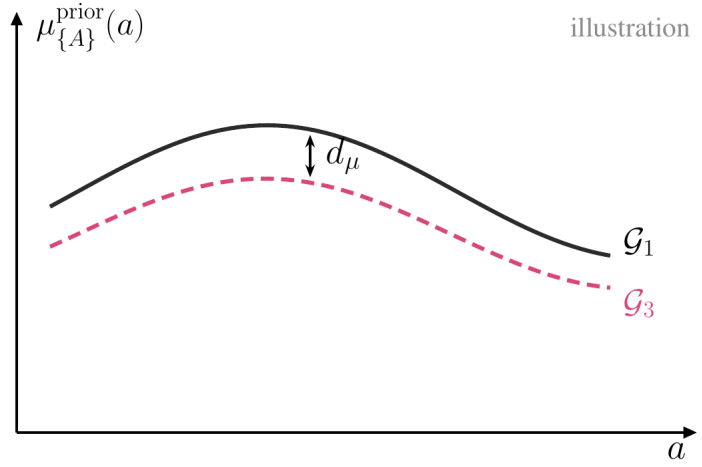

critical only if dμ > δμ

LLMs and guarded

LLMs judge pairwise causal relations well, but full-graph recovery may be wrong and can mistake correlation for causation.

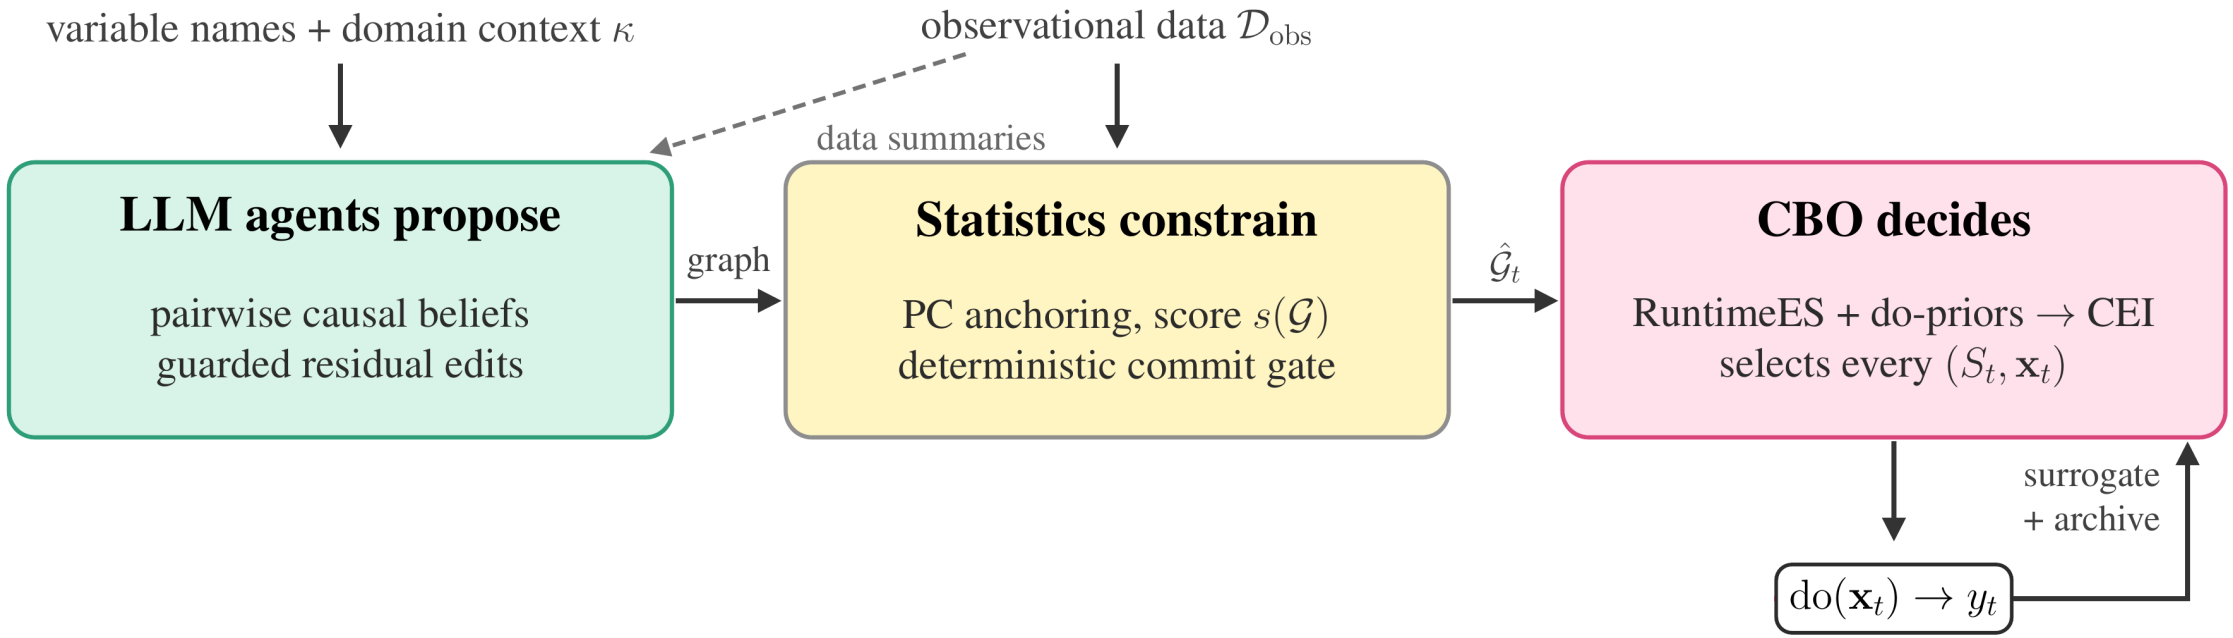

Pipeline

Propose → **Calibrate** → **Repair** → **Commit**

Names + context

Pairwise LLM beliefs

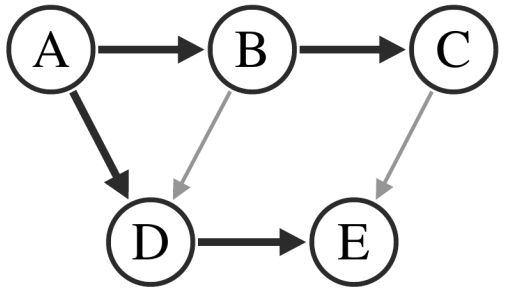

line width $\propto$ LLM belief

+

Observational discovery

PC anchoring

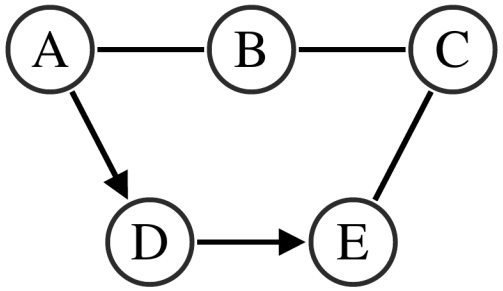

PC on $\mathcal{D}_{\text{obs}}$; plain line = unoriented

→

Structured proposal

Backbone + residuals

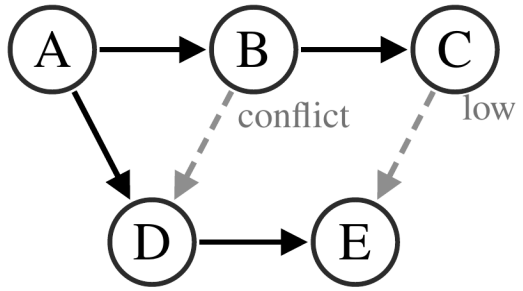

backbone $\mathcal{B}_t$ residual $\mathcal{R}_t$

Output: candidate graph + protected backbone + uncertain pairs

Pipeline

Propose → **Calibrate** → **Repair** → Commit

Calibration

LLM configures statistical discovery

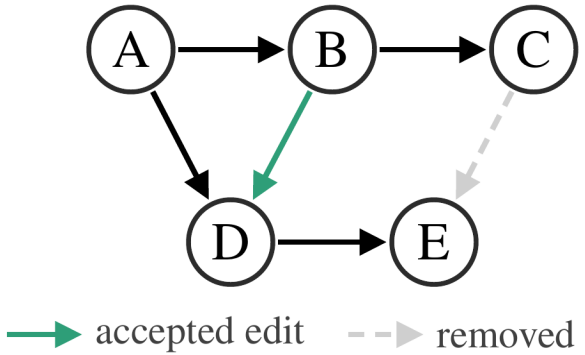

Why: re-check uncertain pairs with a discovery method(PC, GES, DirectLiNGAM) suited to the data

Repair

Local residual edits

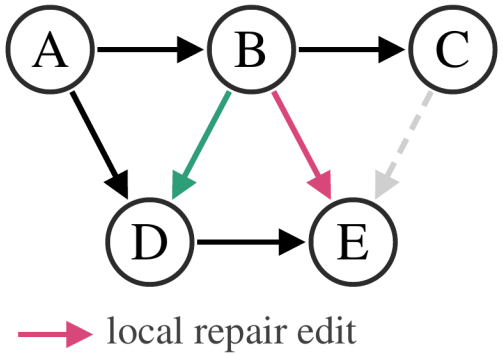

Why: fix the few target-relevant errors left, using CI and confounder checks

Shared guard

Accept only if the score improves

accept if $s(\bar{\mathcal{G}}_t^k) < s(\mathcal{G}_t^{k-1})$

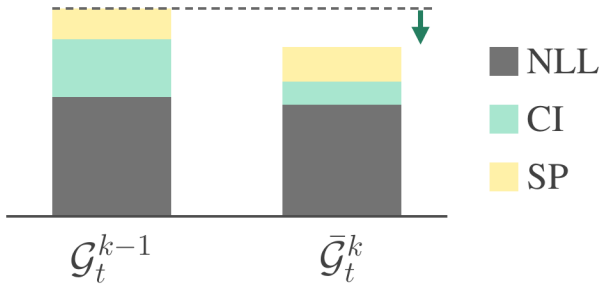

- **NLL**: negative log-likelihood on held-out data
  (Evaluating the goodness of fit of a model to the data)
- **CI**: conditional-independence mismatches
  (Penalize Independent-Looking Yet Correlated Data)
- **SP**: sparsity penalty
  (penalize excessive edge count)

Calibration involves recalculating the logic and performing a full, end-to-end rerun; Repair entails addressing specific, remaining issues one by one.

Pipeline

Propose → Calibrate → Repair → **Commit**

Candidate graph

Confidence check

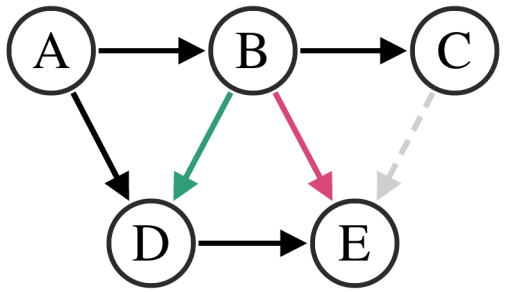

$(\tilde{\mathcal{G}}_t,\ c_t)$

Routing proxies

Ancestry + ancestor nSHD

Commit Gate

- $c_t \geq \tau_c \quad \text{and} \quad \tilde{\mathcal{G}}_t \neq \hat{\mathcal{G}}_{t-1}$
- $\mathcal{A}_Y(\tilde{\mathcal{G}}_t) \neq \emptyset$
- $c_t \geq \min\left(\tau_c + \frac{1}{2}\, n\mathrm{SHD}_t,\ 0.9\right)$

Committed graph

Rebuild ES and priors

ACCEPT

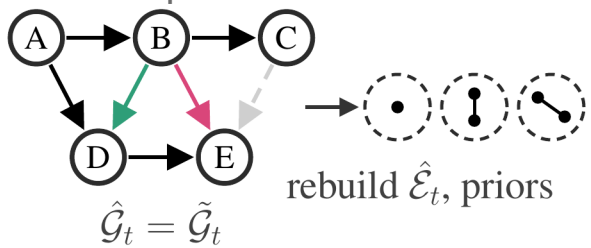

$\hat{\mathcal{G}}_t = \tilde{\mathcal{G}}_t$

rebuild $\hat{\mathcal{E}}_t$, priors

REJECT

keep $\hat{\mathcal{G}}_{t-1}$

- Minimum confidence threshold.
- The set of intervenable ancestors of Y within the candidate graph.
- The more extensive the modifications related to Y, the higher the required confidence level.

Commit is deterministic. Rejected proposal → keep the current graph.

Three variants

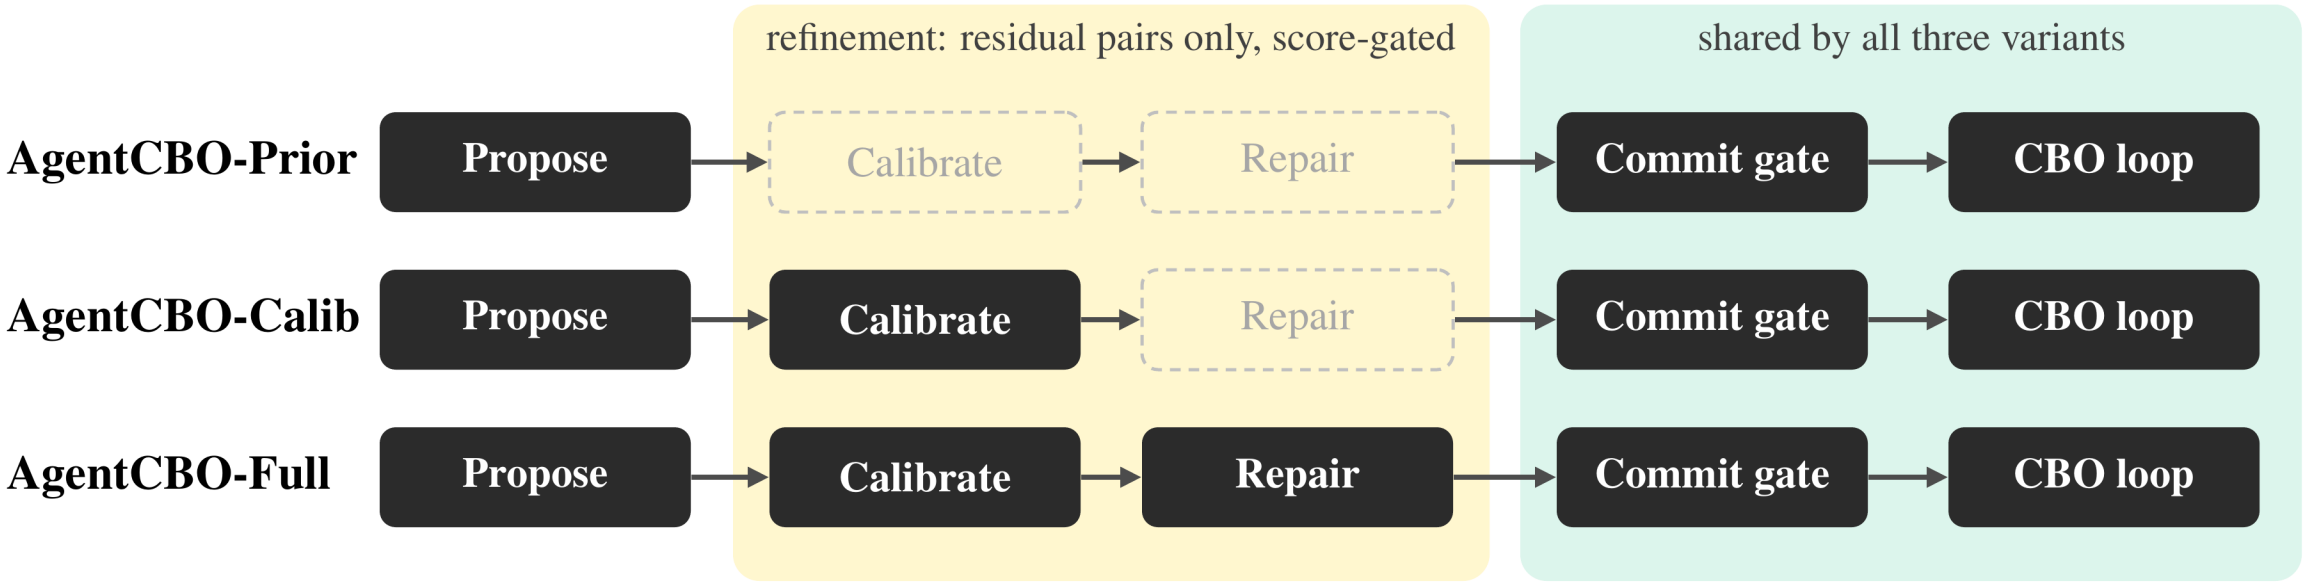

- Each added stage may edit only residual pairs, and only when the structural score improves.
- Choose different pipelines based on the varying difficulty levels of the tasks.

Graph updates keep the experimental history

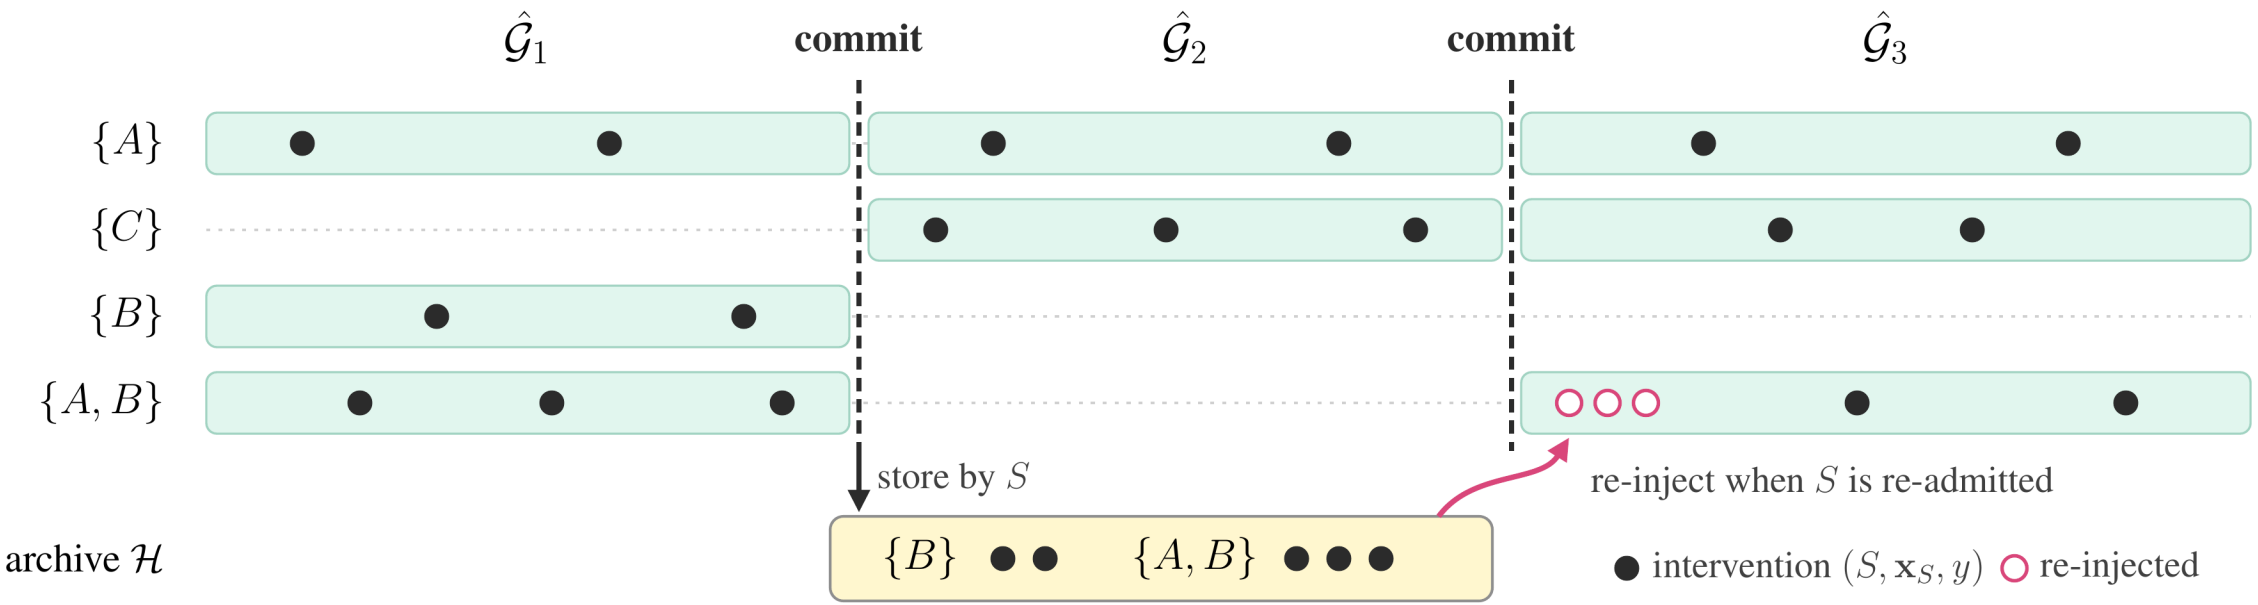

- Archived tuples are keyed by ES signature S and re-injected when a later commit re-admits S.
- Interventions update per-set surrogates and the archive; observational data drive graph discovery.

AgentCBO pipeline

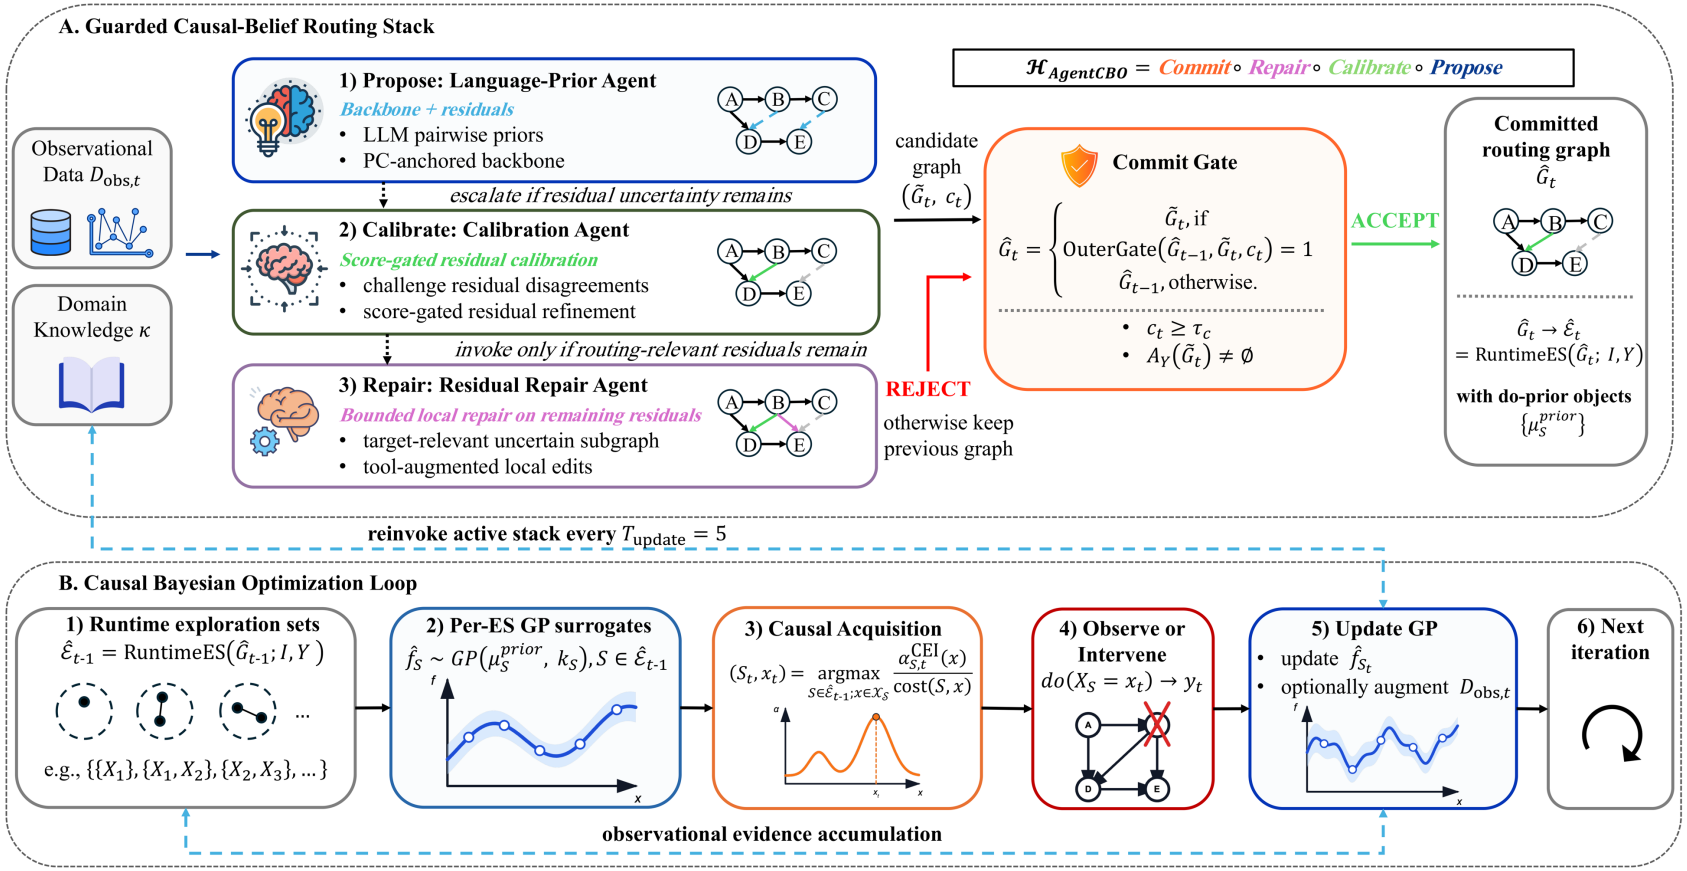

A · Guarded stack

Re-invoked every $T_{update} = 5$ steps

B · CBO loop

RuntimeES + do-priors → CEI → intervention

Results on ten benchmarks

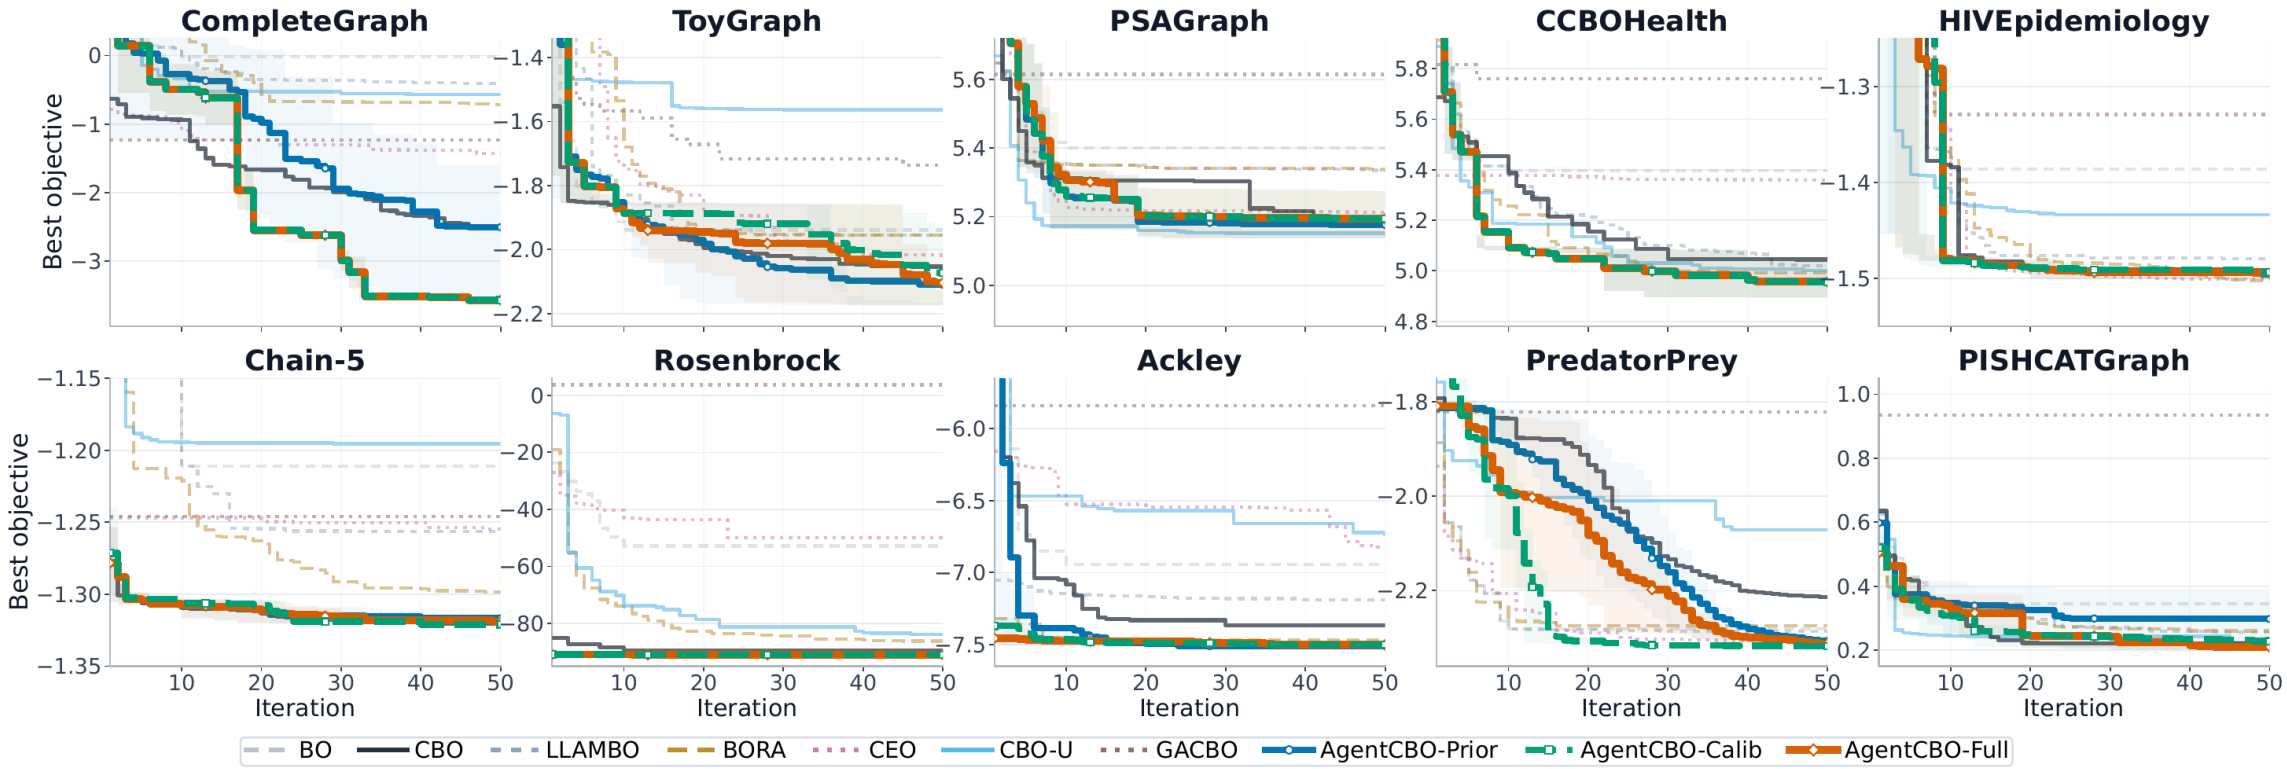

Best objective vs. iteration (T = 50) · 20 seeds with variance bands · lower is better

A learned graph beats the true one

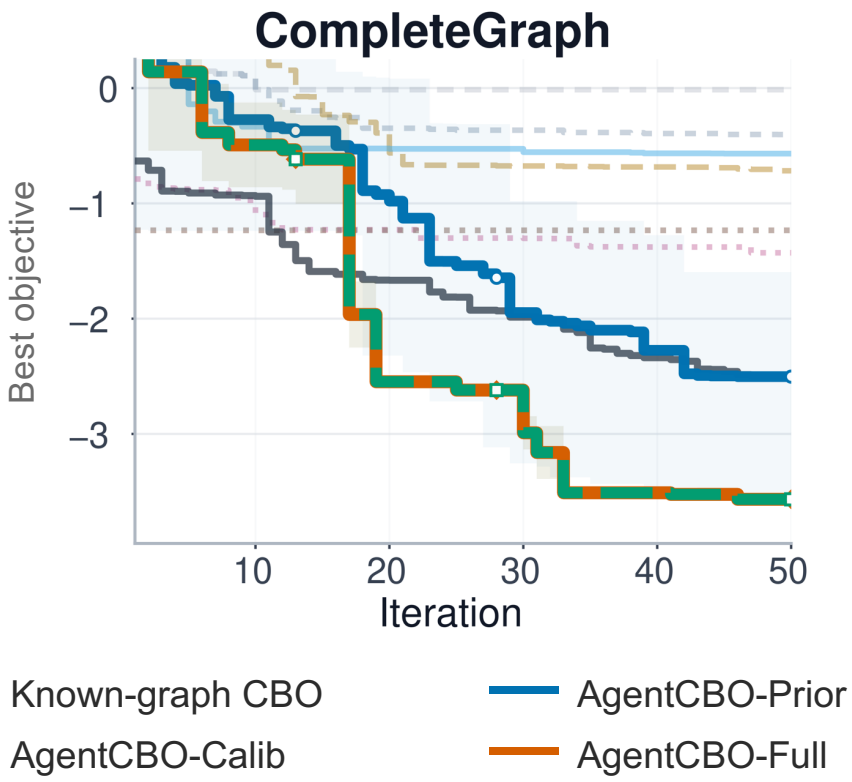

−3.57 **AgentCBO-Calib / Full**
± 0.07

−2.50 **AgentCBO-Prior**
± 1.21

−2.48 **Known-graph CBO**
± 1.14

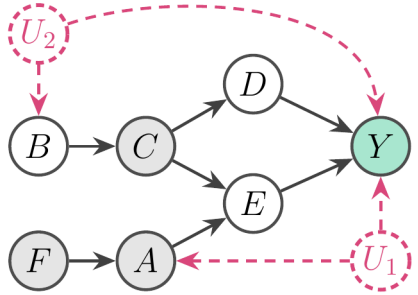

Latent U_1, U_2 confound the observed parents of Y: the learned graph acts as a routing proxy.

1 Sets 7 → 4; {B, D} picks 1/11 → 18/27.
2 Only {B}, {D}, {B, D}: 14/20 seeds hit y ≤ −3 (random pairs: 4/20).

Thanks!

Q&A

Contact: qiyu.wei@postgrad.manchester.ac.uk