

Causal Bayesian Optimization: Foundations, Methods, and Applications

Chenfeng Huang

*Department of Statistics & Data Science
University of California, Los Angeles*

chenfenghuang@ucla.edu

Thuy T. Le

*Department of Mathematics and Statistics
California State University, Long Beach*

Thuy.Le@csulb.edu

Zixuan Ma

*Department of Statistics & Data Science
University of California, Los Angeles*

zixuanma@ucla.edu

Hien Tran

*Department of Mathematics
North Carolina State University*

tran@math.ncsu.edu

Reviewed on OpenReview: <https://openreview.net/forum?id=XT6DC37m5I>

Abstract

Causal Bayesian Optimization (CBO) integrates causal inference with Bayesian optimization to enable sample-efficient intervention selection in systems governed by causal structure. This survey provides a comprehensive and systematic review of the CBO landscape, organizing the growing literature through a unified BO-loop perspective that reveals how causal assumptions shape four core components: intervention search spaces, surrogate construction, acquisition design, and decision policies. We organize methods along recurring design axes, including graph and system-knowledge assumptions, environmental assumptions, intervention representation, surrogate architecture, and decision rules, and we clarify conceptual and notational connections between CBO and adjacent fields, including causal bandits, Bayesian experimental design, safe optimization, policy search, and causal abstraction. To address the lack of standardized evaluation in the field, we introduce a reproducibility-oriented benchmark that covers hard- and soft-intervention settings, implements both the standard GAP metric and a new trajectory-aware Path-Aware GAP (PA-GAP) metric, and evaluates seven CBO methods alongside a non-causal BO baseline under a common scoring protocol. Within this benchmark-specified regime, where the benchmark SCM, intervention domains, and reference optima are fixed and known-graph methods receive the benchmark graph unless explicitly specified otherwise, our empirical study across thirteen datasets, three budget levels, and two metrics reveals that no single method dominates uniformly: rankings depend critically on the dataset, budget, metric, and method-specific use of causal information, and strong non-causal baselines remain competitive in several settings. We further add controlled graph-misspecification and omitted-variable stress tests, which show that rankings can change substantially when the learner-side causal information is perturbed. We therefore identify robustness to causal-assumption violations, scalable unknown-graph optimization, mixed intervention types, realistic cost models, tighter theoretical guarantees, and integration with modern representation learning and causal abstractions as open challenges for moving CBO from controlled benchmarks toward reliable deployment.

1 Introduction

Many scientific and engineering problems require the selection of interventions in systems whose variables are causally coupled. A clinician chooses a drug dosage that propagates through metabolic pathways to affect a clinical endpoint; an engineer tunes process parameters that cascade through a manufacturing pipeline; a policymaker adjusts economic levers whose effects depend on the causal structure of the economy. In each case, the key challenge is the same: the system’s response to an intervention is not governed by statistical association but by a causal mechanism, and purely correlational optimization can recommend actions that appear promising in observational data, yet fail or even cause harm when deployed (Pearl, 2009; Peters et al., 2017).

Bayesian Optimization (BO) offers a basic framework for sample-efficient optimization when evaluations are expensive, noisy, or constrained (Jones et al., 1998; Shahriari et al., 2016). By iteratively fitting a probabilistic surrogate and selecting queries via an acquisition function that balances exploitation against exploration, BO can identify near-optimal solutions with far fewer evaluations than grid or random search. However, standard BO treats the objective as a black box: it is agnostic to the causal relationships among the variables it manipulates. Causal Bayesian Optimization (CBO) closes this gap by embedding structural causal knowledge into the BO loop, enabling the optimizer to reason about *which* variables to intervene on, *how* observational data relate to interventional effects, and *where* the experimental budget is best spent (Aglietti et al., 2020).

Since its introduction, CBO has rapidly expanded into a diverse family of methods, spanning dynamic environments (Aglietti et al., 2021), safety constraints (Aglietti et al., 2023), mechanism-level surrogates with regret guarantees (Sussex et al., 2023), functional and contextual policies (Gultchin et al., 2023; Arsenyan et al., 2026), adversarial non-stationarity (Sussex et al., 2024), high-dimensional scaling (Wu et al., 2024), unknown-graph optimization (Branchini et al., 2023; Mukherjee et al., 2024), noise-robust acquisition with learned priors (Li et al., 2023), and abstraction-based causal decision making (Zennaro et al., 2024; Dyer et al., 2025; Zeitler, 2025). Section 3 organizes these developments along five recurring design axes—graph and system-knowledge assumptions, environment, intervention representation, surrogate architecture, and decision rule (Figure 1, Table 2).

Despite this progress, the field lacks two critical ingredients for maturation. First, a **unified conceptual framework** that reveals the shared structure across CBO variants and connects them to adjacent fields, such as causal bandits, Bayesian experimental design, active causal discovery, and safe reinforcement learning, is missing. Second, **standardized evaluation infrastructure** is absent. Existing papers use different datasets, incompatible codebases, inconsistent intervention conventions, and different metrics, making cross-paper comparisons unreliable. A method that excels when the graph is known and interventions are low-dimensional may behave very differently under soft interventions, unknown graphs, or context-dependent policies, yet these distinctions are obscured when each paper evaluates in isolation.

This survey addresses both gaps. We provide a systematic, technically detailed review of the CBO landscape organized through a unified BO-loop perspective, and we introduce a reproducibility-oriented benchmark that standardizes and repackages inherited CBO tasks together with newly added scenarios under a common execution and scoring protocol. In terms of concreteness, our contributions are as follows.

- **A unified design-space perspective.** We organize CBO methods by graph and system-knowledge assumptions, environmental assumptions, intervention representation, surrogate architecture, and decision rule (Section 3). This taxonomy is intended as a survey-level organization of recurring design patterns and couplings, rather than as an ablation study proving that components are interchangeable.
- **Connections to adjacent fields.** We clarify conceptual and notational links between CBO and causal bandits, Bayesian experimental design, active causal discovery, safe optimization, policy search, and causal abstraction / multi-scale decision making, while distinguishing such connections from formal equivalences or reductions (Section 3.7).
- **A reproducibility-oriented benchmark.** We release a standardized benchmark harness covering eight hard-intervention and five soft-intervention scenarios, combining inherited tasks from prior CBO/function-

network work with newly added or reconstructed scenarios. The benchmark provides standardized GAP, a new trajectory-aware PA-GAP metric, and a common best-so-far trajectory interface that enables fair re-scoring of outputs from heterogeneous implementations (Section 4).

- **A unified empirical comparison.** We evaluate seven CBO methods and a non-causal BO baseline across thirteen datasets, three budget levels, and two metrics. Because the compared methods do not all solve the same optimization problem, we do not pool them into cross-method rank plots: each method–task pair carries a formulation label distinguishing matched comparisons, stress-test comparisons, and descriptive single-method or single-dataset settings, the main text reports per-dataset comparisons, and budget-stratified descriptive rank-reliability checks are reported in the appendix. The results show that no method is statistically separable as uniformly best, rankings depend strongly on dataset–budget–metric interactions, and strong non-causal baselines remain competitive in several settings (Section 4.4–4.5).
- **Structured open problems.** We identify six concrete research directions (robustness to causal assumptions, scalability, richer intervention models, realistic evaluation, theoretical foundations, and integration with representation learning) that must be addressed for CBO to achieve reliable real-world deployment (Section 5).

Paper outline. Section 1 places this survey relative to existing reviews. Section 2 establishes the theoretical foundations in causal inference and Bayesian optimization. Section 3 presents the methodological frameworks organized by design axes. Section 4 describes the benchmark, metrics, datasets, and empirical results. Section 5 discusses open problems, and Section 6 concludes.

Related Surveys and Scope

Several surveys cover topics adjacent to CBO, but none provides the unified treatment of causal optimization methods and standardized empirical evaluation that we pursue here.

Bayesian optimization surveys. Comprehensive reviews of BO cover surrogate modeling, acquisition functions, and scalability (Shahriari et al., 2016; Frazier, 2018; Garnett, 2023), but treat the objective as a black box and do not address causal structure, intervention design, or the observation-intervention trade-off central to CBO.

Causal inference and causal discovery. Textbooks and surveys on causal inference (Pearl, 2009; Peters et al., 2017; Hernán & Robins, 2020; Glymour et al., 2019) provide the framework of the structural causal model that underlies CBO, but focus on the identification and estimation of causal effects rather than on sequential optimization of interventions under budget constraints. Causal discovery surveys (Spirtes et al., 2000; Chickering, 2002; Glymour et al., 2019) address the learning of the graph from data, which is a component of unknown-graph CBO but not its main goal.

Causal bandits and Bayesian experimental design. Causal bandits (Lattimore et al., 2016; Lu et al., 2021; Nair et al., 2021) share CBO’s goal of selecting interventions using causal structure but typically assume discrete action sets with parametric reward models, whereas CBO handles continuous domains with nonparametric surrogates; Bayesian optimal experimental design (Chaloner & Verdinelli, 1995; Foster et al., 2021) selects experiments to maximize information about model parameters rather than to optimize an outcome. Section 3.7 details both relationships.

Scope of this survey. We focus on methods that explicitly combine (i) a structural causal model or causal graph with (ii) a sequential Bayesian optimization loop to select interventions. We exclude pure causal effect estimation without an optimization loop, causal bandits without GP-based surrogates (except when contrasting with CBO), and reinforcement learning methods that learn policies through environment interaction without an explicit causal model. We include both known- and unknown-graph settings, hard and soft interventions, and single- and multi-objective formulations.

Table 1: Summary of notation used throughout this survey.

Symbol	Meaning
$M = \langle U, V, F, P(U) \rangle$	Structural Causal Model
$V = \{V_1, \dots, V_n\}$	Endogenous (observed) variables
U	Exogenous variables (determined outside the system)
Pa_i	Parents of V_i in the causal graph $\mathcal{G}$
$\mathcal{G}$	Causal directed acyclic graph (DAG)
$do(X=\mathbf{x})$	Hard intervention setting X to value $\mathbf{x}$
X, Y	Manipulable variables, target outcome
C	Context variables (used only by contextual extensions, Section 3.3)
$X_s \subseteq X$	Intervention scope (subset of manipulable variables)
$\mathbf{x}_s \in \mathcal{D}(X_s)$	Intervention value in the feasible domain
$f(X_s, \mathbf{x}_s)$	Interventional objective: $\mathbb{E}[Y \mid do(X_s=\mathbf{x}_s)]$ (C enters only in contextual extensions)
$\mathcal{GP}(m, k)$	Gaussian process with mean m and kernel k
$\mu_t(\cdot), \sigma_t^2(\cdot)$	GP posterior mean and variance after t observations
$\alpha(\cdot)$	Acquisition function
T	Total number of interventional trials (budget)
y^*	Reference optimum (for metric computation)
R_t	Best-so-far improvement ratio at trial t

Survey protocol and inclusion criteria. To make the construction of the review corpus explicit, we followed a targeted search-and-screening protocol with a literature cutoff date of June 28, 2026. We searched OpenReview, PMLR, arXiv, Google Scholar, and the proceedings of major machine learning and causality venues using keyword combinations such as “causal Bayesian optimization,” “causal global optimization,” “Bayesian optimization with interventions,” “causal optimization,” and “optimization under causal structure.” Candidate papers were first screened by title and abstract, and then by their methodological formulation. A work was included if its central contribution couples an explicit causal model, such as an SCM, causal graph, or distribution over graphs, with a sequential Bayesian-optimization-style procedure for intervention selection. We further used backward and forward citation tracing from the original CBO paper and subsequent representative extensions to identify closely related work. Papers whose primary objective is only causal effect estimation, graph recovery, bandit learning, experimental design, or reinforcement learning were excluded from the main corpus unless they directly clarify the boundary between CBO and adjacent fields. This protocol yields a focused corpus of methods in which causal assumptions actively modify the Bayesian optimization loop, rather than a broad survey of causal decision-making methods.

2 Theoretical Foundations

This section establishes the conceptual and mathematical foundations underlying Causal Bayesian Optimization. We first review structural causal models, interventions, and identifiability (Section 2.1). We then summarize classical Bayesian optimization with an emphasis on components that CBO modifies (Section 2.2). Finally, we formulate the CBO problem and highlight the key ways in which causality reshapes the BO loop (Section 2.3). Table 1 consolidates the notation used throughout the paper.

2.1 Causal Inference

Most CBO methods adopt the framework of Structural Causal Models (SCMs) (Pearl, 2009), which provide a compact, modular representation of how variables are generated and how they respond to interventions.

Definition 1 (Structural Causal Model (SCM)). *An SCM is a tuple $M = \langle U, V, F, P(U) \rangle$ where U is a set of exogenous variables, $V = \{V_1, \dots, V_n\}$ is a set of endogenous variables, and $F = \{f_1, \dots, f_n\}$ is a collection of structural assignments*

$$V_i := f_i(\text{Pa}_i, U_i), \quad i = 1, \dots, n, \quad (1)$$

where $\text{Pa}_i \subseteq V \setminus \{V_i\}$ denotes the direct causes (parents) of V_i in the causal graph and $U_i \subseteq U$ is the exogenous noise that affects V_i . The distribution $P(U)$ completes the model and, together with F , induces a unique joint distribution on V . The exogenous variables U are variables determined outside the modeled causal system; their probability distribution $P(U)$ is explicitly part of the SCM specification, so they are not “unmodeled” but rather externally determined.

Causal graphs, assumptions, and the causal Markov property. An SCM induces a causal graph $\mathcal{G}$, commonly represented as a directed acyclic graph (DAG), with a node for each endogenous variable and an edge $V_j \rightarrow V_i$ whenever $V_j \in \text{Pa}_i$. The graph encodes a qualitative causal structure in which variables are direct causes of others but does not, by itself, specify functional forms or effect magnitudes. That quantitative information resides in the structural equations F .

Two standard assumptions link the graph to the distribution. The *causal Markov condition* states that each variable is conditionally independent of its non-descendants given its parents (Pearl, 2009; Spirtes et al., 2000). Under the additional assumption that exogenous variables are mutually independent (*causal sufficiency*), the observational distribution factorizes according to the graph.

$$P(V_1, \dots, V_n) = \prod_{i=1}^n P(V_i \mid \text{Pa}_i). \quad (2)$$

The *faithfulness* assumption further requires that all conditional independencies in P are entailed by the Markov condition applied to $\mathcal{G}$, ruling out cancelations that could hide edges. Together, these assumptions enable structure learning from observational data and underpin the identifiability results that CBO exploits.

Hard interventions. In Pearl’s framework, a *hard* (or *perfect*) intervention $do(X=\mathbf{x})$ replaces the structural equation for X with the constant $\mathbf{x}$, severing all incoming edges to X while leaving other mechanisms unchanged. This produces a modified SCM $M_{do(X=\mathbf{x})}$ and an interventional distribution that admits a *truncated factorization*:

$$P(V_1, \dots, V_n \mid do(X=\mathbf{x})) = \prod_{V_i \notin X} P(V_i \mid \text{Pa}_i) \Big|_{X=\mathbf{x}}. \quad (3)$$

Hard interventions are the default in the original CBO formulation and in most Hard intervention benchmarks.

Soft interventions. A *soft* (or *imperfect*) intervention modifies the structural mechanism of a variable without completely severing its incoming edges. Formally, a soft intervention on V_i replaces f_i with a new mechanism $\tilde{f}_i(\text{Pa}_i, U_i; \theta)$ parameterized by an intervention parameter θ . In the broader literature, soft interventions can also encompass structural changes such as dropping or adding parent dependencies, or policy-based functional interventions where θ parameterizes a function mapping contexts to mechanism parameters; the formal summaries in Section 3 state what θ represents for each method. This is important in several CBO variants, notably MCBO (Sussex et al., 2023) and fCBO (Gultchin et al., 2023), that model mechanism modifications rather than variable clamping. Soft interventions yield a richer action space but require modeling how the intervention parameter interacts with the existing mechanism.

Counterfactuals. SCMs also define counterfactual quantities that describe outcomes under hypothetical actions in a specific latent world (indexed by $u \in U$), typically written as $Y_{\mathbf{x}}(u)$ (Pearl, 2009). The methods surveyed here optimize interventional expectations rather than counterfactuals; counterfactual CBO (for example, personalized optimization after observing unit-specific histories, or counterfactual safety constraints) remains a future direction.

Identifiability and the *do*-calculus. In practice, the full SCM is rarely known, so CBO relies on *identifiability*: whether an interventional quantity such as $\mathbb{E}[Y \mid do(X=\mathbf{x})]$ can be expressed purely in terms of the observational distribution given assumptions about the graph. Pearl’s *do*-calculus provides three general inference rules for transforming interventional expressions into observational quantities when identification

is possible (Pearl, 2009; 2012). A common special case is the *backdoor adjustment*: if a set Z satisfies the backdoor criterion relative to (X, Y) , then

$$\mathbb{E}[Y \mid \text{do}(X=\mathbf{x})] = \sum_z \mathbb{E}[Y \mid X=\mathbf{x}, Z=z] P(Z=z). \quad (4)$$

When identification holds but closed-form evaluation is difficult, one can approximate the required expectations via Monte Carlo integration of the identified estimand.

Identifiability as an algorithmic assumption in CBO. In CBO, identifiability is not merely a causal inference condition; it directly shapes the optimizer’s behavior. CBO benefits from identifiability to construct observationally informed priors, which accelerates sample efficiency, but the framework does not strictly rely on it to function. If a causal effect is not identifiable under the given graph, a sound algorithm can recognize this and fall back to a standard uninformative GP prior while still leveraging the causal graph for search-space reduction via POMIS. When an intervention effect is identifiable from observational data, the observational sample can be used to construct informative GP priors, initialize surrogates, or eliminate dominated intervention scopes *before* any costly intervention is performed. Two failure modes should be distinguished: if identifiability fails on a correctly specified graph, a sound adjustment procedure defaults to an uninformative prior, avoiding bias; overconfident, biased priors instead arise when the graph or adjustment assumptions themselves are wrong—a failure of the graph assumption rather than of identifiability. Empirical comparisons should therefore distinguish methods that assume correct identifiable priors from those that rely primarily on interventional data or explicitly model graph and identification uncertainty. Throughout this survey, we follow the common CBO assumption of causal sufficiency and perfect interventions unless otherwise stated.

2.2 Classical Bayesian Optimization

Bayesian Optimization (BO) is a sequential strategy for optimizing an expensive black-box objective $f : \mathcal{X} \rightarrow \mathbb{R}$ on a bounded domain $\mathcal{X} \subset \mathbb{R}^d$ (Kushner, 1964; Zhilinskias, 1975; Mockus et al., 1978; Mockus, 1989; Jones et al., 1998; Shahriari et al., 2016). At each iteration t , BO selects a query $\mathbf{x}_t$, observes $y_t = f(\mathbf{x}_t) + \epsilon_t$ with noise ϵ_t , and updates a probabilistic model to guide future queries. BO is characterized by two interacting components.

Probabilistic surrogate. The most common surrogate is a Gaussian Process (GP) (Rasmussen & Williams, 2006), specified by a prior mean function $m_0(\cdot)$ and a kernel (covariance function) $k(\cdot, \cdot)$. Given t observations $\mathcal{D}_t = \{(\mathbf{x}_i, y_i)\}_{i=1}^t$, the GP posterior of any candidate input $\mathbf{x}$ is available in closed form:

$$\mu_t(\mathbf{x}) = m_0(\mathbf{x}) + \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma_\epsilon^2 \mathbf{I})^{-1} (\mathbf{y}_t - \mathbf{m}_0), \quad (5)$$

$$\sigma_t^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}_t(\mathbf{x})^\top (\mathbf{K}_t + \sigma_\epsilon^2 \mathbf{I})^{-1} \mathbf{k}_t(\mathbf{x}), \quad (6)$$

where $\mathbf{k}_t(\mathbf{x})$ is the vector of covariances between $\mathbf{x}$ and the observed inputs, $\mathbf{K}_t$ is the kernel matrix over the observed inputs, and σ_ϵ^2 is the variance of the observation noise. The posterior mean μ_t provides a point estimate, while σ_t^2 quantifies the predictive uncertainty, enabling the acquisition function to distinguish well-explored from poorly-understood regions.

Acquisition functions. The acquisition function $\alpha(\mathbf{x}; \mathcal{D}_t)$ scores candidate inputs by balancing exploitation (questioning where μ_t is favorable) against exploration (questioning where predictive uncertainty or information gain is large). Two widely used choices are Expected Improvement (EI) and the Upper Confidence Bound (UCB):

$$\alpha_{\text{EI}}(\mathbf{x}) = \mathbb{E}[\max(f(\mathbf{x}) - f^+, 0) \mid \mathcal{D}_t], \quad (7)$$

$$\alpha_{\text{UCB}}(\mathbf{x}) = \mu_t(\mathbf{x}) + \beta_t^{1/2} \sigma_t(\mathbf{x}), \quad (8)$$

where f^+ is the incumbent value (in noisy settings, the best noisy observation y^+ or a posterior estimate of the latent incumbent) and $\beta_t > 0$ is a schedule parameter that controls the exploration–exploitation trade-off (Jones et al., 1998; Srinivas et al., 2010). The BO loop iterates: maximize α to select $\mathbf{x}_{t+1}$, evaluate f , and update the GP posterior.

Regret and sample efficiency. Under regularity conditions in the kernel and the appropriate choice of β_t , GP-UCB achieves a sublinear cumulative regret $\sum_{t=1}^T [f(\mathbf{x}^*) - f(\mathbf{x}_t)] = \mathcal{O}(\sqrt{T\gamma_T})$, where γ_T is the maximum information gain of the kernel (Srinivas et al., 2010; Chowdhury & Gopalan, 2017). This theoretical grounding is relevant to CBO because MCBO (Sussex et al., 2023) and ACBO (Sussex et al., 2024) extend GP-UCB-style regret analysis to causal settings, obtaining bounds that can improve over the non-causal case when the causal graph induces a factored reward structure.

Beyond standard GPs. Although GPs remain the dominant surrogate in CBO, the broader BO literature also uses random forests, Bayesian neural networks, neural processes, and deep kernel learning (Shahriari et al., 2016; Garnett, 2023). These alternatives are not causal by themselves. They become CBO components only when the modeled target is a causal estimand, when observational data enter through an identified causal prior, or when the architecture respects graph or mechanism structure. Thus, surrogate choice encodes both a statistical model class and a causal modeling target: effect-level surrogates model $\theta \mapsto \mathbb{E}[Y \mid \text{do}(X_s = \theta)]$, whereas mechanism-level surrogates model structural equations and propagate uncertainty through the graph.

2.3 Causal Bayesian Optimization

In CBO, the manipulable variables are linked to the target through structural mechanisms rather than through a black-box input–output map. This causal structure requires and enables algorithmic choices that standard BO does not face: which variables to intervene on (the intervention scope), at what level, and how abundant observational data relate to interventional effects through causal identification. This coupling of combinatorial scope selection with continuous value optimization, guided by a causal graph, is what distinguishes CBO from black-box BO.

Section 3.1 formalizes the CBO problem, defines the key concepts of minimal and possibly-optimal intervention sets, presents the surrogate and acquisition design, and identifies the limitations that motivate the extensions reviewed in the remainder of Section 3.

3 Methodological Frameworks

Causal Bayesian Optimization (CBO) is best viewed as a design space rather than a single algorithm. All CBO methods instantiate a common loop: maintain a probabilistic model of candidate interventional effects, score candidate interventions with an acquisition or decision rule, select the next intervention or observation action, and update beliefs using newly obtained observational and/or interventional data. What differentiates methods is where the causal structure is injected into this loop and which uncertainty sources are modeled.

A unifying lens: design axes and assumptions. In the literature, most variants can be understood through five interacting axes (Figure 1):

1. **Graph and system-knowledge assumptions.** Known graph, unknown graph, graph uncertainty, causal sufficiency, hidden confounding, and identifiability assumptions.
2. **Environmental assumptions.** Static, dynamic, contextual, constrained, adversarial or non-stationary, and multi-objective settings.
3. **Intervention representation.** Hard interventions, soft mechanism modifications, intervention scopes, functional policies, context-dependent actions, and mixed discrete–continuous actions.
4. **Surrogate architecture.** Direct effect-level surrogates, mechanism-level surrogates, graph-uncertain surrogates, multi-output models, and neural or deep-kernel variants.
5. **Decision rule and budget allocation.** EI-/UCB-style acquisition, information-gain search, optimism over plausible models, bandit allocation across scopes, and online-learning rules.

We use these axes to highlight what assumptions each method relaxes or introduces, what mathematical object changes relative to baseline CBO, and what statistical efficiency or computational cost follows.




Figure 1: Causal Bayesian Optimization development map organized by the primary bottleneck addressed.

How we summarize methods. For each approach, we emphasize: (i) assumptions (graph knowledge, intervention type, observed variables); (ii) the main algorithmic lever (search-space reduction, surrogate design, acquisition/decision policy); and (iii) failure modes and scaling constraints. This avoids re-stating paper narratives verbatim and instead clarifies how the methods relate at the level of design choices in the CBO loop. To make this comparison fully systematic, every surveyed method closes with a *formal summary* block whose four slots restate, as mathematics, the four columns of Table 2: (i) assumptions, (ii) objective, (iii) surrogate, and (iv) acquisition. Slots whose defining equation already appears in the text point to it rather than repeating it. Unless a statement is explicitly tied to benchmark evidence in Section 4, the “main challenge” and limitation remarks in this section are *analytical*: they follow from each method’s stated assumptions, surrogate level, and decision rule (see also the identification-strategy audit in Table 4), and should not be read as new empirical findings of this survey.

Common notation for method summaries. To make the method summaries precise and comparable, we use a single notation throughout this section: $\mathcal{G}$ denotes the causal graph, $X_s \subseteq X$ an intervention scope, θ an intervention value or policy parameter, Y the target outcome, and $\mathcal{D}_{\text{obs}}$ and $\mathcal{D}_{\text{int}}$ the observational and interventional datasets. Effect-level surrogates model the map $g_s(\theta) = \mathbb{E}[Y \mid \text{do}(X_s = \theta)]$ directly, whereas mechanism-level surrogates place a model $\hat{f}_j$ on each structural equation $f_j(\text{Pa}_j, U_j)$ and obtain the effect on Y by propagating samples or uncertainty through $\mathcal{G}$. Actions are written in one of three equivalent forms depending on the method family: hard scope-value interventions $(X_s, \mathbf{x}_s)$ (CBO, DCBO, cCBO, HCBO, MO-CBO, CEO, GACBO, CNEI), mechanism-shift actions $\theta \in \Theta$ that enter the structural equations as inputs (MCBO, ACBO), and policies π (fCBO, CoCaBO). Objectives follow the minimization convention of Definition 2 unless a method is inherently a maximization or regret formulation (cCBO, CoCaBO, GACBO, ACBO), in which case the conversion is immediate.

3.1 Causal Bayesian Optimization: Core Framework

The original CBO framework (Aglietti et al., 2020) formalizes causal optimization as a search over interventions in a known causal graph. Let X denote manipulable variables, C non-manipulable context variables, and Y the target outcome. An action consists of an intervention scope $X_s \subseteq X$ and an assignment $\mathbf{x}_s$ in the feasible intervention domain $\mathcal{D}(X_s)$.

Definition 2 (CBO objective and action space). *A CBO action is a pair $(X_s, \mathbf{x}_s)$ where $X_s \subseteq X$ is an intervention scope and $\mathbf{x}_s \in \mathcal{D}(X_s)$ is an assignment in the feasible intervention domain. The interventional objective is*

$$f(X_s, \mathbf{x}_s) = \mathbb{E}[Y \mid \text{do}(X_s = \mathbf{x}_s)]. \quad (9)$$

The global causal optimization problem is

$$(X_s^*, \mathbf{x}_s^*) \in \arg \min_{X_s \subseteq X, \mathbf{x}_s \in \mathcal{D}(X_s)} f(X_s, \mathbf{x}_s), \quad (10)$$

with $\arg \min$ replaced by $\arg \max$ for maximization tasks.

Definition 2 states the baseline CBO problem strictly according to its original formulation (Aglietti et al., 2020), which does not involve context variables. Context variables C are introduced only in Section 3.3 when discussing contextual extensions such as CoCaBO (Equation 28). The objective in Definition 2 is an interventional expectation: none of the methods surveyed in this section optimizes counterfactual quantities $Y_{\mathbf{x}}(u)$, which we treat only as a future direction (Section 2.1).

The key point is that the optimizer is free to choose the scope of the intervention, not only the level. This yields a mixed discrete–continuous search problem and motivates two core ideas in CBO: graph-based pruning of candidate scopes and observationally informed priors over interventional effects.

Why causality changes the BO problem. A naive approach would flatten $(X_s, \mathbf{x}_s)$ into a single input vector and run the standard BO on the interventional objective. This ignores the structural information encoded in the causal graph. Causal knowledge enables two advantages that standard BO cannot exploit: (i) *scope reduction*: identifying which variables need to be manipulated at all, thereby shrinking the combinatorial search; and (ii) *informative priors*: converting plentiful observational data into prior beliefs about interventional effects, which is possible only when identifiability holds (Section 2.1).

Scope reduction: minimal intervention sets. A central insight of CBO is that the causal graph can prune intervention scopes that are structurally redundant. Two key notions formalize this:

Definition 3 (Minimal Intervention Set (MIS) (Lee & Bareinboim, 2018; 2019)). *An intervention scope $X_s \subseteq X$ is a minimal intervention set for the target Y if (i) there exists a directed path from at least one variable in X_s to Y in the causal graph $\mathcal{G}$ that is not blocked by the intervention on the remaining variables in X_s , and (ii) no proper subset $X'_s \subsetneq X_s$ satisfies condition (i). Equivalently, X_s is a minimal set of manipulable variables whose joint intervention can affect Y through the graph.*

Definition 4 (Possibly-Optimal MIS (POMIS) (Lee & Bareinboim, 2018; 2019)). *A MIS X_s is possibly optimal if, given the observational distribution $P(V)$ and the causal graph $\mathcal{G}$, one cannot determine from observational evidence alone that $\max_{\mathbf{x}_s} \mathbb{E}[Y \mid \text{do}(X_s = \mathbf{x}_s)]$ is dominated by the optimal objective achievable under some other MIS $X_{s'}$. Formally, the set of POMIS is the subset of MISs that remains after eliminating scopes whose optimal interventional value is provably dominated under the available graph and distributional information.*

By restricting the search to POMIS, CBO avoids wasting budget on scopes that are causally irrelevant (no path to Y) or observationally dominated (another scope provably achieves a better objective). This reduction is not merely computational; it changes the statistical problem by concentrating evaluations on scopes that the graph predicts can meaningfully move the outcome, and it is one of the main reasons CBO can be statistically more efficient than standard BO. It is also complementary to causal abstraction, which relates decisions across graphs at different granularities (Section 3.7).

Surrogate modeling: effect-level GPs with causal priors. The CBO models the interventional objective for each candidate scope with an effect-level GP. The key innovation is the prior construction: when the interventional effect is identifiable, observational data are used to build an initial estimate of the interventional mean and its uncertainty:

$$g_s(\mathbf{x}_s) := \mathbb{E}[Y \mid \text{do}(X_s = \mathbf{x}_s)], \quad (11)$$

$$g_s(\cdot) \sim \mathcal{GP}(m_s(\cdot), k_s(\cdot, \cdot)), \quad (12)$$

$$m_s(\mathbf{x}_s) \approx \widehat{\mathbb{E}}[Y \mid \text{do}(X_s = \mathbf{x}_s)], \quad (13)$$

$$k_s(\mathbf{x}_s, \mathbf{x}'_s) = k_{\text{RBF}}(\mathbf{x}_s, \mathbf{x}'_s) + \widehat{\sigma}_s(\mathbf{x}_s) \widehat{\sigma}_s(\mathbf{x}'_s), \quad (14)$$

Algorithm 1 Generic Causal Bayesian Optimization Loop**Require:** Causal graph $\mathcal{G}$ (or prior over graphs), observational data $\mathcal{D}_{\text{obs}}$, budget T **Ensure:** Best intervention $(X_s^*, \mathbf{x}_s^*)$

- 1: **Scope reduction** [causal]: Use $\mathcal{G}$ to identify candidate scopes (MIS/POMIS)
- 2: **Prior construction** [causal]: Use $\mathcal{D}_{\text{obs}}$ and identifiability to set GP prior means $\{m_s\}$
- 3: **for** $t = 1, \dots, T$ **do**
- 4: **Surrogate update:** Update GP posteriors $\{\mu_t^{(s)}, \sigma_t^{(s)}\}$ for each scope s
- 5: **Acquisition** [causal]: Select $(X_s, \mathbf{x}_s) = \arg \max_{s, \mathbf{x}} \alpha_s(\mathbf{x}; \mathcal{D}_t)$ across scopes
- 6: **Evaluate:** Observe $y_t = f(X_s, \mathbf{x}_s) + \epsilon_t$ via intervention
- 7: **Update:** Add $(X_s, \mathbf{x}_s, y_t)$ to interventional dataset $\mathcal{D}_t$
- 8: **end for**
- 9: **return** $(X_s^*, \mathbf{x}_s^*)$ with best observed objective

where $k_{\text{RBF}}(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2 / (2\ell^2))$ and $\hat{\sigma}_s(\mathbf{x}_s)$ is the uncertainty in the observationally-derived effect estimate. The observational data thus serve as prior information, not as additional labels, giving the surrogate a reasonable global shape before costly interventions begin.

Acquisition and the observation–intervention trade-off. CBO uses a cost-aware expected improvement to choose between pairs (scope, value):

$$\text{CEI}_s(\mathbf{x}_s) = \frac{\mathbb{E}[(y^+ - g_s(\mathbf{x}_s))_+]}{\text{Co}(X_s, \mathbf{x}_s)}, \quad (15)$$

where y^+ is the incumbent (the best value observed so far; we reserve y^* for the reference optimum used in the evaluation metrics of Section 4) and $\text{Co}(\cdot)$ is an intervention cost. Beyond the standard exploration–exploitation trade-off, CBO introduces an *observation–intervention trade-off*: the algorithm can decide whether to allocate budget to additional observational samples (cheap but potentially biased for interventional decisions) or to new interventions (expensive but causally informative). This framing is foundational: many later variants can be interpreted as changing *what is uncertain* and therefore changing *which data is worth acquiring*.

Formal summary (CBO).

$$(\text{assumptions}) \quad \mathcal{G} \text{ known DAG; } X_s \in \text{POMIS}(\mathcal{G}) \text{ (Defs. 3–4); } m_s \text{ identifiable from } \mathcal{D}_{\text{obs}} \quad (16a)$$

$$(\text{objective}) \quad \text{global causal optimum over hard interventions } (X_s, \mathbf{x}_s), \text{ Equation 10} \quad (16b)$$

$$(\text{surrogate}) \quad \text{effect-level GP with causal prior mean and inflated kernel, Equation 12} \quad (16c)$$

$$(\text{acquisition}) \quad \text{cost-aware causal EI with observe-or-intervene choice, Equation 15} \quad (16d)$$

The generic CBO loop. At each iteration, the CBO loop proceeds in three stages: (i) the surrogate predicts the causal effect for each candidate intervention (scope, value), (ii) the acquisition function scores candidates using the posterior predictions, and (iii) the highest-scoring intervention is selected and evaluated. Algorithm 1 summarizes the generic CBO procedure, highlighting where causal knowledge enters. Steps 1–2 exploit the graph *before* the optimization loop; Steps 4–5 use a causal structure *within* each iteration. This template applies, with variations, to all methods reviewed in the remainder of this section.

Limitations that motivate extensions. The modularity of CBO also exposes its bottlenecks: (i) reliance on a correctly known graph and identifiable observational priors; (ii) limited robustness to noise and prior misspecification; and (iii) scaling challenges due to exponentially many candidate scopes and separate surrogates per scope. The remainder of this section can be read as a sequence of responses to these bottlenecks, organized by the design axis that each method primarily addresses.

3.2 Modeling Extensions

The original CBO formulation assumes a static causal system, a single objective, and an intervention space that remains tractable after graph-based pruning. Modeling extensions relax these assumptions by changing what the optimizer must represent: time variation that renders the objective non-stationary, feasibility that restricts attention to a safe set, dimensionality that enlarges the scope of the space, or multiple competing goals that shift the target from a single optimum to a Pareto set, in each case keeping the same surrogate-acquisition loop while adding structure that decomposes the problem into smaller subproblems that share statistical strength.

Time-evolving causal systems (DCBO). Dynamic Causal Bayesian Optimization (DCBO) addresses environments where the best intervention is non-stationary, typically because mechanisms drift or the system state evolves (Aglietti et al., 2021). The main contribution is not a new acquisition rule, but a transfer mechanism. Under invariance assumptions such as a fixed graph and additive-noise structure, interventional knowledge from earlier time steps can be mapped into an informative prior for the current step. This treats time as a source of reusable evidence so that exploration can be amortized rather than restarted. The trade-off is the sensitivity of the assumption. If the graph or key mechanisms shift in ways not captured by the transfer recursion, the prior can become overconfident and bias the search, so DCBO is strongest when temporal variation is structured and not adversarial.

Formal summary (DCBO).

$$(assumptions) \quad \mathcal{G}_t = \mathcal{G}_0 \quad \forall t; \quad Y_t = f_Y^p(\text{Pa}^T(Y_t)) + f_Y^{np}(\text{Pa}^{NT}(Y_t)) + \epsilon_t; \quad \text{no hidden confounders in } \mathcal{G}_{0:T} \quad (17a)$$

$$(objective) \quad (X_{s,t}^*, \mathbf{x}_{s,t}^*) \in \arg \min_{X_{s,t} \subseteq X_t, \mathbf{x} \in \mathcal{D}(X_{s,t})} \mathbb{E}[Y_t \mid do(X_{s,t}=\mathbf{x}), I_{0:t-1}], \quad I_{0:t-1} = \bigcup_{i=0}^{t-1} do(X_{s,i}^*=\mathbf{x}_{s,i}^*) \quad (17b)$$

$$(surrogate) \quad g_{s,t}(\mathbf{x}) = f_Y^p(\mathbf{f}_{t-1}^*) + \mathbb{E}_{p(\mathbf{w} \mid do(X_{s,t}=\mathbf{x}), I_{0:t-1})} [f_Y^{np}(\mathbf{x}, \mathbf{w})], \quad g_{s,t} \sim \mathcal{GP}(m_{s,t}, k_{s,t}) \quad (17c)$$

$$(acquisition) \quad \text{CEI}_{s,t}(\mathbf{x}) = \frac{\mathbb{E}[(y_t^+ - g_{s,t}(\mathbf{x}))_+]}{\text{Co}(X_{s,t}, \mathbf{x})} \quad (17d)$$

Here $\text{Pa}^T(Y_t) \subseteq \{Y_0, \dots, Y_{t-1}\}$ and $\text{Pa}^{NT}(Y_t)$ partition the parents of Y_t into past targets and non-target variables, $\mathbf{f}_{t-1}^* = \{\mathbb{E}[Y_i \mid do(X_{s,i}^*=\mathbf{x}_{s,i}^*), I_{0:i-1}]\}_{Y_i \in \text{Pa}^T(Y_t)}$ collects the optimal targets achieved at earlier steps, $\mathbf{w}$ denotes the remaining non-intervened variables, and y_t^+ is the best target value observed at time t across scopes. The additive split in Equation 17a is exactly what licenses the prior-transfer recursion in Equation 17c: the prior mean $m_{s,t}$ is obtained by plugging $\mathbf{f}_{t-1}^*$ and observational estimates into Equation 17c, and the kernel $k_{s,t}$ mirrors the static causal prior of Equation 12.

Constraints and safe causal optimization (cCBO). Constrained Causal Bayesian Optimization (cCBO) targets settings where interventions must satisfy safety or feasibility constraints (Aglietti et al., 2023). The key insight from the modeling is that constraints change the definition of a useful intervention scope. A scope that can move the objective but cannot maintain feasibility is effectively irrelevant; cCBO therefore extends graph-based pruning to constraint-relevant pathways through constrained minimal intervention sets, and it couples objective and constraints through multi-output surrogates so evidence about feasibility shapes where the objective surrogate should be trusted. Formally, cCBO solves

$$(X_s^*, \mathbf{x}_s^*) \in \arg \max_{X_s, \mathbf{x}_s} \mathbb{E}[Y \mid do(X_s=\mathbf{x}_s)] \quad \text{s.t.} \quad \Pr(c_j(X_s, \mathbf{x}_s) \leq 0 \quad \forall j) \geq 1-\delta, \quad (18)$$

where each c_j is a constraint function and δ controls the tolerated feasibility violation. In the original formulation, the constraints are stated as thresholds on interventional constraint means, $\mathbb{E}[C_j \mid do(X_s=\mathbf{x}_s)] \geq \lambda_j$; Equation 18 recovers this form with $c_j(X_s, \mathbf{x}_s) := \lambda_j - \mathbb{E}[C_j \mid do(X_s=\mathbf{x}_s)]$. In practice, this often acts as a regularizer that prevents spending the budget in regions that look promising under the objective surrogate but are unlikely to be deployable. The main limitation is scaling with the number of constraints and candidate scopes, although structural pruning mitigates this in graphs with sparse causal connectivity.

Formal summary (cCBO).

$$\begin{aligned} (\text{assumptions}) \quad & \mathcal{G} \text{ known DAG; constraint variables } C_1, \dots, C_J \text{ embedded in } \mathcal{G}; \\ & \text{scopes pruned to constraint-aware minimal intervention sets} \end{aligned} \quad (19a)$$

$$(\text{objective}) \quad \text{constrained interventional optimum, Equation 18} \quad (19b)$$

$$(\text{surrogate}) \quad (g_s, c_{s,1}, \dots, c_{s,J}) \sim \mathcal{GP}(\mathbf{m}_s, \mathbf{K}_s), \quad [\mathbf{K}_s]_{ab} = \text{Cov}(h_a(\mathbf{x}_s), h_b(\mathbf{x}'_s)) \quad (19c)$$

$$(\text{acquisition}) \quad \text{cEI}_s(\mathbf{x}_s) = \frac{\mathbb{E} \left[(g_s(\mathbf{x}_s) - y_{\text{feas}}^+) + \prod_{j=1}^J \mathbb{I}\{c_{s,j}(\mathbf{x}_s) \leq 0\} \right]}{\text{Co}(X_s, \mathbf{x}_s)} \quad (19d)$$

In Equation 19c, $h_0 = g_s$ and $h_j = c_{s,j}$ index the outputs of the multi-output GP. The expectation in Equation 19d is taken under this joint posterior, whose cross-output covariances are either learned (multi-task GP) or derived from the shared structural equations under intervention ($\mathcal{G}$ -informed multi-task GP); y_{feas}^+ is the best *feasible* incumbent in $\mathcal{D}_{\text{int}}$, and the cost normalization $\text{Co}(\cdot)$ replaces the scope-size normalization $|X_s|$ used in the original paper, for uniformity with Equation 15.

High-dimensional intervention spaces (HCBO). High-Dimensional Causal Bayesian Optimization (HCBO) addresses the combinatorial bottleneck of scope selection in large graphs (Wu et al., 2024). Its central move is to replace exact enumeration of minimal or possibly-optimal scopes with an efficient coverage-based approximation. Specifically, HCBO selects a scope family $\mathcal{S}^* \subseteq 2^X$ by solving a coverage problem that approximately maximizes the fraction of causal pathways to Y that are “covered” by at least one scope in $\mathcal{S}^*$. Let $\text{Paths}(\mathcal{G}, Y)$ denote the set of directed paths from manipulable variables to Y . HCBO seeks

$$\mathcal{S}^* \in \arg \max_{\mathcal{S} \subseteq 2^X, |\mathcal{S}| \leq K} |\{p \in \text{Paths}(\mathcal{G}, Y) : \exists X_s \in \mathcal{S}, X_s \cap p \neq \emptyset\}|, \quad (20)$$

where K is a user-specified scope budget. This reframes high-dimensional CBO as a hierarchical problem: select a promising scope family, then run continuous BO within the chosen scopes. HCBO also introduces normalized cross-scope scoring so that surrogates trained on different scopes can be compared in a meaningful way despite scale and data imbalance. The cost of scalability is that any coverage proxy can miss narrow, high-impact pathways, so HCBO is most reliable when the system exhibits causal sparsity or a low intrinsic causal dimension.

Formal summary (HCBO).

$$\begin{aligned} (\text{assumptions}) \quad & \mathcal{G} \text{ known DAG; causal sparsity: the intrinsic causal dimension is small} \\ & \text{(the minimal covering } K \text{ in Equation 20)} \end{aligned} \quad (21a)$$

$$(\text{objective}) \quad \text{Equation 10 restricted to the scope family } \mathcal{S}^* \text{ of Equation 20} \quad (21b)$$

$$(\text{surrogate}) \quad \text{per-scope effect-level GPs (Equation 12) with posteriors } (\mu_t^{(s)}, \sigma_t^{(s)}) \quad (21c)$$

$$(\text{acquisition}) \quad (X_s, \mathbf{x}_s)_{t+1} \in \arg \max_{X_s \in \mathcal{S}^*, \mathbf{x}_s \in \mathcal{D}(X_s)} \frac{\mu_t^{(s)}(\mathbf{x}_s) + \beta_t^{1/2} \sigma_t^{(s)}(\mathbf{x}_s) - b_t^{(s)}}{r_t^{(s)}} \quad (21d)$$

Equation 21d states the cross-scope scoring rule in schematic form: $b_t^{(s)}$ and $r_t^{(s)}$ are per-scope location and scale statistics that place UCB scores from surrogates fitted on different scopes and data volumes on a common scale, following the normalized scope-comparison scheme of Wu et al. (2024).

Multiple outcomes and Pareto structure (MO-CBO). Multi-Objective Causal Bayesian Optimization (MO-CBO) generalizes CBO to settings with multiple targets, where the goal is to approximate a Pareto front rather than a single optimum (Bhatija et al., 2025). The key observation is that causal structure is especially valuable when objectives compete. Intervening on everything can obscure trade-offs by coupling variables that need not be coupled. MO-CBO uses the graph to reduce the search to a minimal collection of local multi-objective subproblems (Equation 22b) and then allocates evaluation budget across them using a cross-scope criterion based on relative hypervolume improvement (Equation 22d). Within each local

problem, standard multi-objective BO components can be used, for example, DGEMO (Konakovic Lukovic et al., 2020). The main challenge remaining is combinatorial. The number of local subproblems can still grow quickly with dense graphs or many objectives, motivating stronger cross-scope sharing and joint surrogates that exploit causal factorization across objectives.

Formal summary (MO-CBO).

$$(assumptions) \quad \mathcal{G} \text{ known multi-target DAG; } \mathbf{Y} = (Y_1, \dots, Y_m); \quad \boldsymbol{\mu}(X_s, \mathbf{x}_s) = (\mathbb{E}[Y_i \mid do(X_s = \mathbf{x}_s)])_{i=1}^m \quad (22a)$$

$$(objective) \quad \mathcal{P}^c(\mathcal{S}) = \{\boldsymbol{\mu}(X_s, \mathbf{x}_s) : \nexists (X_{s'}, \mathbf{x}_{s'}) \text{ with } \boldsymbol{\mu}(X_{s'}, \mathbf{x}_{s'}) \prec \boldsymbol{\mu}(X_s, \mathbf{x}_s)\}, \quad \mathcal{P}^c(\mathcal{S}) \subseteq \bigcup_{X_s \in \mathcal{S}} \mathcal{P}^\ell(X_s) \quad (22b)$$

$$(surrogate) \quad \hat{\mu}_i^{(s)} \sim \mathcal{GP}(m_i^{(s)}, k_i^{(s)}) \quad \text{independently for each objective } i \text{ and each scope } X_s \in \mathcal{S} \quad (22c)$$

$$(acquisition) \quad \text{RHVI}(X_s, B_s) = \frac{\mathcal{H}(\mathcal{P}^\ell(X_s) \cup \boldsymbol{\mu}(X_s, B_s)) - \mathcal{H}(\mathcal{P}^\ell(X_s))}{\mathcal{H}(\mathcal{P}^\ell(X_s))}, \quad B^* \in \arg \max_{X_s \in \mathcal{S}, B_s} \text{RHVI}(X_s, B_s) \quad (22d)$$

Here $\prec$ denotes Pareto dominance under minimization (weakly smaller in every coordinate of $\boldsymbol{\mu}$, strictly smaller in at least one), $\mathcal{S}$ is the family of possibly-Pareto-optimal minimal intervention sets (the multi-target analogue of POMIS, characterized graphically in Bhatija et al., 2025), $\mathcal{P}^\ell(X_s)$ is the local Pareto front of the subproblem restricted to scope X_s , B_s is a candidate evaluation batch, and $\mathcal{H}$ is the hypervolume indicator with respect to a fixed reference point.

3.3 Intervention Strategy Developments

Beyond modeling extensions, another major line of work changes what an intervention looks like and how decisions are made during optimization. The common motivation is that treating each intervention as a single black-box query can waste information. These methods modify the BO loop so it can reuse causal structure inside the system, express richer actions such as policies, and stay meaningful under contextual variation or non-stationarity.

Two orthogonal design choices appear in this section and should not be conflated: expanding the *intervention representation* (fCBO, CoCaBO) and changing the *surrogate architecture* (MCBO, ACBO) are independent axes—mechanism-level surrogates can be paired with only hard interventions, and effect-level surrogates with soft interventions. Relatedly, which uncertainty dominates a method is not a design choice but a property *induced* by the system-knowledge assumptions (known versus unknown graph, causal sufficiency, known feasible set); the taxonomy therefore lists system-knowledge assumptions as the primary axis, and the discussion below identifies the induced uncertainty for each method.

Two further clarifications: CoCaBO and ACBO also differ on the *environmental* axis (contextual versus adversarial or non-stationary), so grouping them here reflects shared machinery, not identical assumptions. And a closely related formulation outside the SCM tradition is Bayesian optimization of function networks (Astudillo & Frazier, 2021), which optimizes over a known computational DAG with observed intermediate values; it anticipates the mechanism-level, soft-intervention view of MCBO and ACBO, its benchmarks are inherited by our soft-intervention benchmark (Section 4.5), and its key difference from SCM-based CBO is that function-network inputs are causally independent decision variables rather than variables embedded in a causal system.

Mechanism-level modeling with guarantees (MCBO). Model-based Causal Bayesian Optimization (MCBO) models the causal system at the level of structural mechanisms rather than only the intervention-to-outcome effect (Sussex et al., 2023). This changes data reuse: observations of intermediate variables update upstream mechanisms and therefore tighten uncertainty about downstream outcomes, which can reduce the number of costly interventions needed to identify good actions. The main limitations are practical. Mechanism posteriors must be maintained for multiple equations, uncertainty propagation increases computational cost, and acquisition optimization can be challenging. MCBO also assumes a known graph and a mechanism model class that is expressive enough for the system at hand.

Formally, each mechanism receives an independent GP prior, $\hat{f}_j \sim \mathcal{GP}(m_j, k_j)$, and the outcome of an intervention $\boldsymbol{\theta}$ (a hard clamp or soft mechanism shift) is obtained by propagating through the graph in

topological order:

$$V_j = \hat{f}_j(\text{Pa}_j(\boldsymbol{\theta}), U_j) \text{ for each non-intervened node, } \hat{Y}(\boldsymbol{\theta}) = \hat{f}_Y(\text{Pa}_Y(\boldsymbol{\theta}), U_Y), \quad (23)$$

where each parent value $\text{Pa}_j(\boldsymbol{\theta})$ is itself generated by upstream mechanism models under the intervention. The decision rule is optimism-based over a calibrated set of plausible mechanism vectors:

$$\mathcal{F}_t = \left\{ \tilde{F} = (\tilde{f}_j)_j : \tilde{f}_j \in \mathcal{H}_{k_j}, \|\tilde{f}_j\|_{k_j} \leq B_j, |\tilde{f}_j(\cdot) - \mu_{j,t-1}(\cdot)| \leq \beta_t \sigma_{j,t-1}(\cdot) \forall j \right\}, \quad (24)$$

$$\boldsymbol{\theta}_t \in \arg \max_{\boldsymbol{\theta} \in \Theta} \text{ucb}_t(\boldsymbol{\theta}), \quad \text{ucb}_t(\boldsymbol{\theta}) = \max_{\tilde{F} \in \mathcal{F}_t} \mathbb{E}_U[\hat{Y}(\boldsymbol{\theta}) | \tilde{F}], \quad (25)$$

where $(\mu_{j,t-1}, \sigma_{j,t-1})$ are the mechanism-GP posteriors, β_t is calibrated so that the true mechanism vector lies in $\mathcal{F}_t$ with probability at least $1 - \delta$, and the inner maximization is made tractable by the reparameterization $\tilde{f}_j = \mu_{j,t-1} + \beta_t \sigma_{j,t-1} \eta_j$ with $\eta_j(\cdot) \in [-1, 1]$. This optimism-based rule, applied to the propagated outcome of Equation 23, yields the cumulative-regret guarantees established in [Sussex et al. \(2023\)](#).

Compared to effect-level CBO, MCBO has two additional trade-offs that deserve explicit discussion. First, maintaining $|V|$ separate mechanism-level GPs and propagating posterior uncertainty through the full graph at every acquisition step is computationally far more expensive than fitting a single effect-level GP, limiting current implementations to graphs with tens of nodes (see also Section 5.2). Second, mechanism-level models are highly vulnerable to unobserved confounders between intermediate nodes: if a hidden confounder links two internal variables, the GP for the downstream mechanism will be fitted on systematically biased inputs, and this bias propagates through the graph. Effect-level CBO can mitigate the same issue more simply by discarding the biased observational prior and proceeding with an uninformative GP prior.

A related nuance concerns intermediate observations under effect-level models. The claim that effect-level models “cannot reuse intermediate observations” is only strictly true for the GP posterior itself: a direct $X_s \rightarrow Y$ surrogate does not condition on intermediate node values. However, intermediate observations collected during interventional trials can be appended to the observational dataset and used to re-estimate or refine the observationally-derived prior mean $m_s(\cdot)$ at each step. This indirect reuse does not provide the same per-mechanism uncertainty reduction as MCBO, but it is a practical pathway that existing effect-level implementations could exploit.

Formal summary (MCBO).

(assumptions) $\mathcal{G}$ known DAG; causal sufficiency (independent noises U_j);

mechanisms $f_j \in \mathcal{H}_{k_j}$ in the plausible set $\mathcal{F}_t$ of Equation 24 (26a)

(objective) $\max_{\boldsymbol{\theta} \in \Theta} \mathbb{E}[Y | \boldsymbol{\theta}]$ (hard clamps or soft mechanism shifts; static environment) (26b)

(surrogate) mechanism-level GPs propagated through $\mathcal{G}$ in topological order, Equation 23 (26c)

(acquisition) optimistic causal UCB over $\mathcal{F}_t$ with regret guarantees, Equation 25 (26d)

Policies as interventions (fCBO). Functional Causal Bayesian Optimization (fCBO) expands the action space from selecting a value to selecting a policy that maps the observed context to an action ([Gultchin et al., 2023](#)). This is motivated by domains such as personalized treatment, where a single fixed intervention is rarely optimal for all individuals. fCBO makes policy search tractable by defining a kernel over policies using RKHS tools and running Bayesian optimization in this function space with a functional expected improvement rule. The key contribution is that it unifies causal optimization and policy learning within a single sequential procedure, where the surrogate models policy-to-outcome effects and the acquisition balances exploration across qualitatively different decision rules. The main challenge is representation. Because the policy space is large, empirical performance depends strongly on the chosen parameterization and on whether the kernel captures a meaningful similarity between policies.

Formal summary (fCBO).

$$\begin{aligned} (\text{assumptions}) \quad & \mathcal{G} \text{ known DAG; functional interventions } do(X_s = \pi(Z)) \text{ replace the mechanism of } X_s; \\ & \pi \text{ in a bounded ball } \Pi_s \subset \mathcal{H}_{\kappa_s} \end{aligned} \quad (27a)$$

$$(\text{objective}) \quad \pi^* \in \arg \min_{\pi \in \Pi_s} \mathbb{E}[Y \mid do(X_s = \pi(Z))] \quad \text{jointly with } X_s \text{ over non-redundant scopes} \quad (27b)$$

$$(\text{surrogate}) \quad g_s(\pi) \sim \mathcal{GP}(m_s(\pi), k_s(\pi, \pi')), \quad k_s(\pi, \pi') = k_{\text{RBF}}(\|\pi - \pi'\|_{\mathcal{H}_{\kappa_s}}) \quad (27c)$$

$$(\text{acquisition}) \quad \text{fEI}_s(\pi) = \frac{\mathbb{E}[(y^+ - g_s(\pi))_+]}{\text{Co}(X_s, \pi)}, \quad (X_s, \pi)_{t+1} \in \arg \max_{X_s, \pi \in \Pi_s} \text{fEI}_s(\pi) \quad (27d)$$

The original formulation optimizes jointly over scope and policy and allows mixed scopes that combine hard components with functional components, $\Pi_s = \mathbb{R}^{|s_{\text{hard}}|} \times \mathcal{B}(\mathcal{H}_{\kappa_s})$; correspondingly, the policy kernel in Equation 27c is in general a composite of a Euclidean distance on hard components and the RKHS distance on functional components (Gultchin et al., 2023).

Context-dependent scope selection (CoCaBO). Contextual Causal Bayesian Optimization (CoCaBO) targets settings where the exogenous context affects outcomes and where the best intervention depends on the realized context (Arsenyan et al., 2026). The additional difficulty is that one must decide not only which action to take, but also which contextual variables should be used to condition decisions. CoCaBO separates these levels: an outer bandit allocates budget across candidate *mixed policy scopes* $\mathcal{M}_k = (X_{s_k}, C_{s_k})$, each pairing an intervention scope with a context subset, while an inner mixed-variate Bayesian optimization routine such as HEBO (Cowen-Rivers et al., 2022) optimizes within the selected scope. CoCaBO optimizes a context-dependent objective of the form

$$\max_{\pi: \mathcal{C} \rightarrow \mathcal{X}} \mathbb{E}_C[\mathbb{E}[Y \mid do(X_{s_k} = \pi(C_{s_k})), C]], \quad (28)$$

where π maps observed context values to intervention values within the chosen scope, and the outer expectation is over the context distribution. This architecture makes explicit that context introduces a discrete model selection problem on top of continuous optimization. The main limitations are derived from the hierarchy. The outer allocation must spend trials to distinguish scopes, and improvements learned within one scope do not necessarily transfer to others.

Formal summary (CoCaBO).

$$\begin{aligned} (\text{assumptions}) \quad & \mathcal{G} \text{ known DAG; observed exogenous context } C; \\ & \text{actions are mixed policy scopes } \mathcal{M}_k = (X_{s_k}, C_{s_k}) \text{ pruned to possibly-optimal scopes} \end{aligned} \quad (29a)$$

$$(\text{objective}) \quad \text{contextual policy objective, Equation 28} \quad (29b)$$

$$(\text{surrogate}) \quad g_k(\mathbf{x}, \mathbf{c}) \sim \mathcal{GP}(0, k_k((\mathbf{x}, \mathbf{c}), (\mathbf{x}', \mathbf{c}')))) \quad \text{for each candidate scope } \mathcal{M}_k \quad (29c)$$

$$(\text{acquisition}) \quad k_t \in \arg \max_k \left[\bar{U}_k(t) + \rho_k(n_k(t)) / \sqrt{n_k(t)} \right], \quad \text{then mixed-variate BO within } \mathcal{M}_{k_t} \quad (29d)$$

In Equation 29d, $n_k(t)$ is the number of rounds allocated to $\mathcal{M}_k$, $\bar{U}_k(t)$ is the running average of the inner optimizer’s UCB values under $\mathcal{M}_k$ (the arm statistic averages optimistic scores rather than raw rewards), and ρ_k is an exploration width chosen to obtain $\tilde{O}(\sqrt{mT})$ cumulative regret over m candidate scopes (Arsenyan et al., 2026).

Non-stationarity and adversarial environments (ACBO). Adversarial Causal Bayesian Optimization (ACBO) addresses cases where the system response changes over time or is influenced by an adversary that can also intervene in the system (Sussex et al., 2024). ACBO combines multiplicative weight updates from adversarial online learning (Littlestone & Warmuth, 1994; Freund & Schapire, 1997) with causal modeling to construct calibrated optimistic reward estimates via a causal UCB oracle. A central benefit of using the causal graph is that it can reduce regret dependence on the nominal action dimension when rewards

factor through causal compositions, improving over non-causal analogs such as GP-MW (Sessa et al., 2019). The main challenges are computational and statistical. The oracle can be expensive to optimize, and in non-stationary regimes, the method relies on uncertainty sets that must track drift accurately to avoid misleading optimism.

Formally, writing θ'_t for the adversary’s action at round t and $r(\theta, \theta') = \mathbb{E}[Y \mid \theta, \theta']$ for the expected reward, ACBO targets the adversarial regret

$$R(T) = \max_{\theta \in \Theta} \sum_{t=1}^T r(\theta, \theta'_t) - \sum_{t=1}^T r(\theta_t, \theta'_t), \quad (30)$$

and selects actions from a multiplicative-weights distribution over a discretized action set Θ , fed by optimistic causal-UCB reward estimates:

$$\theta_t \sim w_t, \quad w_{t+1}(\theta) \propto w_t(\theta) \exp\left(\tau \text{UCB}_t^{\mathcal{G}}(\theta, \theta'_t)\right), \quad \text{UCB}_t^{\mathcal{G}}(\theta, \theta') = \max_{\tilde{F} \in \mathcal{F}_t} \mathbb{E}_U[Y \mid \tilde{F}, \theta, \theta'], \quad (31)$$

with learning rate $\tau = \sqrt{8 \log |\Theta| / T}$ and the plausible mechanism set $\mathcal{F}_t$ of Equation 24, whose mechanisms now take the adversary’s action as an additional input; the inner maximization uses the same mechanism-confidence reparameterization as MCBO (Sussex et al., 2024). No-regret means $R(T)/T \rightarrow 0$ for any adversary sequence.

Formal summary (ACBO).

- (assumptions) $\mathcal{G}$ known DAG; mechanism-level model class and plausible set $\mathcal{F}_t$ (Equation 24),
extended by the adversary input θ' (32a)
- (objective) no-regret against the best fixed intervention, Equation 30 (32b)
- (surrogate) mechanism-level propagation, Equation 23, inputs extended by (θ, θ') (32c)
- (acquisition) multiplicative weights with optimistic causal-UCB rewards, Equation 31 (32d)

3.4 Unknown causal structure

A major practical barrier for Causal Bayesian Optimization is that many applications do not provide a trusted causal graph in advance. In this regime, the optimization loop must allocate a limited intervention budget across two coupled goals: improving the outcome and reducing the ambiguity about which causal explanations remain consistent with the data. This introduces an additional uncertainty source, structure uncertainty, on top of uncertainty about causal mechanisms and observation noise.

To keep terminology consistent throughout the paper, we use *unknown graph* for the assumption class in which the learner is not given the causal graph; *graph uncertainty* (used interchangeably with *structure uncertainty*) for the learner’s remaining uncertainty over graph structure, however it is represented (a posterior, a confidence set, or a candidate list); and *graph misspecification* for the distinct situation in which the learner is given, and trusts, a single graph that differs from the true one. The methods in this subsection address the first two; graph misspecification is instead probed by the stress tests in Appendix B.1. Existing approaches mainly differ in how they represent uncertainty over candidate graphs and in whether they explicitly target the decision-relevant structure, that is, the parts of the graph that materially change which intervention should be selected.

Information-seeking joint learning and optimization (CEO). Causal Entropy Optimization, CEO, couples structure learning and optimization using an explicitly information-seeking decision rule (Branchini et al., 2023). The method maintains a posterior over candidate graphs and averages interventional predictions across it (Equation 33c). The next intervention is then chosen to maximize the expected information gain about the joint of the optimal outcome value and the graph (Equation 33d), rather than to maximize the predicted reward alone. This focus has an important practical consequence. CEO is not optimized to

recover a globally accurate graph. It spends budget on experiments that resolve the structural ambiguities that matter to decide how to act so that it can focus on a decision-relevant subgraph for the target Y .

Formal summary (CEO).

$$\begin{aligned} (\text{assumptions}) \quad & \mathcal{G} \text{ unknown; } G \in \{\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(R)}\}, \text{ prior } P(G); \\ & V_j = f_j(\text{Pa}_j^G) + \epsilon_j, \quad f_j \sim \mathcal{GP}(0, k_j), \quad \epsilon_j \sim \mathcal{N}(0, \sigma_j^2) \end{aligned} \quad (33a)$$

$$(\text{objective}) \quad \text{Equation 10 in the (unknown) true graph; } P(G | \mathcal{D}) \propto p(\mathcal{D} | G) P(G) \quad (33b)$$

$$\begin{aligned} (\text{surrogate}) \quad & m_s(\mathbf{x}) = \mathbb{E}_{P(G|\mathcal{D})}[\mathbb{E}[Y | \text{do}(X_s=\mathbf{x}), G]], \\ & \hat{V}(\mathbf{x}) = \mathbb{E}_{P(G|\mathcal{D})}[\hat{V}[Y | \text{do}(X_s=\mathbf{x}), G]] + \mathbb{V}_{P(G|\mathcal{D})}[\mathbb{E}[Y | \text{do}(X_s=\mathbf{x}), G]] \end{aligned} \quad (33c)$$

$$\begin{aligned} (\text{acquisition}) \quad & \alpha_{\text{CES}}(X_s, \mathbf{x}) = \frac{1}{\text{Co}(X_s, \mathbf{x})} \mathbb{E}_{p(\mathbf{v}_Y, y | \text{do}(X_s=\mathbf{x}), \mathcal{D})} \left[\mathbb{H}[p(y^*, G | \mathcal{D})] \right. \\ & \left. - \mathbb{H}[p(y^*, G | \mathcal{D} \cup \{(\mathbf{x}, \mathbf{v}_Y, y)\})] \right] \end{aligned} \quad (33d)$$

The graph posterior in Equation 33b is available in closed form because the GP marginal likelihoods are tractable. The entropy $\mathbb{H}$ in Equation 33d is over the *joint* of the global optimal value y^* and the graph G , so a single experiment is credited both for shrinking structure uncertainty and for localizing the optimum; $\mathbf{v}_Y$ denotes the observed non-target variables of the interventional sample. Because $p(y^* | \mathcal{D})$ has no closed form, CEO approximates it by a mixture over which scope hosts the optimum,

$$p(y^* | \mathcal{D}) \approx \sum_{X_s \in \text{ES}} P(X_s = X^*) p(y_s^* | \mathcal{D}), \quad (34)$$

with set-level weights maintained by a UCB-style bandit rule over the exploration set ES of candidate scopes.

The main limitation is computational overhead. Approximating a posterior over graphs typically requires sampling in a large discrete space, and estimating information gain must be repeated across many candidate interventions (the outer expectation in Equation 33d). As a result, CEO is most practical when the graph is small to medium or when strong structural priors substantially narrow the set of plausible graphs.

Optimism over plausible models (GACBO). Graph-Agnostic Causal Bayesian Optimization, GACBO, follows a more algorithmic route that emphasizes scalability (Mukherjee et al., 2024). Rather than choosing interventions by entropy reduction, it maintains a confidence set of causal models that remain statistically plausible given the data and selects interventions that could be optimal under at least one model in this set. Formally, let $\mathcal{C}_t$ denote the confidence set of plausible causal models (graph-mechanism pairs) at round t . GACBO selects the next intervention by solving

$$(X_s, \mathbf{x}_s) \in \arg \max_{X_s, \mathbf{x}_s} \max_{M \in \mathcal{C}_t} \hat{f}_M(X_s, \mathbf{x}_s), \quad (35)$$

where $\hat{f}_M$ is the predicted interventional objective under model M , i.e., $\hat{f}_M(X_s, \mathbf{x}_s) = \mathbb{E}[Y | \tilde{F}^{\mathcal{G}'}, \text{do}(X_s=\mathbf{x}_s)]$ for a graph-mechanism pair $M = (\mathcal{G}', \tilde{F}^{\mathcal{G}'})$. This “optimism in the face of uncertainty” principle selects interventions that look promising under *some* plausible causal explanation, which drives exploration toward experiments that can eliminate incorrect models. The confidence set itself is constructed from per-node GP posteriors under each candidate graph:

$$\begin{aligned} \mathcal{C}_t &= \left\{ (\mathcal{G}', \tilde{F}) : \mathcal{G}' \in G_t, \quad \tilde{f}_j \in \mathcal{H}_{k_j}, \quad |\tilde{f}_j(\cdot) - \mu_{j,t}^{\mathcal{G}'}(\cdot)| \leq \beta_t \sigma_{j,t}^{\mathcal{G}'}(\cdot) \quad \forall j \right\}, \\ G_t &= \{ \mathcal{G}' \in G_{t-1} : \exists \tilde{F} \text{ satisfying the above bounds under } \mathcal{G}' \}, \end{aligned} \quad (36)$$

where $(\mu_{j,t}^{\mathcal{G}'}, \sigma_{j,t}^{\mathcal{G}'})$ are per-node GP posteriors fitted under the parent sets of $\mathcal{G}'$ and β_t is chosen so that the true graph-mechanism pair remains in $\mathcal{C}_t$ with probability at least $1 - \delta$; candidate graphs are eliminated once no mechanism vector remains plausible under them (Mukherjee et al., 2024). Compared with CEO, GACBO replaces posterior averaging with best-case evaluation over models consistent with the observations,

avoiding explicit entropy computations; this aligns closely with optimistic bandit and Bayesian optimization principles and can make it easier to implement and scale.

The main risk is over-optimism early in the learning process. When the confidence set is wide, the method may conduct trials on interventions that look good only under structural hypotheses that have not yet been ruled out. Performance, therefore, depends on the calibration of the confidence set and on how quickly interventional evidence eliminates incorrect graphs and mechanisms. Recent work extends this direction by exploring alternative uncertainty decompositions, including approaches that emphasize learning exogenous distributions and broader formulations of unknown-graph problems (Ren et al., 2024; Durand et al., 2025).

Formal summary (GACBO).

(assumptions) $\mathcal{G}$ unknown; causal sufficiency and acyclicity;
mechanisms in RKHS balls $\mathcal{H}_{k_j}$; all nodes observed after each intervention (37a)

(objective) Equation 10 with the graph unknown (37b)

(surrogate) graph-mechanism confidence set $\mathcal{C}_t$ and surviving graphs G_t , Equation 36 (37c)

(acquisition) optimism over plausible models, Equation 35 (37d)

Summary. Unknown-graph CBO turns intervention design into an adaptive experimental design problem in which structure learning is not a preprocessing step but part of the optimization loop. The CEO and GACBO represent two complementary ways to couple these goals. The CEO prioritizes information gain about the best intervention under a Bayesian posterior, while GACBO prioritizes optimistic improvement under a frequentist-style set of plausible models. The choice between them largely reflects an information-versus-optimism trade-off and the computational budget available for representing and updating uncertainty over graphs.

3.5 Robustness to noise and priors

Even with a known graph, practical CBO pipelines face two recurring robustness issues. First, interventional outcomes can be noisy, heteroscedastic, or affected by unmodeled disturbances, which can destabilize improvement-based acquisitions. Second, observationally derived priors may be misspecified, either because identifiability assumptions fail, or because the estimand is hard to estimate accurately from finite data, or because observational correlations leak confounding into a prior mean.

Counter-noise Expected Improvement (CNEI) (Li et al., 2023) targets both issues through a noise-aware acquisition and a learned prior construction. CNEI modifies EI-style selection to account for the noise level (Equation 38d) so that the acquisition does not reward an apparent improvement dominated by measurement variance. In parallel, it estimates a surrogate prior mean by fitting predictive models to observational data and using the resulting predictor as a warm-start mean for the GP (Equation 38c). This can improve the early-stage search by giving the surrogate a reasonable global shape before sufficient interventional data accumulate.

Formal summary (CNEI).

(assumptions) $\mathcal{G}$ known DAG; $y_s = g_s(\mathbf{x}_s) + \epsilon$, $\epsilon \sim \mathcal{N}(0, \delta^2)$; $X_s \in \text{ES}$ (MIS or POMIS of $\mathcal{G}$) (38a)

(objective) Equation 10 (baseline CBO problem, unchanged) (38b)

(surrogate) $g_s \sim \mathcal{GP}(m_s, k_s)$, $m_s(\mathbf{x}_s) = \sum_{\mathbf{z}} \hat{p}(\mathbf{z}) \hat{r}(\mathbf{x}_s, \mathbf{z})$, $\hat{r}(\mathbf{x}_s, \mathbf{z}) \approx \mathbb{E}[Y \mid X_s = \mathbf{x}_s, Z = \mathbf{z}]$ fit on $\mathcal{D}_{\text{obs}}$ (38c)

(acquisition) $\alpha_{\text{CNEI}}^{(s)}(\mathbf{x}_s) = y_s^+ - \mathbb{E}[(g_s(\mathbf{x}_s^+) - g_s(\mathbf{x}_s))_+ \mid \mathcal{D}]$, $(X_s, \mathbf{x}_s)_{t+1} \in \arg \min_{X_s \in \text{ES}, \mathbf{x}_s} \alpha_{\text{CNEI}}^{(s)}(\mathbf{x}_s)$ (38d)

Here $\hat{r}$ is a supervised regressor (support vector regression, ridge regression, or a random forest) plugged into the adjustment formula of Equation 38c, $\mathbf{x}_s^+$ is the within-scope incumbent, and y_s^+ its observed value, which places per-scope improvements on a common scale for cross-scope comparison. The “counter-noise” element

Table 2: Classification of major CBO methods by key assumptions and modeling choices. Each row lists the graph/system-knowledge assumption, surrogate modeling target, decision strategy, and intervention or environmental feature introduced by the method. Superscripts mark causal-sufficiency assumptions: ‡ indicates that the method’s formulation can accommodate hidden confounders in the graph (e.g., via identification-based adjustment or averaging over structures), whereas † indicates that the method models individual mechanisms or rewards directly and therefore assumes causal sufficiency (fully observed parent–child relations). Unmarked methods follow the common known-graph setting in which causal sufficiency is typically assumed but is not central to the method’s design.

Method	Graph/system assumptions	Surrogate target	Acquisition / strategy	Environment / intervention feature
CBO (2020)	Known DAG ‡	Effect-level GP for Y with causal prior	Causal EI; cost-aware; observe vs. intervene	Scope + value, hard interventions
DCBO (2021)	Known DAG ‡	Dynamic effect-level GP	Time-aware causal EI	Intervention sequences over time
MCBO (2023)	Known DAG †	Mechanism-level GPs, full SCM	UCB / optimism, regret bounds	Hard/soft interventions; mechanism modeling
cCBO (2023)	Known DAG ‡	Multi-output GP, objective and constraints	Constrained EI, feasibility-weighted	Safety and feasibility constraints
fCBO (2023)	Known DAG	GP over policies, RKHS function space	Functional EI	Functional and policy interventions
CoCaBO (2026)	Known DAG	BO within scope plus bandit over scopes	Two-layer, MAB-UCB then BO within scope	Contextual mixed interventions
ACBO (2024)	Known DAG †	Mechanism-level modeling, full SCM	MW-style online learning plus causal UCB oracle	Adversarial and non-stationary environments
MO-CBO (2025)	Known DAG	GP per objective, local MOBO subproblems	Relative hypervolume improvement	Multi-objective Pareto optimization
HCBO (2024)	Known DAG	Multiple effect-level GPs, selected scopes	Normalized UCB-style scope scoring	High-dimensional scope selection
CEO (2023)	Unknown DAG ‡	Bayesian model average over graphs, mixture	Entropy and information gain	Learns graph structure during optimization
GACBO (2024)	Unknown DAG †	Confidence set over graphs and mechanisms	Optimism / UCB over plausible models	Structure discovery via optimistic exploration
CNEI (2023)	Known DAG	GP for Y with learned supervised prior	Counter-noise EI	Noise-robust acquisition; learned priors

is that the improvement in Equation 38d is defined on the *latent* objective g_s under the joint posterior of $(g_s(\mathbf{x}_s), g_s(\mathbf{x}_s^+))$ rather than on the noisy observation $y_s^+ = g_s(\mathbf{x}_s^+) + \epsilon$; the variance $\mathbb{V}[g_s(\mathbf{x}_s) - g_s(\mathbf{x}_s^+)] = k_s(\mathbf{x}_s, \mathbf{x}_s) + k_s(\mathbf{x}_s^+, \mathbf{x}_s^+) - 2k_s(\mathbf{x}_s, \mathbf{x}_s^+)$ retains the covariance with the incumbent, so apparent improvements driven by measurement variance are discounted.¹

¹We state the noise-robust acquisition in the survey’s minimization convention as the expected best value after the candidate intervention; the source alternates sign conventions, and we normalize the improvement variance to the standard variance-of-difference form. The generalization from the source’s single adjustment variable to an adjustment set Z in Equation 38c is a survey-level paraphrase.

Table 3: Datasets used in empirical evaluations across CBO variants (as reported in the corresponding papers). We add a category column to make the dataset’s family structure explicit and use light grid lines to improve readability. The dataset-usage table is restricted to single-objective CBO evaluations.

Category	Dataset	CBO	cCBO	fCBO	MCBO	HCBO	CEO	GACBO	CNEI	ACBO	CoCaBO
Synthetic (hard)	ToyGraph	✓	✓	✓	✓		✓	✓			
	Synthetic	✓							✓		
	Synthetic-2		✓								
	Chain-hard			✓							
Synthetic (soft)	Ackley				✓			✓		✓	
	Rosenbrock				✓			✓		✓	
	Dropwave				✓			✓		✓	
	Alpine2				✓			✓		✓	
	Chain-soft			✓							
Real / fitted SCMs (hard)	Ecology	✓				✓					
	Healthcare	✓	✓	✓	✓	✓	✓				✓
	Protein-reconstructed		✓								
	Epidemiology						✓	✓			

A critical distinction concerns the causal status of this prior. CNEI’s learned prior is an adjustment-type plug-in (Equation 38c) whose inner regression $\hat{r}$ is a purely supervised, correlational fit; its causal validity therefore rests entirely on the assumed adjustment set. When that set is invalid or hidden confounding is present, the prior degrades to a correlational predictor $\mathbb{E}[Y | X]$ and can encode spurious associations that a valid causal adjustment would have removed. CNEI therefore trades causal soundness for practical robustness, relying on the noise-aware acquisition to prevent the potentially biased prior from dominating the search as interventional evidence accumulates.

More broadly, observationally learned priors should serve as weak guidance with appropriately inflated uncertainty rather than as a substitute for interventional evidence, motivating robust CBO procedures that explicitly control how strongly observational information can influence acquisition decisions under noise, misspecification, and possible hidden confounding.

3.6 Comparison tables and unifying observations

This subsection consolidates the cross-method synthesis in one place: rather than closing each preceding subsection with its own summary, we compare the method families once, here, along the design axes introduced at the start of Section 3. Tables 2 and 3 summarize the field from two complementary angles. Table 2 organizes methods by key design dimensions: assumptions about the causal graph, the surrogate modeling target (effect-level, mechanism-level, or graph-uncertainty) and the acquisition or decision strategy. Table 3 maps empirical coverage and reveals substantial fragmentation: beyond a small core of shared benchmarks (ToyGraph, Healthcare), many datasets appear in only one or two papers, which complicates cross-paper performance comparison.

We additionally include an identification-strategy audit table. In particular, none of the benchmarked released pipelines exposes a general front-door or full-ID implementation in the benchmark path used here; where observational adjustment is used, it is method-specific and should not be interpreted as a complete ID-algorithm implementation.

Together, these tables highlight two orthogonal axes of variation. The *modeling axis* ranges from effect-level surrogates (CBO, DCBO, cCBO, HCBO, CNEI) through mechanism-level models with uncertainty propagation (MCBO, ACBO) to explicit graph-uncertainty methods (CEO, GACBO). The *action axis* ranges from scope-and-value selection (CBO, DCBO, HCBO) through policy- or context-dependent decisions (fCBO, CoCaBO) to adversarial interaction (ACBO). The identification-strategy audit (Table 4) additionally clarifies that the benchmark compares released pipelines as complete systems, not isolated acquisition functions. Recognizing these interactions is important for practitioners choosing a method and for researchers identifying gaps in the design space.

Table 4: Identification-strategy audit for benchmarked pipelines. “Method-level identification” describes the causal-effect identification strategy assumed or implemented by the released method. “Benchmark effect source” describes how the benchmark wrapper supplies or approximates effects when re-running the released pipeline.

Method	Surrogate level	Method-level identification	Benchmark effect source
BO	Target-level GP	None; non-causal baseline	Direct black-box evaluations
CoCaBO	Mixed-input BO	No explicit ID estimand in our wrapper; graph/scope-based policy selection	Benchmark SCM / task oracle evaluations
CBO	Effect-level GP	Observational adjustment when identifiable; not full ID in released code	Benchmark wrapper uses SCM/oracle effect estimates where available
cCBO	Effect-level / constrained GP	Observational adjustment with constraints; not full ID in released code	Benchmark wrapper uses SCM/oracle effect estimates where available
DCBO	Dynamic effect-level / temporal GP	Temporal adjustment under known dynamic-graph assumptions; not full ID	Benchmark dynamic SCM / oracle evaluations
HCBO	Scope-level / high-dimensional CBO	Graph-based scope selection; not full ID	Benchmark SCM / SEM evaluations
CEO	Graph-uncertain / oracle SEM search	Structure-learning over candidate graphs; not full ID over arbitrary latent structures	Candidate graph / SEM evaluations
MCBO	Mechanism-level GP	None required for effect identification; mechanisms are modeled directly	Mechanism-level simulation / forward propagation

3.7 Connections to Adjacent Fields

CBO does not exist in isolation. Several established research areas share overlapping goals, tools, or assumptions. Understanding these connections clarifies what CBO inherits, where it innovates, and what cross-pollination opportunities remain.

Causal abstractions and multi-scale decision making. Causal abstraction studies how high-level causal models relate to lower-level mechanistic models through abstraction maps. This perspective is directly relevant to CBO whenever interventions can be chosen at multiple resolutions. Causally Abstracted Multi-Armed Bandits (CAMAB) (Zennaro et al., 2024) formalizes decision making across causally linked abstraction levels, while AT-UCB (Dyer et al., 2025) shows that exploiting abstraction hierarchies can reduce cumulative regret in complex simulators, including epidemiological models—precisely the kind of complex simulation setting where CBO’s aspiration for real-world deployment meets computational tractability limits. For CBO, this suggests a multi-scale analogue of scope reduction: rather than only pruning intervention scopes within one graph, the optimizer may transfer or prune decisions across graphs at different granularities. A first proposal in this direction for CBO itself is MFACBO (Zeitler, 2025), which connects causal abstraction levels to multi-fidelity Bayesian optimization; we note that this proposal is a workshop contribution and not yet peer-reviewed. We therefore treat causal abstraction learning as distinct from causal representation learning from raw observations, and as a promising route for CBO over latent, coarse-grained, or hierarchical causal variables.

Causal bandits. The causal bandit literature (Lattimore et al., 2016; Lu et al., 2021; Nair et al., 2021) also selects interventions using causal structure, but typically assumes discrete action sets, parametric reward models, and focuses on cumulative regret in stationary environments. CBO extends this setting in three directions: (i) continuous intervention domains requiring nonparametric surrogates, (ii) mixed discrete–continuous action spaces (scope value + and (iii) more robust uncertainty modeling through GP. Conversely, causal bandits offer tighter finite-sample regret bounds and principled treatments of partial graph knowledge that CBO methods have only begun to adopt (e.g., MCBO’s regret analysis draws on GP-bandit theory). A promising direction is to unify the two frameworks by developing CBO methods with bandit-style theoretical guarantees for continuous domains.

Bayesian experimental design. The Bayesian optimal experimental design (BOED) (Chaloner & Verdinelli, 1995; Foster et al., 2021) selects experiments to maximize information about the model pa-

rameters. CEO (Branchini et al., 2023) can be viewed as a BOED method applied to intervention selection: choose experiments to maximize the information gained about the identity of the best intervention. The key distinction is that standard BOED targets parameter estimation without an optimization objective, whereas CEO jointly optimizes information and reward. This connection suggests that advances in amortized BOED (Foster et al., 2021) could accelerate information-seeking CBO methods.

Active causal discovery. Active structure learning (Spirtes et al., 2000; Chickering, 2002; Eberhardt & Scheines, 2007) selects interventions to learn the causal graph itself, while CBO selects interventions to optimize an outcome. Unknown-graph CBO methods (CEO, GACBO) bridge these goals by learning graph structure as a means to better optimization. The distinction is that active discovery aims for global graph accuracy, whereas CBO needs only *decision-relevant* structure, the parts of the graph that change which intervention is optimal. This focus on decision-relevant learning is a truly novel aspect of unknown-graph CBO.

Safe and constrained optimization. Safe Bayesian optimization (Sui et al., 2015; Berkenkamp et al., 2019) maintains feasibility during optimization, which is closely related to cCBO’s (Aglietti et al., 2023) constrained causal optimization. The causal contribution is that safety constraints can be analyzed through the graph: a scope that cannot affect a constraint variable need not model constraint feasibility, reducing the multi-output modeling burden. Future work could combine safe BO’s theoretical safety guaranties with CBO’s causal constraint analysis.

Policy search and reinforcement learning. Functional CBO (fCBO) (Gultchin et al., 2023) and contextual CBO (CoCaBO) (Arsenyan et al., 2026) can be viewed as policy search methods operating within an explicit causal model. Unlike model-free reinforcement learning, these methods exploit the known graph to decompose the policy-to-outcome mapping and to construct informative priors from observational data. The main limitation relative to RL is the assumption of a single-step (or few-step) decision problem; extending CBO to sequential multi-stage intervention policies remains an open challenge.

4 Experiments

Although this paper is primarily a survey, we include a benchmark study to support two of the survey’s main claims. First, existing CBO papers are often evaluated in heterogeneous codebases and on only partially overlapping datasets, making cross-paper comparison difficult. Second, commonly reported efficiency metrics can be sensitive to how optimization trajectories are summarized, especially when improvements occur late in the evaluation budget. Our benchmark is therefore designed to separate algorithm execution from evaluation: methods can run in their native implementations, but their outputs are exported into a common best-so-far trajectory format and re-scored with identical metric code.

4.1 CBO benchmark and evaluation protocol

We release a CBO benchmark at <https://github.com/chenfeng-huang/CBO-Benchmark-TMLR-2026>. The benchmark is a standardized repackaging and extension of inherited CBO/function-network tasks, rather than an entirely new collection of datasets. It includes eight curated hard-intervention scenarios and a separate set of soft-intervention benchmarks, with dataset provenance and code-authoritative domains documented in Appendix A.4. It provides standardized implementations of GAP and PA-GAP. Because existing CBO methods often rely on incompatible software environments, the benchmark does not force all algorithms into a single execution stack. Instead, each method is run in its native or vendored implementation, and its trajectory is exported and re-scored using the same evaluation pipeline. This design enables reproducible comparisons across heterogeneous implementations and makes it easy to add new methods, datasets, or metrics.

Because the implemented methods target different problem formulations, we do not pool them into a single cross-method leaderboard: competitive statements are restricted to methods solving the same formulation, and cross-method aggregation appears only as descriptive rank-reliability tables in Appendix B.3. The

benchmark uses native or vendored implementations with released defaults, common trajectory export, and common scoring; it is therefore a reproducibility harness rather than a fully controlled reimplementations study. Appendix A.2 documents method settings and wrapper choices.

To make the formulation labels of Table 5 usable as fixed reference points rather than user-configured settings, each task in the repository ships with a frozen configuration: the intervention domains, optimization direction, observational sample sizes, constraint functions and thresholds for the constrained tasks, the policy space for Chain-soft, and the reference optimum used for scoring. A new method can therefore be evaluated on a given formulation without any task-design decisions being left to the user.

Execution protocol. For each data set, method, and trial limit, we run 20 random seeds. A seed fixes the observational sample, the initial interventional design, and the stochastic components of the optimizer. All methods are evaluated under the same trial limits, intervention domains, optimization direction, and scoring code. An interventional evaluation counts as one trial. Observational data used to construct priors, learn graphs, or initialize surrogates are fixed before the optimization loop and are not counted as interventional trials unless otherwise stated. Each method is run with its released default hyperparameters unless explicitly stated in Appendix A.2; we do not tune hyperparameters separately for each data set.

Trajectory export and scoring. The output of each method is converted to a common trajectory format with two required fields: trial number and best-so-far objective value in the natural optimization direction of the task. For minimization tasks, this is the minimum value observed up to the trial t ; for maximization tasks, it is the maximum value observed up to the trial t . The trajectory includes an initial value followed by interventional trials. Thus, if the budget is B interventions, the exported trajectory contains $B + 1$ entries: one initial best value and B post-intervention best-so-far values. GAP and PA-GAP are then computed from these exported trajectories using the same evaluation script for all methods.

Reference optimum. Both GAP and PA-GAP require a reference optimum y^* . In our benchmark, y^* is computed offline with oracle access to the benchmark SCM, function generator, or precomputed interventional data. This value is used only for scoring and is never provided to any optimizer. For low-dimensional or discrete intervention domains, we enumerate or densely evaluate admissible intervention scopes and values. For continuous domains, we use high-budget offline search with multi-start optimization or stored oracle intervention sets. The exact offline budgets, grids, restarts, and dataset-specific procedures are reported in the Appendix A.3.

Failure handling. When a method proposes an infeasible or out-of-domain intervention, we follow the behavior of the released implementation. If the implementation projects the proposals back to the admissible domain, the projected intervention is evaluated. Otherwise, the trial is marked as a failed proposal, and the trajectory records the current best-so-far value. Runs with numerical errors, missing outputs, or non-finite objective values are logged and are not silently removed. Runtime is not used as a primary metric because implementations differ substantially in language, hardware support, and dependency stacks.

Cost model. Unless otherwise stated, each intervention has a unit cost. This choice isolates the efficiency of the sample but does not evaluate the full cost-aware behavior of the CBO acquisition functions. In real applications, the cost of the intervention can depend on the target variable, the number of variables manipulated, the magnitude of the intervention, safety constraints, or whether the action is observational or experimental. We therefore report the current results as unit-cost benchmarks and treat realistic intervention-cost modeling as a separate evaluation dimension.

4.2 Performance metrics

We evaluated each run using two numerical efficiency metrics, GAP and PA-GAP, and one visualization metric based on the best-so-far objective trajectory. All scalar metrics are computed from the same exported trajectories described above.

We emphasize that GAP and PA-GAP are auxiliary normalized efficiency summaries that complement, rather than replace, standard BO metrics. Best-so-far trajectories are simple-regret-equivalent (see the trajectory-visualization paragraph below), and normalized average-reward trajectories, the analogue of cumulative-regret reporting, are provided in Appendix B.4. The reason for additionally reporting GAP and PA-GAP is aggregation: raw simple regret is not comparable across datasets with different objective scales, initializations, and optimization directions, whereas GAP and PA-GAP normalize improvement relative to the initial value and reference optimum, enabling the cross-dataset rank summaries used in this benchmark. Readers interested only in per-dataset behavior can rely on the regret-equivalent trajectory plots.

GAP. GAP was introduced in DCBO to quantify the efficiency of optimization by combining the final improvement with the speed with which the best value is discovered (Aglietti et al., 2021). Let R_t denote the best improvement ratio so far in the trial t , computed in the natural optimization direction of the task. For maximization tasks,

$$R_t = \min \left\{ \frac{y(\mathbf{x}_t^*) - y(\mathbf{x}_{\text{init}})}{y^* - y(\mathbf{x}_{\text{init}})}, 1 \right\}, \quad (39)$$

and for minimization tasks,

$$R_t = \min \left\{ \frac{y(\mathbf{x}_{\text{init}}) - y(\mathbf{x}_t^*)}{y(\mathbf{x}_{\text{init}}) - y^*}, 1 \right\}. \quad (40)$$

Here, $\mathbf{x}_t^*$ denotes the best point so far in the trial t . The clipping prevents numerical artifacts when an observed value exceeds the offline reference optimum. Let R_T be the final improvement ratio and let t^* be the first trial in which R_T is achieved. If no improvement occurs, we set $t^* = T$. GAP is then

$$\text{GAP} = \left[\underbrace{R_T}_{\text{Final improvement}} + \underbrace{\frac{T - t^*}{T}}_{\text{Discovery efficiency}} \right] / \underbrace{\left(1 + \frac{T - 1}{T} \right)}_{\text{Normalization}}. \quad (41)$$

Under this convention, GAP lies in $[0, 1]$, with larger values indicating larger final improvement and earlier discovery.

Path-Aware GAP (PA-GAP). GAP depends only on the final best value and the first trial at which it is reached. This can under-value runs that improve late or make steady progress over the full budget. To better reflect the full optimization trajectory, we use Path-Aware GAP:

$$\text{PA-GAP} = \frac{1}{T} \sum_{t=1}^T \left(\underbrace{R_t}_{\text{Improvement ratio}} \cdot \underbrace{\frac{T - (t - 1)}{T}}_{\text{Efficiency weight}} \right). \quad (42)$$

PA-GAP multiplies the best-so-far improvement by a time-dependent efficiency weight, so later improvements are down-weighted but not discarded, and averages these weighted improvements over the full trajectory. Since $R_t \in [0, 1]$, raw PA-GAP lies in

$$\left[0, \frac{T + 1}{2T} \right].$$

The upper bound is attained when the reference optimum is found in the first trial and remains the best value so far thereafter. Thus, PA-GAP is a raw trajectory-weighted efficiency score rather than a unit-normalized metric. Larger values indicate faster and stronger trajectory-level improvement, and raw PA-GAP values should be compared directly only under the same trial budget. Appendix A.1.2 gives a counterexample where GAP ranks a worse final trajectory above a better late-improving one, while PA-GAP resolves the ordering.

Trajectory visualization. In addition to scalar metrics, we plot best-so-far objective trajectories over trials. For a fixed dataset and optimization direction, the best-so-far trajectory is equivalent to a simple-regret trajectory up to an additive constant and, for maximization tasks, a sign reversal. We therefore use best-so-far plots as simple-regret-equivalent visualizations in the original objective scale. We also report normalized average-reward trajectories in Appendix B.4 as a diagnostic complement to GAP and PA-GAP.

4.3 Datasets and task conventions

Table 3 summarizes the datasets used in the CBO literature. Most evaluations rely on SCM-generated data, where observational samples are drawn from the unintervened SCM and interventional samples are generated by applying hard or soft interventions to the SCM mechanisms. The optimization direction is data-dependent. For hard-intervention SCM benchmarks, we follow the standard CGO/CBO convention and treat the target as a minimization objective unless explicitly stated otherwise. ToyGraph, Synthetic, Synthetic-2, Chain-hard, Healthcare, Protein-reconstructed, and Epidemiology are minimization tasks, while Ecology is a maximization task. For the soft-intervention benchmarks for the function-network, Ackley, Rosenbrock, Dropwave, and Alpine2 are treated as reward-maximization tasks following the convention used in the BO, MCBO, and ACBO function-network. Chain-soft is a minimization task because the objective is to reduce the target variable Y .

4.3.1 Hard-intervention SCM benchmarks

Hard-intervention benchmarks evaluate optimization when intervened variables are clamped to fixed values. They include small synthetic SCMs, a reconstructed protein-signaling SCM, and real-world-inspired SCMs from ecology, healthcare, and epidemiology. The benchmark also provides optional high-dimensional generators, but the main cross-method comparison focuses on the eight curated scenarios for which multiple implementations can be evaluated under the same protocol.

The eight scenarios are as follows; per-dataset structural equations, causal graphs, intervention domains, and provenance are documented in Appendix A.4 (Tables 11 and 12) and the per-dataset appendix subsections. **ToyGraph** (Aglietti et al., 2020) is a minimal nonlinear three-node chain $X \rightarrow Z \rightarrow Y$ with manipulable X, Z , serving as a sanity check (Appendix A.5.1). **Synthetic** (Aglietti et al., 2020) has seven observed variables plus two latent confounders, with manipulable B, D, E , combining nonlinearity, latent structure, and multiple intervention options (Appendix A.5.2). **Synthetic-2** (Aglietti et al., 2023) is a lightweight nonlinear three-node chain with Gaussian exogenous noise (Appendix A.5.3). **Chain-hard** (Gultchin et al., 2023) has four observed variables with manipulable W, Z clamped in $[-1, 1]$ and a non-manipulable context X whose interaction with Z drives the target (Appendix A.5.4). **Ecology** (Aglietti et al., 2020; Courtney et al., 2017) is a Bermuda-reef ecosystem model in which net ecosystem calcification is maximized under environmental interventions, using the SCM and data released in the public CBO repository (Appendix A.6.1). **Protein-reconstructed** is based on the protein-signaling data of Sachs et al. (2005) and the cCBO graph (Aglietti et al., 2023): Erk is minimized by perturbing Mek, PKC, PKA, and Akt under biological constraints on PKC and PKA; because the fitted SCM used in the original cCBO experiments is unreleased, this task is a transparent reconstruction fitted on 852 observational samples rather than an exact reproduction (Appendix A.6.2). **Healthcare** (Aglietti et al., 2020; Thompson, 2019; Ferro et al., 2015) minimizes PSA through statin and aspirin interventions, with the common continuous relaxation of the binary treatment variables. **Epidemiology** (Branchini et al., 2023; Havercroft & Didelez, 2012) minimizes HIV viral load in a modified epidemiology SCM with code-authoritative intervention domains on L and B (Appendices A.6.3 and A.6.4).

4.3.2 Soft-intervention function-network benchmarks

Soft-intervention benchmarks evaluate settings where interventions modify structural mechanisms rather than simply clamping variables to fixed values. Ackley, Rosenbrock, Dropwave, and Alpine2 are classical global-optimization functions represented as function-network BO benchmarks introduced by Astudillo & Frazier (2021) and adopted in MCBO and ACBO (Sussex et al., 2023; 2024). We preserve the inherited task domains after any internal normalization used by wrappers: Ackley and Rosenbrock use $[-2, 2]$, Dropwave uses $[-5.12, 5.12]$, and Alpine2 uses $[0, 10]$. Thus, the optimizer may operate on normalized coordinates internally, but the benchmark task domains are not changed. In these tasks, the output node is treated as a reward to be maximized, even when this corresponds to a sign-flipped version of a standard minimization benchmark.

We explicitly distinguish these *computational-graph benchmarks* from *SCM benchmarks*. In the function networks, the manipulable input variables are mutually independent, orthogonal search dimensions: there

Table 5: Formulation labels used to contextualize benchmark results. Several labels are not independently rankable because they contain only one method or one dataset; therefore the statistical rank analysis is reported by budget rather than as separate formulation-track leaderboards.

Label	What it identifies	How it is used
Unconstrained hard SCM	Static hard-intervention tasks without explicit feasibility constraints	Descriptive rank summaries by budget (Appendix B.3); comparisons are qualified because some methods are stress-test or graph-uncertain variants
Constrained hard SCM	Tasks where cCBO-style feasibility constraints are part of the intended formulation	Descriptive only: no fully matched constraint-adapted baseline is available across all datasets
Contextual / policy	Settings where scope or intervention choice is policy/context dependent	Descriptive only: CoCaBO is the only currently integrated method in this formulation
Soft computational graph	Ackley, Rosenbrock, Dropwave, and Alpine2 function-network tasks	Descriptive rank summaries by budget (Appendix B.3); explicitly not treated as SCM benchmarks with confounding or cascading manipulable variables
Soft SCM / policy	Chain-soft	Descriptive only: single soft-SCM dataset in the current repository

are no causal relationships among the inputs themselves, so an intervention on one input cannot sever incoming edges of, or cascade into, another manipulable variable. These tasks therefore probe optimization over a known computational composition, but they do not test confounding or cascading effects among manipulable variables, which are core challenges of causal optimization tested by SCM benchmarks such as Healthcare or Ecology. Chain-soft is the only soft-intervention SCM task in the current repository, and results on the two families are labeled and interpreted separately (Table 5).

The five scenarios are as follows; equations, graphs, and visualizations are given in Appendices A.7.1–A.7.5. **Ackley** ($D = 6$) decomposes the highly multimodal Ackley objective into a squared-magnitude component and a cosine component combined at the output node. **Rosenbrock** ($D \in \{3, 5, 7\}$) is a chain-structured accumulation of adjacent-variable-pair terms with a narrow curved valley. **Dropwave** ($D = 2$) computes the radial quantity $\sqrt{x_1^2 + x_2^2}$ at an intermediate node before the oscillatory output transformation. **Alpine2** ($D = 6$) is a multiplicative chain of per-variable Alpine2 terms, separable and multimodal. **Chain-soft** (Gultchin et al., 2023) extends Chain-hard by letting the mechanism of Z be modified through a context-dependent policy $\pi_0(X)$ while W remains under hard interventions in $[-1, 1]$; because the target depends on the interaction between Z and X , the optimal intervention on Z varies with X , making it a benchmark for context-adaptive intervention strategies (Appendix A.7.5).

4.4 Hard-intervention benchmark

To contextualize the benchmark results, Table 5 defines the formulation labels used throughout this section. These labels clarify which methods are solving matched optimization problems and which are included as stress tests or descriptive comparisons.

Compared methods. We compare a non-causal baseline, BO, with seven CBO-family methods: CBO, cCBO, DCBO, MCBO, HCBO, CoCaBO, and CEO (Aglietti et al., 2020; 2023; 2021; Sussex et al., 2023; Wu et al., 2024; Arsenyan et al., 2026; Branchini et al., 2023). We label comparisons by how closely the benchmark matches the method’s intended setting. CBO and CoCaBO are matched known-graph hard-intervention methods for the static SCM tasks. cCBO is included as a controlled constrained-CBO variant; on unconstrained tasks it should be interpreted as a modeling comparison rather than a pure constraint comparison. In particular, cCBO does not reduce to CBO when constraints are inactive: the released cCBO pipeline uses a different surrogate construction and exploration-set handling than the released CBO pipeline, so cCBO-versus-CBO differences on unconstrained tasks reflect these modeling and implementation




Figure 2: Best-so-far trajectories on four synthetic hard-intervention datasets: ToyGraph, Synthetic, Synthetic-2, and Chain-hard. Solid lines show the mean across 20 seeds and shaded regions show the standard deviation. Lower values indicate better objective values for these minimization tasks.

differences rather than an effect of constraint handling. DCBO is designed for time-evolving SCMs, so its use on static hard-intervention tasks is a transfer/stress-test comparison. MCBO is a mechanism-level method and is most directly matched to function-network/mechanism-modeling settings. CEO is an unknown-graph reference method and spends part of its budget resolving structure uncertainty, whereas known-graph methods receive the benchmark graph. BO is a graph-free GP-EI baseline and is used as a robustness reference rather than a causal competitor. We exclude fCBO, CNEI, and GACBO because we did not obtain compatible public implementations of the benchmark protocol. We exclude ACBO and MO-CBO from this hard-intervention benchmark because their released implementations or task settings are not aligned with the single-objective hard-intervention protocol considered here.

Selected datasets. We evaluated eight hard-intervention datasets: ToyGraph, Synthetic, Synthetic-2, Chain-hard, Ecology, Protein-reconstructed, Healthcare, and Epidemiology. These datasets cover small synthetic SCMs, real or fitted-SCM benchmarks, and settings that are unevenly shared across prior CBO papers. Protein-reconstructed is included as a transparent reconstruction of the cCBO Protein benchmark, rather than as an exact reproduction of the unreleased fitted SCM used in the original cCBO experiments.

Results summary. The hard-intervention benchmark shows that no method dominates uniformly across datasets, budgets, and metrics within the benchmark-specified regime. In the main tables, the environment SCM, intervention domains, and reference optima are fixed by the benchmark, and graph-using methods are evaluated with the benchmark graph unless they are explicitly unknown-graph methods. Under this protocol, the descriptive rankings vary across datasets, budgets, and metrics, and the rank-reliability analysis in Appendix B.3 identifies no statistically separable uniformly best method: CoCaBO attains the lowest




Figure 3: Best-so-far trajectories on four real or fitted-SCM hard-intervention datasets: Ecology, Protein-reconstructed, Healthcare, and Epidemiology. Solid lines show the mean across 20 seeds and shaded regions show the standard deviation. Ecology is a maximization task, while Protein-reconstructed, Healthcare, and Epidemiology are minimization tasks.

descriptive average rank at $T = 100$ and $T = 50$ and BO at $T = 20$, but the Nemenyi critical difference exceeds the rank spread at every budget. We interpret CoCaBO’s strong cells (Synthetic variants, Healthcare, Epidemiology) as conditional patterns on datasets where mixed scope selection and contextualized search align with the benchmark structure, rather than as evidence that CoCaBO is intrinsically strongest across CBO settings; DCBO’s strength on ToyGraph and MCBO’s on Protein-reconstructed are likewise dataset-conditional. BO remains competitive on Chain-hard, Ecology, and Healthcare, where early improvements earn favorable trajectory-level (PA-GAP) scores even when BO is not the best final (GAP) performer. These results show that causal structure can improve optimization, but its benefit depends on how the method’s assumptions match the dataset and on whether performance is measured by final discovery or by the full optimization path.

Trajectory-level behavior. Figures 2 and 3 complement the scalar table by showing how the methods reach their final scores. In synthetic datasets, many methods improve through a few large drops followed by long plateaus, suggesting that much of the gain comes from identifying a useful intervention scope rather than from gradual local refinement; this plateau structure is one source of GAP/PA-GAP disagreement (Section 4.2). In Synthetic, CoCaBO shows strong path-level progress, consistent with its high scores under both metrics. On ToyGraph, DCBO is strong under GAP, while PA-GAP shows that early progress with CBO and HCBO-style can also matter. In Chain-hard, BO is highly competitive, indicating that a flexible black-box optimizer can perform well when the effective intervention landscape is simple or when the causal prior does not provide a decisive advantage.

Table 6: Hard-intervention performance across different trial limits and datasets. Each value is the average $\pm$ standard error across 20 random seeds and different initializations of the observational and interventional data. Higher is better for both GAP and PA-GAP. Best mean values in each row are in **bold**.

Trial limit	Metric	Dataset	BO	CBO	cCBO	DCBO	MCBO	HCBO	CoCaBO	CEO
100	GAP	ToyGraph	0.346 $\pm$.046	0.810 $\pm$.044	0.730 $\pm$.023	0.949$\pm$.048	0.547 $\pm$.032	0.727 $\pm$.030	0.625 $\pm$.054	0.654 $\pm$.040
		Synthetic	0.401 $\pm$.038	0.658 $\pm$.023	0.691 $\pm$.028	0.482 $\pm$.007	0.652 $\pm$.030	0.559 $\pm$.024	0.757$\pm$.049	0.565 $\pm$.038
		Synthetic-2	0.000 $\pm$.000	0.544 $\pm$.021	0.549 $\pm$.017	0.485 $\pm$.003	0.000 $\pm$.000	0.174 $\pm$.016	0.562$\pm$.045	0.378 $\pm$.038
		Chain-hard	0.721 $\pm$.032	0.688 $\pm$.014	0.734$\pm$.025	0.542 $\pm$.027	0.502 $\pm$.046	0.135 $\pm$.024	0.205 $\pm$.018	0.502 $\pm$.034
		Ecology	0.799$\pm$.041	0.511 $\pm$.017	0.697 $\pm$.027	0.637 $\pm$.026	0.000 $\pm$.000	0.334 $\pm$.030	0.000 $\pm$.000	0.336 $\pm$.032
		Protein-reconstructed	0.751 $\pm$.030	0.449 $\pm$.043	0.000 $\pm$.000	0.000 $\pm$.000	0.878$\pm$.064	0.412 $\pm$.032	0.664 $\pm$.052	0.275 $\pm$.021
		Healthcare	0.310 $\pm$.020	0.577 $\pm$.024	0.842 $\pm$.030	0.487 $\pm$.001	0.510 $\pm$.023	0.335 $\pm$.007	0.964$\pm$.044	0.696 $\pm$.066
		Epidemiology	0.634 $\pm$.049	0.844 $\pm$.029	0.000 $\pm$.000	0.375 $\pm$.028	0.000 $\pm$.000	0.557 $\pm$.024	0.898$\pm$.059	0.308 $\pm$.021
	PA-GAP	ToyGraph	0.224 $\pm$.031	0.069 $\pm$.009	0.235 $\pm$.030	0.337$\pm$.043	0.071 $\pm$.010	0.285 $\pm$.036	0.087 $\pm$.012	0.295 $\pm$.037
		Synthetic	0.198 $\pm$.035	0.172 $\pm$.028	0.184 $\pm$.027	0.036 $\pm$.004	0.162 $\pm$.023	0.107 $\pm$.020	0.411$\pm$.003	0.358 $\pm$.045
		Synthetic-2	0.000 $\pm$.000	0.055 $\pm$.008	0.065 $\pm$.008	0.002 $\pm$.000	0.000 $\pm$.000	0.099 $\pm$.012	0.129$\pm$.021	0.035 $\pm$.005
		Chain-hard	0.495$\pm$.002	0.280 $\pm$.042	0.244 $\pm$.032	0.049 $\pm$.006	0.111 $\pm$.015	0.026 $\pm$.005	0.038 $\pm$.004	0.002 $\pm$.000
		Ecology	0.303$\pm$.039	0.117 $\pm$.006	0.204 $\pm$.027	0.144 $\pm$.019	0.000 $\pm$.000	0.210 $\pm$.026	0.000 $\pm$.000	0.016 $\pm$.002
		Protein-reconstructed	0.407$\pm$.044	0.254 $\pm$.036	0.000 $\pm$.000	0.000 $\pm$.000	0.045 $\pm$.006	0.000 $\pm$.000	0.240 $\pm$.030	0.014 $\pm$.004
		Healthcare	0.125 $\pm$.014	0.175 $\pm$.020	0.170 $\pm$.022	0.000 $\pm$.000	0.100 $\pm$.021	0.010 $\pm$.001	0.329$\pm$.045	0.200 $\pm$.025
		Epidemiology	0.299 $\pm$.038	0.346 $\pm$.021	0.000 $\pm$.000	0.175 $\pm$.022	0.000 $\pm$.000	0.138 $\pm$.017	0.392$\pm$.047	0.005 $\pm$.001
50	GAP	ToyGraph	0.296 $\pm$.043	0.676 $\pm$.045	0.723 $\pm$.022	0.896$\pm$.041	0.513 $\pm$.024	0.643 $\pm$.030	0.629 $\pm$.048	0.690 $\pm$.036
		Synthetic	0.531 $\pm$.018	0.597 $\pm$.024	0.565 $\pm$.023	0.422 $\pm$.003	0.630 $\pm$.022	0.432 $\pm$.019	0.793$\pm$.056	0.372 $\pm$.034
		Synthetic-2	0.000 $\pm$.000	0.459 $\pm$.021	0.531$\pm$.014	0.467 $\pm$.007	0.000 $\pm$.000	0.160 $\pm$.020	0.438 $\pm$.023	0.213 $\pm$.043
		Chain-hard	0.919$\pm$.018	0.462 $\pm$.009	0.712 $\pm$.023	0.535 $\pm$.024	0.215 $\pm$.027	0.302 $\pm$.016	0.106 $\pm$.016	0.505 $\pm$.032
		Ecology	0.798$\pm$.037	0.363 $\pm$.015	0.680 $\pm$.023	0.624 $\pm$.021	0.000 $\pm$.000	0.246 $\pm$.016	0.000 $\pm$.000	0.151 $\pm$.018
		Protein-reconstructed	0.656 $\pm$.027	0.553 $\pm$.038	0.000 $\pm$.000	0.000 $\pm$.000	0.752$\pm$.041	0.320 $\pm$.017	0.731 $\pm$.041	0.377 $\pm$.025
		Healthcare	0.366 $\pm$.019	0.598 $\pm$.019	0.758 $\pm$.027	0.474 $\pm$.000	0.315 $\pm$.019	0.155 $\pm$.002	0.927$\pm$.036	0.694 $\pm$.046
		Epidemiology	0.750 $\pm$.043	0.644 $\pm$.024	0.000 $\pm$.000	0.596 $\pm$.021	0.000 $\pm$.000	0.461 $\pm$.018	0.793$\pm$.041	0.096 $\pm$.001
	PA-GAP	ToyGraph	0.177 $\pm$.029	0.068 $\pm$.008	0.233 $\pm$.030	0.316 $\pm$.040	0.063 $\pm$.010	0.262 $\pm$.033	0.357$\pm$.050	0.296 $\pm$.037
		Synthetic	0.130 $\pm$.021	0.131 $\pm$.025	0.191 $\pm$.017	0.033 $\pm$.004	0.150 $\pm$.022	0.050 $\pm$.013	0.349 $\pm$.040	0.357$\pm$.032
		Synthetic-2	0.000 $\pm$.000	0.043 $\pm$.008	0.063 $\pm$.008	0.002 $\pm$.000	0.000 $\pm$.000	0.091$\pm$.011	0.051 $\pm$.014	0.033 $\pm$.004
		Chain-hard	0.500$\pm$.002	0.178 $\pm$.030	0.233 $\pm$.031	0.049 $\pm$.006	0.078 $\pm$.013	0.012 $\pm$.003	0.034 $\pm$.004	0.002 $\pm$.000
		Ecology	0.306$\pm$.039	0.085 $\pm$.013	0.194 $\pm$.027	0.140 $\pm$.019	0.000 $\pm$.000	0.208 $\pm$.025	0.000 $\pm$.000	0.015 $\pm$.002
		Protein-reconstructed	0.339$\pm$.031	0.192 $\pm$.034	0.000 $\pm$.000	0.000 $\pm$.000	0.045 $\pm$.006	0.000 $\pm$.000	0.157 $\pm$.022	0.003 $\pm$.001
		Healthcare	0.120 $\pm$.014	0.131 $\pm$.018	0.263 $\pm$.038	0.000 $\pm$.000	0.034 $\pm$.011	0.009 $\pm$.001	0.295$\pm$.042	0.202 $\pm$.025
		Epidemiology	0.247 $\pm$.039	0.222 $\pm$.020	0.000 $\pm$.000	0.169 $\pm$.021	0.000 $\pm$.000	0.127 $\pm$.015	0.354$\pm$.041	0.000 $\pm$.000
20	GAP	ToyGraph	0.383 $\pm$.022	0.510 $\pm$.044	0.702 $\pm$.023	0.729$\pm$.038	0.406 $\pm$.020	0.372 $\pm$.022	0.539 $\pm$.043	0.508 $\pm$.028
		Synthetic	0.239 $\pm$.014	0.392 $\pm$.020	0.367 $\pm$.022	0.231 $\pm$.002	0.558 $\pm$.016	0.166 $\pm$.002	0.756$\pm$.052	0.390 $\pm$.028
		Synthetic-2	0.000 $\pm$.000	0.285 $\pm$.017	0.474$\pm$.018	0.408 $\pm$.002	0.000 $\pm$.000	0.186 $\pm$.017	0.000 $\pm$.000	0.282 $\pm$.023
		Chain-hard	0.819$\pm$.014	0.136 $\pm$.003	0.641 $\pm$.022	0.510 $\pm$.018	0.360 $\pm$.023	0.000 $\pm$.000	0.065 $\pm$.011	0.513 $\pm$.028
		Ecology	0.795$\pm$.039	0.208 $\pm$.003	0.625 $\pm$.020	0.585 $\pm$.016	0.000 $\pm$.000	0.271 $\pm$.012	0.000 $\pm$.000	0.340 $\pm$.043
		Protein-reconstructed	0.503 $\pm$.022	0.508 $\pm$.034	0.000 $\pm$.000	0.000 $\pm$.000	0.540 $\pm$.036	0.028 $\pm$.016	0.648$\pm$.041	0.168 $\pm$.016
		Healthcare	0.604 $\pm$.015	0.392 $\pm$.021	0.490 $\pm$.019	0.432 $\pm$.000	0.000 $\pm$.000	0.172 $\pm$.001	0.810$\pm$.017	0.689 $\pm$.036
		Epidemiology	0.646 $\pm$.030	0.452 $\pm$.022	0.000 $\pm$.000	0.457 $\pm$.017	0.000 $\pm$.000	0.154 $\pm$.015	0.791$\pm$.028	0.000 $\pm$.000
	PA-GAP	ToyGraph	0.088 $\pm$.021	0.392$\pm$.051	0.227 $\pm$.029	0.260 $\pm$.034	0.040 $\pm$.008	0.200 $\pm$.040	0.190 $\pm$.029	0.300 $\pm$.036
		Synthetic	0.026 $\pm$.008	0.034 $\pm$.012	0.147 $\pm$.017	0.024 $\pm$.003	0.116 $\pm$.021	0.000 $\pm$.000	0.334 $\pm$.035	0.361$\pm$.042
		Synthetic-2	0.000 $\pm$.000	0.016 $\pm$.005	0.057 $\pm$.007	0.002 $\pm$.000	0.000 $\pm$.000	0.074$\pm$.008	0.000 $\pm$.000	0.027 $\pm$.004
		Chain-hard	0.357$\pm$.035	0.032 $\pm$.007	0.200 $\pm$.029	0.047 $\pm$.006	0.034 $\pm$.009	0.000 $\pm$.000	0.031 $\pm$.004	0.002 $\pm$.000
		Ecology	0.316$\pm$.039	0.030 $\pm$.008	0.166 $\pm$.025	0.126 $\pm$.019	0.000 $\pm$.000	0.205 $\pm$.024	0.000 $\pm$.000	0.015 $\pm$.002
		Protein-reconstructed	0.250 $\pm$.024	0.070 $\pm$.020	0.000 $\pm$.000	0.290$\pm$.036	0.047 $\pm$.006	0.000 $\pm$.000	0.182 $\pm$.045	0.001 $\pm$.000
		Healthcare	0.121 $\pm$.014	0.034 $\pm$.018	0.143 $\pm$.007	0.000 $\pm$.000	0.000 $\pm$.000	0.008 $\pm$.001	0.268$\pm$.044	0.208 $\pm$.025
		Epidemiology	0.158 $\pm$.033	0.078 $\pm$.021	0.000 $\pm$.000	0.152 $\pm$.019	0.000 $\pm$.000	0.102 $\pm$.011	0.331$\pm$.034	0.000 $\pm$.000

The real or fitted-SCM datasets show even stronger metric dependence. Some methods remain nearly flat on particular datasets, which explains the zero or near-zero values in Table 6; Appendix A.2 defines these scores and reports the per-case diagnosis of all zero-score entries, including the cCBO Protein-reconstructed and MCBO small-budget cases. Ecology and Healthcare also show that strong performance can arise through different paths: some methods improve early and then plateau, while others discover a better intervention only after many trials.

Comparability across methods. The comparisons above are reported at the per-dataset level because the compared pipelines do not all solve the same optimization problem: cCBO solves a strictly harder constrained formulation, CoCaBO conditions its search on observed contextual information, DCBO runs here as a static-task stress test outside its temporal design setting, CEO spends part of its budget resolving graph uncertainty, and BO is a graph-free reference baseline. A single pooled ranking over these formulations would penalize methods for solving harder problems and reward methods that exploit additional information, so we restrict statements of algorithmic superiority to methods solving the same formulation (Table 5). Readers who want a cross-method overview can consult Appendix B.3.1, which reports budget-stratified average ranks with bootstrap confidence intervals and Friedman/Nemenyi tests as purely *descriptive* summaries: no pair of methods is statistically separable by average rank at any budget (the rank spread stays below the Nemenyi

critical difference throughout, and the conservative per-dataset Friedman test never rejects), subject to the non-independence caveats stated there.

4.5 Soft-intervention benchmark

For soft interventions, we compare BO, MCBO, and ACBO in Ackley, Rosenbrock, Dropwave, Alpine2 and Chain-soft (Sussex et al., 2023; 2024). BO serves as a non-causal baseline, while MCBO and ACBO represent mechanism-level and adversarial CBO-style approaches. Because ACBO was designed for adversarial or non-stationary settings, its use on these stationary function-network tasks is interpreted as a stress-test/reference result rather than a claim of intrinsic dominance in its native setting. fCBO also supports policy interventions, but it is excluded from the comparison because we did not obtain a compatible reference implementation (Gultchin et al., 2023).

Results summary. Because the soft benchmark combines two formulations: four computational-graph/function-network tasks and one SCM-based policy task, we report per-dataset comparisons rather than a pooled ordering. ACBO is especially strong on Ackley and Chain-soft, where it obtains the best GAP across budgets and often achieves strong PA-GAP values, indicating fast and sustained progress. A plausible explanation for why an adversarially designed method performs well on these stationary tasks is that, when the environment does not change, ACBO’s CBO-MW procedure effectively reduces to optimistic causal exploration: the causal UCB oracle exploits the known function-network structure much as MCBO does, while the multiplicative-weights layer over a discretized action set concentrates quickly on high-reward actions once rewards are stable, since no adversary forces re-exploration. The discretized action set may additionally act as a coarse global search grid that is well suited to highly multimodal landscapes such as Ackley, where continuous acquisition optimization can stall in local optima. We present this as a mechanism-level hypothesis consistent with the trajectories in Figure 4, not as a measured causal explanation. BO remains a highly competitive baseline on Rosenbrock and Alpine2. In these landscapes, standard black-box optimization often reaches strong final values and also receives favorable PA-GAP scores when it improves early and then maintains its best-so-far value. MCBO is most competitive on Dropwave, where it achieves the best GAP and PA-GAP at 50 and 20 trials; this pattern again reflects the metric distinction of Section 4.2. Overall, the soft-intervention results reinforce the main conclusion from the hard-intervention benchmark: causal modeling can improve optimization, but its advantage depends on how well the method’s assumptions match the intervention type, objective landscape, and evaluation metric.

Trajectory-level behavior. Figure 4 shows that the three methods behave differently in objective landscapes. ACBO improves quickly on Ackley and Chain-soft, which explains its strong scalar scores on these datasets. BO remains highly competitive in Rosenbrock and Alpine2, where its best-so-far trajectory often improves early enough to receive strong PA-GAP values as well as strong GAP values. Dropwave shows a complementary pattern: MCBO and ACBO both improve substantially, but MCBO is stronger under tighter budgets after applying the same PA-GAP clipping rule used throughout the benchmark. This correction caps the improvement ratio at one when a noisy observed best value exceeds the recorded expected-value reference optimum, ensuring that GAP and PA-GAP remain within their stated ranges. These trajectory patterns show why per-dataset conclusions differ across the two metrics: PA-GAP gives additional credit to methods that improve early and maintain progress, which is why BO remains competitive with ACBO on several tasks despite being a non-causal baseline.

Comparability across methods. As in the hard-intervention benchmark, statements of algorithmic superiority are restricted to matched formulations: Chain-soft supports only within-dataset ordering, and ACBO runs outside its adversarial design setting. The per-dataset results in Table 7 and Figure 4 are therefore the primary evidence, while the budget-stratified descriptive rank-reliability tables in Appendix B.3.2 provide a cross-method overview in which no pair of methods is statistically separable at any budget (here even the omnibus Friedman test does not reject).




Figure 4: Best-so-far trajectories on five soft-intervention datasets: Ackley, Rosenbrock, Dropwave, Alpine2, and Chain-soft. Solid lines show the mean over 20 random seeds and shaded regions denote the standard deviation across runs.

Table 7: Soft-intervention performance across different trial limits and datasets. Each value is the average $\pm$ standard error across 20 random seeds and initializations of the observational and interventional data. Higher is better; best mean values per row are in **bold**.

Trial limit	Metric	Dataset	BO	MCBO	ACBO
100	GAP	Ackley	0.328 $\pm$.022	0.310 $\pm$.003	0.922$\pm$.082
		Rosenbrock	0.825$\pm$.055	0.507 $\pm$.041	0.623 $\pm$.047
		Dropwave	0.024 $\pm$.005	0.558 $\pm$.065	0.690$\pm$.094
		Alpine2	0.536$\pm$.046	0.522 $\pm$.052	0.178 $\pm$.014
		Chain-soft	0.634 $\pm$.028	0.512 $\pm$.022	0.693$\pm$.052
	PA-GAP	Ackley	0.033 $\pm$.004	0.012 $\pm$.003	0.403$\pm$.006
		Rosenbrock	0.308 $\pm$.046	0.217 $\pm$.004	0.412$\pm$.007
		Dropwave	0.008 $\pm$.001	0.355 $\pm$.041	0.379$\pm$.005
50	GAP	Ackley	0.090 $\pm$.005	0.072 $\pm$.002	0.843$\pm$.071
		Rosenbrock	0.761$\pm$.042	0.515 $\pm$.049	0.631 $\pm$.052
		Dropwave	0.235 $\pm$.024	0.803$\pm$.085	0.318 $\pm$.040
		Alpine2	0.403$\pm$.034	0.085 $\pm$.007	0.175 $\pm$.009
		Chain-soft	0.515 $\pm$.023	0.000 $\pm$.000	0.655$\pm$.033
	PA-GAP	Ackley	0.018 $\pm$.003	0.000 $\pm$.000	0.320$\pm$.042
		Rosenbrock	0.239 $\pm$.041	0.053 $\pm$.017	0.344$\pm$.047
		Dropwave	0.008 $\pm$.001	0.318$\pm$.040	0.171 $\pm$.024
20	GAP	Ackley	0.166 $\pm$.013	0.000 $\pm$.000	0.593$\pm$.049
		Rosenbrock	0.559$\pm$.035	0.000 $\pm$.000	0.499 $\pm$.047
		Dropwave	0.252 $\pm$.023	0.701$\pm$.067	0.293 $\pm$.041
		Alpine2	0.142$\pm$.013	0.000 $\pm$.000	0.000 $\pm$.000
		Chain-soft	0.000 $\pm$.000	0.000 $\pm$.000	0.534$\pm$.024
	PA-GAP	Ackley	0.003 $\pm$.001	0.000 $\pm$.000	0.151$\pm$.017
		Rosenbrock	0.082 $\pm$.018	0.000 $\pm$.000	0.216$\pm$.038
		Dropwave	0.007 $\pm$.001	0.253$\pm$.030	0.130 $\pm$.021
		Alpine2	0.049$\pm$.006	0.000 $\pm$.000	0.000 $\pm$.000
		Chain-soft	0.000 $\pm$.000	0.000 $\pm$.000	0.116$\pm$.019

4.6 Cross-cutting analysis

The hard and soft-intervention benchmarks together span 78 dataset–budget–metric settings. Rather than summarizing each setting individually, we distill the main patterns into four cross-sectional observations.

When does causal structure help? Causal methods provide the clearest advantage when the graph enables effective scope reduction or when observational priors are well-calibrated. In the hard-intervention setting, CoCaBO and cCBO benefit from structured scope selection on datasets with multiple intervenable variables and clear causal pathways (Synthetic, Healthcare, Epidemiology). In the soft-intervention setting, ACBO benefits from mechanism-level modeling on datasets with complex intermediate structure (Ackley, Chain-soft). Conversely, causal structure provides little or no advantage on datasets where the effective intervention landscape is simple (Chain-hard) or where causal priors are poorly calibrated (Ecology under some methods). BO remains competitive precisely in these settings, suggesting that the value of causal modeling is conditional on the alignment between the method’s structural assumptions and the problem’s actual causal complexity.

Budget sensitivity. Rankings are moderately sensitive to budget. At tight budgets ($T = 20$), methods that rely on observational priors or scope reduction (CBO, cCBO, DCBO) tend to perform relatively better because their prior information is most valuable when interventional data are scarce. At larger budgets ($T = 100$), methods with richer exploration mechanisms (CoCaBO, ACBO) can overtake prior-driven methods by accumulating sufficient interventional evidence to overcome initial advantages from prior information.

Metric sensitivity. GAP and PA-GAP can produce different rankings for the same dataset and budget. GAP rewards the final best value and its discovery time, so methods that make a single large improvement early are favored. PA-GAP rewards steady trajectory-level progress, so methods that gradually improve throughout the budget can rank well even without the best final value. This discrepancy is most pronounced when comparing methods that plateau early (e.g., DCBO on Synthetic-2) versus methods that improve late (e.g., HCBO on Synthetic-2). We recommend that future CBO evaluations report both metrics, as well as trajectory plots, to avoid misleading conclusions from any single scalar summary.

Threats to validity. Several factors limit the generalizability of our benchmark findings. First, the benchmark is a native-wrapper reproducibility harness rather than a fully controlled reimplementations of every method; implementation details, released defaults, and dependency choices may affect performance. Second, the benchmark uses unit-cost interventions, which does not capture realistic cost heterogeneity. Third, the main hard-intervention tables assume benchmark-specified graphs and causal sufficiency for most graph-using methods. To probe this limitation rather than only state it, Appendix B.1 reports controlled learner-side graph-misspecification stress tests under edge addition, edge deletion, and edge reversal, and a companion omitted-variable stress test for hidden confounding. In all of these diagnostics, the true environment SCM and reference optimum remain fixed, while only the learner’s causal information is perturbed. Fourth, we use released default hyperparameters for all methods; per-dataset tuning could change relative rankings, as illustrated by the MCBO β sensitivity analysis. Relatedly, the benchmark does not control for differences in identification strategy across pipelines (Table 4): the released methods differ in whether and how they construct observational priors, none exposes a general front-door or full-ID implementation in the benchmark path used here, and Monte Carlo approximation settings for identified estimands are not uniformly exposed. On datasets with latent variables (most notably Synthetic, whose confounders U_1, U_2 make the backdoor criterion inapplicable for some effects), performance differences therefore cannot be cleanly attributed to acquisition or decision design alone; they may partly reflect differences in identification completeness and prior-construction quality. Fifth, GAP and PA-GAP are useful efficiency summaries but should be interpreted alongside best-so-far trajectories, simple-regret-equivalent summaries, average-reward diagnostics, and rank-reliability tests. Finally, the set of compared methods is limited by the availability of compatible public implementations; fCBO, CNEI, GACBO, and ACBO for hard interventions are not included in the hard-intervention comparison.

5 Open Problems and Future Directions

The benchmark results and methodological analysis presented in this survey reveal that CBO is a maturing but incomplete field: no single method dominates across settings (Section 4.6). To make the deployment gap concrete, consider using CBO to tune a treatment policy, industrial controller, or ecological intervention from observational data and a partially trusted causal graph. In such a setting, hidden confounding, graph errors, measurement noise, heterogeneous intervention costs, and safety constraints are not secondary details; any one of them can make an apparently promising intervention unreliable or harmful. In the following, we organize the most important open challenges into six concrete research directions. Each direction is stated relative to this deployment scenario: robustness to causal assumptions addresses the partially trusted graph and possible hidden confounding; scalability addresses the fact that realistic systems have far more variables and candidate scopes than current benchmarks; richer intervention models address mixed hard/soft/stochastic actions and heterogeneous intervention costs; realistic evaluation addresses the gap between benchmark-specified conditions and deployment conditions such as distribution shift; theoretical foundations address the guarantees a practitioner would need before acting on a recommended intervention; and integration with representation learning addresses settings where the causal variables themselves are not directly measured.

5.1 Robustness to causal assumptions

Hidden confounding. When unobserved confounders are present, observational priors constructed via backdoor adjustment or the g -formula can be systematically biased, leading the acquisition function to favor interventions that appear promising under confounded associations but fail under true causal effects. Developing CBO methods that explicitly model sensitivity to hidden confounding, for example, by propagating bounds on the causal effect under varying degrees of unmeasured confounding (Pearl, 2009), is an important open direction.

Graph misspecification. Even when a graph is provided, edges may be missing, spurious, or incorrectly oriented. The impact of such errors on CBO performance has been largely unexplored. Future work should characterize how graph misspecification degrades scope reduction and prior quality and develop methods that are robust to bounded graph errors.

Prior misspecification. CNEI (Li et al., 2023) addresses noise robustness, but the broader problem of controlling how strongly observational information influences acquisition decisions under potential misspecification remains open. Methods that adaptively downweight observational priors as interventional data accumulate, analogous to Bayesian robustness techniques in other domains, would be valuable.

5.2 Scalability

High-dimensional graphs. HCBO (Wu et al., 2024) addresses the selection of the scope in larger graphs, but the combinatorial explosion of candidate scopes remains a fundamental bottleneck. Future methods should take advantage of causal sparsity, hierarchical decomposition, or learned embeddings of the intervention scopes to avoid exhaustive enumeration. Notably, hierarchical decomposition across levels of causal granularity has already been formalized in bandit settings by the causal abstraction literature (Section 3.7); transferring these abstraction-based pruning and transfer mechanisms to GP-based CBO with continuous intervention domains is a concrete and promising path toward scalability.

Unknown-graph scalability. CEO (Branchini et al., 2023) and GACBO (Mukherjee et al., 2024) demonstrate unknown-graph CBO on small graphs, but maintaining a posterior or confidence set over graphs becomes intractable as the number of variables grows. Scalable alternatives might include local structure learning focused on the decision-relevant subgraph, amortized inference over graph posteriors, or graph neural network-based surrogates that learn causal structure implicitly.

Computational cost. Mechanism-level methods (MCBO, ACBO) require maintaining and updating multiple GPs and propagating uncertainty through the graph at each iteration. The per-iteration cost scales with the number of mechanism models and the depth of the graph. Sparse GP approximations, variational inference, or neural process surrogates could reduce this computational burden.

5.3 Richer intervention models

Mixed intervention types. Many practical systems involve a mix of hard interventions (e.g. turning a device on or off), soft interventions (e.g., adjusting a controller’s parameters) and stochastic interventions (e.g., randomizing a treatment assignment probability). No existing CBO method handles arbitrary mixtures of these types within a single optimization loop.

Multi-agent and sequential interventions. ACBO (Sussex et al., 2024) considers adversarial environments, but the more general setting of multi-agent intervention, where multiple decision-makers intervene in the same system, possibly with conflicting objectives, remains unexplored. Similarly, extending CBO to sequential multi-stage intervention policies, where the intervention at each stage depends on the outcomes of previous stages, would bridge CBO and causal reinforcement learning.

Intervention cost heterogeneity. Our benchmark uses unit-cost interventions to isolate sample efficiency, but real interventions differ substantially in cost. The cost may depend on the target variable, the magnitude of the intervention, the safety requirements, or the logistics of implementation. Future CBO benchmarks should incorporate explicit cost models and future methods should optimize cost-adjusted efficiency rather than pure sample efficiency.

5.4 Realistic evaluation

Standardized benchmarks. The fragmentation of the data set documented in Table 3 makes the comparison between papers unreliable. The community would benefit from a shared and versioned benchmark suite with standardized SCMs, intervention domains, observational datasets, and evaluation code, analogous to established benchmarks in reinforcement learning or supervised learning. In this spirit, the benchmark released with this survey should be read as a unified repository of existing environments together with an execution and scoring harness, not as a definitive benchmark suite. Establishing the latter would additionally require high-dimensional, complex, nonlinear simulators with rich causal structure among manipulable variables (e.g., realistic epidemiological or climate models) to stress-test scalability, combinatorial scope reduction, and robustness; building such simulators is an important direction for future work.

Evaluation under distribution shift. Most benchmarks assume that the test-time SCM matches the training-time SCM. In practice, the system can change between the observational data collection phase and the intervention deployment phase. Evaluating CBO under a controlled distribution shift would test the robustness of observational priors and structure assumptions.

Trajectory-aware metrics. Our analysis shows that GAP and PA-GAP can produce different rankings, highlighting the need for trajectory-aware evaluation. Future work should develop principled metrics that capture the full optimization trajectory, potentially incorporating cost, risk, and constraint satisfaction alongside objective improvement.

5.5 Theoretical foundations

Finite-sample regret bounds. MCBO (Sussex et al., 2023) provides regret bounds under mechanism-level GP models, and ACBO (Sussex et al., 2024) extends these to adversarial settings. However, regret bounds for effect-level CBO with observational priors, for unknown-graph CBO, and for constrained CBO are largely missing. Developing such bounds would clarify when causal structure provably improves over black-box BO and under what conditions.

Identifiability-aware acquisition. Current methods either assume full identifiability or ignore identifiability entirely. An intermediate approach would design acquisition functions that explicitly account for the degree of identifiability: using observational information aggressively when identification is strong and falling back to interventional exploration when it is weak. This would connect CBO to the literature on partial identification and sensitivity analysis in causal inference.

Sample complexity of scope reduction. The POMIS-based scope reduction in CBO eliminates dominated scopes using observational evidence, but the sample complexity of this elimination, i.e., how many observational samples are needed to reliably identify the correct POMIS, has not been analyzed. Such an analysis would clarify when scope reduction is reliable and when it may prematurely discard useful scopes.

5.6 Integration with modern representation learning

Causal representation learning. Methods for learning causal variables and structure from high-dimensional observations (Schölkopf et al., 2021) could enable CBO in settings where the causal graph operates over latent variables rather than directly observed quantities. This would extend CBO to image- or text-based intervention settings. A complementary and arguably more direct route is causal abstraction learning, which relates a fine-grained causal model to a coarser one through abstraction maps; integrating learned abstraction maps into the CBO loop (Section 3.7) is a concrete instance of the latent-variable optimization advocated here.

Large language models for causal reasoning. Large language models (LLMs) have shown preliminary ability to reason about causal relationships from text descriptions. An intriguing direction is to use LLMs to elicit or refine causal graphs from domain experts, providing CBO with better structural priors without requiring formal causal discovery from data.

Neural surrogates. Replacing GP-based surrogates with neural network-based models (e.g., neural processes, transformers) could improve scalability and flexibility, particularly for high-dimensional mechanism functions or non-stationary environments. The challenge is maintaining the well-calibrated uncertainty quantification that makes GP-based CBO effective.

6 Conclusion

This survey has provided a comprehensive review of Causal Bayesian Optimization, organizing the growing literature through a unified design-space perspective that separates graph and system-knowledge assumptions, environmental assumptions, intervention representation, surrogate architecture, and decision rules. By analyzing each major CBO variant through this lens, we have shown that the field’s diversity reflects systematic responses to specific bottlenecks of the original CBO framework (temporal dynamics, safety constraints, high dimensionality, unknown graphs, soft interventions, context dependence, and adversarial non-stationarity) rather than ad hoc extensions.

Our benchmark study reinforces a central finding: the value of causal structure in optimization is conditional. Causal methods provide clear advantages when their structural assumptions align with the problem, when scope reduction eliminates irrelevant interventions, when observational priors are well-calibrated, or when mechanism-level modeling captures the system’s compositional structure. However, when these conditions are not met, strong non-causal BO baselines remain competitive, and method rankings depend substantially on the choice of dataset, budget, and evaluation metric. The introduction of PA-GAP as a trajectory-aware complement to GAP highlights that even the definition of “good performance” in CBO is not straightforward: final-value metrics and trajectory-level metrics can favor different methods.

Looking ahead, the most pressing challenges are not purely algorithmic. Our benchmark does not establish that CBO methods are robust in high-stakes deployment settings; rather, it shows that causal methods can be useful under benchmark-specified assumptions and that performance can change substantially when learner-side causal assumptions are perturbed. Robustness to graph misspecification, hidden confounding, imperfect interventions, measurement error, heterogeneous intervention costs, and safety constraints remains

a prerequisite for CBO to move from controlled benchmarks to real-world deployment. We hope that the unified perspective, benchmark, stress-test protocols, and structured open problems presented in this survey will help researchers navigate the CBO landscape and focus their efforts on directions with practical impact.

Broader impact and deployment caution. CBO is increasingly positioned for deployment in sensitive domains such as healthcare dosing, resource allocation, and public policy. Deploying CBO with misspecified graphs, hidden confounding, or miscalibrated causal priors can select harmful physical interventions in real-world systems. A practitioner reading this survey should not conclude that CBO methods offer reliable advantages in such settings; the evidence applies only to the correctly specified, benchmark-controlled regime. Causal assumptions and failure logging should be treated as part of deployment risk, not merely as technical details. The graph-misspecification and omitted-variable stress tests in Appendix B.1 are a first step toward the robustness evaluation needed for deployment, but comprehensive robustness across the dimensions listed above remains an open evaluation dimension.

References

- Virginia Aglietti, Xiaoyu Lu, Andrei Paleyes, and Javier González. Causal Bayesian optimization. In Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3155–3164. PMLR, 26–28 Aug 2020.
- Virginia Aglietti, Neil Dhir, Javier González, and Theodoros Damoulas. Dynamic causal Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 34, pp. 10549–10560. Curran Associates, Inc., 2021.
- Virginia Aglietti, Alan Malek, Ira Ktena, and Silvia Chiappa. Constrained causal Bayesian optimization. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 304–321. PMLR, 23–29 Jul 2023.
- Vahan Arsenyan, Antoine Grosnit, Haitham Bou Ammar, and Arnak S. Dalalyan. Contextual causal Bayesian optimisation. In *International Conference on Learning Representations*. OpenReview.net, 2026. ICLR 2026 Poster.
- Raul Astudillo and Peter I. Frazier. Bayesian optimization of function networks. In *Advances in Neural Information Processing Systems*, volume 34, pp. 14463–14475, 2021.
- Felix Berkenkamp, Angela P. Schoellig, and Andreas Krause. No-regret Bayesian optimization with unknown hyperparameters. *Journal of Machine Learning Research*, 20(50):1–24, 2019.
- Shriya Bhatija, Paul-David Zuercher, Jakob Thumm, and Thomas Bohné. Multi-objective causal Bayesian optimization. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 4172–4195. PMLR, 13–19 Jul 2025.
- Nicola Branchini, Virginia Aglietti, Neil Dhir, and Theodoros Damoulas. Causal entropy optimization. In Francisco Ruiz, Jennifer Dy, and Jan-Willem van de Meent (eds.), *Proceedings of the 26th International Conference on Artificial Intelligence and Statistics*, volume 206 of *Proceedings of Machine Learning Research*, pp. 8586–8605. PMLR, 25–27 Apr 2023.
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995.
- David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, 2002.

- Sayak Ray Chowdhury and Aditya Gopalan. On kernelized multi-armed bandits. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pp. 844–853. PMLR, 2017.
- Travis A. Courtney, Mario Lebrato, Nicholas R. Bates, Andrew Collins, Samantha J. de Putron, Rebecca Garley, Rod Johnson, Juan-Carlos Molinero, Timothy J. Noyes, Christopher L. Sabine, and Andreas J. Andersson. Environmental controls on modern scleractinian coral and reef-scale calcification. *Science Advances*, 3(11):e1701356, 2017. doi: 10.1126/sciadv.1701356.
- Alexander I. Cowen-Rivers, Wenlong Lyu, Rasul Tutunov, Zhi Wang, Antoine Grosnit, Ryan Rhys Griffiths, Alexandre Max Maraval, Hao Jianye, Jun Wang, Jan Peters, and Haitham Bou-Ammar. Hebo: Pushing the limits of sample-efficient hyper-parameter optimisation. *Journal of Artificial Intelligence Research*, 74: 1269–1349, 2022. doi: 10.1613/jair.1.13643.
- Jean Durand, Yashas Annadani, Stefan Bauer, and Sonali Parbhoo. Causal Bayesian optimization with unknown graphs, 2025.
- Joel Dyer, Nicholas Bishop, Anisoara Calinescu, Michael Wooldridge, and Fabio Massimo Zennaro. Using causal abstractions to accelerate decision-making in complex bandit problems. *arXiv preprint arXiv:2509.04296*, 2025. doi: 10.48550/arXiv.2509.04296.
- Frederick Eberhardt and Richard Scheines. Interventions and causal inference. *Philosophy of Science*, 74(5): 981–995, 2007. doi: 10.1086/525638.
- Ana Ferro, Francisco Pina, Milton Severo, Pedro Dias, Francisco Botelho, and Nuno Lunet. Use of statins and serum levels of prostate specific antigen. *Acta Urológica Portuguesa*, 32(2):71–77, 2015. doi: 10.1016/j.acup.2015.02.002.
- Adam Foster, Desi R. Ivanova, Ilyas Malik, and Tom Rainforth. Deep adaptive design: Amortizing sequential Bayesian experimental design. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 3384–3395. PMLR, 2021.
- Peter I. Frazier. A tutorial on Bayesian optimization. *arXiv preprint arXiv:1807.02811*, 2018. doi: 10.48550/arXiv.1807.02811.
- Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997. doi: 10.1006/jcss.1997.1504.
- Roman Garnett. *Bayesian Optimization*. Cambridge University Press, 2023. doi: 10.1017/9781108348973.
- Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10:524, 2019. doi: 10.3389/fgene.2019.00524.
- Limor Gultchin, Virginia Aglietti, Alexis Bellot, and Silvia Chiappa. Functional causal Bayesian optimization. In Robin J. Evans and Ilya Shpitser (eds.), *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pp. 756–765. PMLR, 31 Jul–04 Aug 2023.
- William G. Havercroft and Vanessa Didelez. Simulating from marginal structural models with time-dependent confounding. *Statistics in Medicine*, 31(30):4190–4206, 2012. doi: 10.1002/sim.5472.
- Miguel A. Hernán and James M. Robins. *Causal Inference: What If*. Chapman and Hall/CRC, 2020.
- Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient global optimization of expensive black-box functions. *Journal of Global Optimization*, 13(4):455–492, 1998. doi: 10.1023/A:1008306431147.
- Mina Konakovic Lukovic, Yunsheng Tian, and Wojciech Matusik. Diversity-guided multi-objective Bayesian optimization with batch evaluations. In *Advances in Neural Information Processing Systems*, volume 33, pp. 17708–17720. Curran Associates, Inc., 2020.

- Harold J. Kushner. A new method of locating the maximum point of an arbitrary multipeak curve in the presence of noise. *Journal of Basic Engineering*, 86(1):97–106, 1964. doi: 10.1115/1.3653121.
- Finnian Lattimore, Tor Lattimore, and Mark D. Reid. Causal bandits: Learning good interventions via causal inference. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- Sanghack Lee and Elias Bareinboim. Structural causal bandits: Where to intervene? In *Advances in Neural Information Processing Systems*, volume 31, pp. 2573–2583, 2018.
- Sanghack Lee and Elias Bareinboim. Structural causal bandits with non-manipulable variables. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 4164–4172, 2019. doi: 10.1609/aaai.v33i01.33014164.
- Zheng Li, Yongxuan Li, Zongqiang Lian, and Rui Zheng. Improved causal Bayesian optimization algorithm with counter-noise acquisition function and supervised prior estimation. *Journal of Physics: Conference Series*, 2547(1):012017, 2023. doi: 10.1088/1742-6596/2547/1/012017.
- Nick Littlestone and Manfred K. Warmuth. The weighted majority algorithm. *Information and Computation*, 108(2):212–261, 1994. doi: 10.1006/inco.1994.1009.
- Yangyi Lu, Amirhossein Meisami, and Ambuj Tewari. Causal bandits with unknown graph structure. In *Advances in Neural Information Processing Systems*, volume 34, pp. 24817–24828. Curran Associates, Inc., 2021.
- Jonas Mockus. *Bayesian Approach to Global Optimization*. Kluwer Academic Publishers, 1989. doi: 10.1007/978-94-009-0909-0.
- Jonas Mockus, Vytautas Tiesis, and Antanas Zilinskas. The application of Bayesian methods for seeking the extremum. In *Towards Global Optimization 2*, pp. 117–129. North-Holland, 1978.
- Sumantrak Mukherjee, Mengyan Zhang, Seth Flaxman, and Sebastian Josef Vollmer. Graph agnostic causal Bayesian optimisation, 2024.
- Vineet Nair, Vishakha Patil, and Gaurav Sinha. Budgeted and non-budgeted causal bandits. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pp. 2017–2025. PMLR, 2021.
- Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2 edition, 2009.
- Judea Pearl. The do-calculus revisited. In *Proceedings of the Twenty-Eighth Conference on Uncertainty in Artificial Intelligence*, 2012.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, 2017.
- Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.
- Shaogang Ren, Zihao Wang, Yuzhou Chen, and Xiaoning Qian. Causal Bayesian optimization via exogenous distribution learning, 2024.
- Karen Sachs, Omar Perez, Dana Pe’er, Douglas A. Lauffenburger, and Garry P. Nolan. Causal protein-signaling networks derived from multiparameter single-cell data. *Science*, 308(5721):523–529, 2005. doi: 10.1126/science.1105809.
- Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. Toward causal representation learning. *Proceedings of the IEEE*, 109(5): 612–634, 2021. doi: 10.1109/JPROC.2021.3058954.

- Pier Giuseppe Sessa, Ilija Bogunovic, Maryam Kamgarpour, and Andreas Krause. No-regret learning in unknown games with correlated payoffs. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- Bobak Shahriari, Kevin Swersky, Ziyu Wang, Ryan P. Adams, and Nando de Freitas. Taking the human out of the loop: A review of Bayesian optimization. *Proceedings of the IEEE*, 104(1):148–175, 2016. doi: 10.1109/JPROC.2015.2494218.
- Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2 edition, 2000.
- Niranjan Srinivas, Andreas Krause, Sham M. Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proceedings of the 27th International Conference on Machine Learning*, pp. 1015–1022, 2010.
- Yanan Sui, Alkis Gotovos, Joel W. Burdick, and Andreas Krause. Safe exploration for optimization with Gaussian processes. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pp. 997–1005. PMLR, 2015.
- Scott Sussex, Anastasia Makarova, and Andreas Krause. Model-based causal Bayesian optimization. In *International Conference on Learning Representations*. OpenReview.net, 2023. ICLR 2023, notable top 25% (spotlight).
- Scott Sussex, Pier Giuseppe Sessa, Anastasia Makarova, and Andreas Krause. Adversarial causal Bayesian optimization. In *International Conference on Learning Representations*. OpenReview.net, 2024. ICLR 2024.
- Clay Thompson. Causal graph analysis with the CAUSALGRAPH procedure. In *Proceedings of the SAS Global Forum 2019 Conference*. SAS Institute Inc., 2019. Paper SAS2998-2019.
- Yupeng Wu, Weiye Wang, Yangwenhui Zhang, Mingjia Li, Yuanhao Liu, Hong Qian, and Aimin Zhou. High-dimensional causal Bayesian optimization. In *ECAI 2024 – 27th European Conference on Artificial Intelligence*, volume 392 of *Frontiers in Artificial Intelligence and Applications*, pp. 2990–2997. IOS Press, 2024. doi: 10.3233/FAIA240839.
- Jakob Zeitler. Towards MFACBO: Multi-fidelity abstraction causal Bayesian optimization in the context of the abstraction-fidelity connection. Workshop on Causal Abstractions and Representations (CAR) at UAI 2025 Poster, 2025.
- Fabio Massimo Zennaro, Nicholas Bishop, Joel Dyer, Yorgos Felekis, Anisoara Calinescu, Michael Wooldridge, and Theodoros Damoulas. Causally abstracted multi-armed bandits. In Negar Kiyavash and Joris M. Mooij (eds.), *Proceedings of the Fortieth Conference on Uncertainty in Artificial Intelligence*, volume 244 of *Proceedings of Machine Learning Research*, pp. 4109–4139. PMLR, 15–19 Jul 2024.
- A. G. Zhilinskias. Single-step bayesian search method for an extremum of functions of a single variable. *Cybernetics*, 11(1):160–166, 1975. doi: 10.1007/BF01069961.

Table of Contents

A	Technical Appendix	41
A.1	Metrics analysis	41
A.1.1	GAP value range	41
A.1.2	Why PA-GAP can rank trajectories more consistently	41
A.2	Implementation protocol	43
A.3	Reference optima	45
A.4	Consolidated dataset summary	46
A.5	Synthetic dataset SCMs	49
A.5.1	ToyGraph	49
A.5.2	Synthetic.	49
A.5.3	Synthetic-2	50
A.5.4	Chain-hard	51
A.6	Real dataset SCMs	51
A.6.1	Ecology	52
A.6.2	Protein-reconstructed	52
A.6.3	Healthcare	53
A.6.4	Epidemiology	54
A.7	Soft-intervention datasets	55
A.7.1	Ackley	56
A.7.2	Rosenbrock	56
A.7.3	Dropwave	57
A.7.4	Alpine2.	58
A.7.5	Chain-soft	58
B	Additional Results	60
B.1	Robustness Stress Tests.	60
B.1.1	Graph-misspecification stress test	60
B.1.2	Hidden confounding (omitted-variable stress test).	67
B.2	MCBO β sensitivity	68
B.3	Ranking Statistical Testing	68
B.3.1	Hard-intervention	69
B.3.2	Soft-intervention	69
B.4	Average Reward Visualization	69

A Technical Appendix

A.1 Metrics analysis

DCBO introduced the GAP metric as a scalar that combines how much an optimizer improves on its initial value with how early it reaches its best value (Aglietti et al., 2021). For a minimization problem, let $y(\mathbf{x}) = \mathbb{E}[Y \mid \text{do}(X = \mathbf{x})]$, let y^* denote the reference optimum (see Appendix A.3), let $\mathbf{x}_{\text{init}}$ be the initial best point, and let $\mathbf{x}_t^*$ be the best point found up to trial t . Let T be the total number of trials, and let t^* be the first trial at which the final best value $y(\mathbf{x}_T^*)$ is reached. GAP is

$$\text{GAP} = \left[\underbrace{\frac{y(\mathbf{x}_T^*) - y(\mathbf{x}_{\text{init}})}{y^* - y(\mathbf{x}_{\text{init}})}}_{\text{Improvement term}} + \underbrace{\frac{T - t^*}{T}}_{\text{Efficiency term}} \right] / \underbrace{\left(1 + \frac{T - 1}{T}\right)}_{\text{Normalization}}. \quad (43)$$

A.1.1 GAP value range

When the global optimum is found in the first trial ($t^* = 1$), the efficiency term is equal to $\frac{T-1}{T}$, and normalization makes the maximum GAP equal to 1. However, DCBO sets $t^* = 0$ when no improvement is observed. In that case, the improvement term is 0, but the efficiency term becomes 1, which produces a nonzero GAP even without improvement. Specifically, GAP is then lower-bounded by

$$\frac{1}{1 + \frac{T-1}{T}} = \frac{T}{2T - 1}. \quad (44)$$

To correct for this, we set $t^* = T$ when no improvement occurs. Then both the improvement term and the efficiency term are 0, so the GAP ranges over $[0, 1]$ as intended.

A.1.2 Why PA-GAP can rank trajectories more consistently

GAP depends only on the final best value and the trial in which that final best value is first reached. This can bias the rankings when two optimization runs share the same early progress but differ in later improvements. Figure 5 illustrates this with a two-model example: both models reach the same improved value at trial 10, but only Model 2 later reaches the global optimum (at trial 30 of $T = 40$), and GAP nonetheless scores Model 1 higher because it evaluates only the final best value and its discovery time rather than the full optimization path.

To better reflect the full optimization process, PA-GAP computes a weighted improvement value at each trial and then averages these values over the trajectory. For minimization tasks, the best-so-far improvement ratio in the trial t is

$$R_t = \min \left\{ \frac{y(\mathbf{x}_{\text{init}}) - y(\mathbf{x}_t^*)}{y(\mathbf{x}_{\text{init}}) - y^*}, 1 \right\}, \quad (45)$$

and for maximization tasks, the numerator and denominator are reversed:

$$R_t = \min \left\{ \frac{y(\mathbf{x}_t^*) - y(\mathbf{x}_{\text{init}})}{y^* - y(\mathbf{x}_{\text{init}})}, 1 \right\}. \quad (46)$$

The clipping prevents numerical artifacts when an observed best value exceeds the offline reference optimum. The PA-GAP score used in our benchmark is then

$$\text{PA-GAP} = \frac{1}{T} \sum_{t=1}^T \left(\underbrace{R_t}_{\text{Improvement ratio}} \cdot \underbrace{\frac{T - (t - 1)}{T}}_{\text{Efficiency weight}} \right). \quad (47)$$

Here T denotes the number of interventional trials, excluding the initial value in the best-so far trajectory.




Figure 5: Demo comparison illustrating the ranking bias of GAP. The left panel shows best-so-far trajectories, while the right panel shows the corresponding per-trial PA-GAP (PA-GAP@t) contributions. The two models share the same early improvement, but Model 2 later reaches the global optimum. GAP can favor the earlier but worse final trajectory, while PA-GAP assigns credit to later improvement by averaging weighted best-so-far progress across the full path.




Figure 6: HCBO and DCBO on Synthetic-2. The left panel shows best-so-far trajectories, while the right panel shows the corresponding per-trial PA-GAP (PA-GAP@t) contributions. DCBO reaches its own best-so-far value early, which gives it a higher GAP score, but HCBO reaches better objective values over the trajectory and receives the higher PA-GAP score.

This definition is not normalized to have a maximum value 1. Since $R_t \in [0, 1]$ and the efficiency weight decreases linearly from 1 at the first trial to $1/T$ at the last trial, the raw PA-GAP range is

$$0 \leq \text{PA-GAP} \leq \frac{1}{T} \sum_{t=1}^T \frac{T - (t - 1)}{T} = \frac{T + 1}{2T}. \quad (48)$$

The upper bound is attained when the global optimum is found in the first trial and remains the best value so far for the rest of the budget. Thus, PA-GAP should be interpreted as a raw trajectory-weighted efficiency score, not as a unit-normalized score. Larger values indicate faster and stronger best-so-far progress, but raw PA-GAP values should be compared directly only within the same trial budget.

Figure 6 shows the same issue on the Synthetic-2 benchmark: DCBO quickly reaches a stable best-so-far value, so GAP rewards its early discovery ($0.485 \pm .012$ versus $0.174 \pm .071$ for HCBO), but that early value is much worse than the values HCBO eventually reaches, and PA-GAP reverses the ranking ($0.002 \pm .001$ for DCBO versus $0.099 \pm .054$ for HCBO), better reflecting HCBO’s sustained progress toward the reference optimum.

Summary and formal properties. GAP is simple and interpretable, but can produce counterintuitive rankings when methods differ primarily in their late-stage behavior or trajectory shape. PA-GAP comple-

Table 8: Formal comparison of GAP and PA-GAP evaluation properties.

Property	GAP	PA-GAP
What is evaluated	Final best value + discovery time	Full best-so-far trajectory
Value range	$[0, 1]$ (after correction)	$[0, \frac{T+1}{2T}]$
Normalization	Unit-normalized	Raw; budget-dependent upper bound
Sensitivity to late improvement	Low: only the time of the <i>final</i> best matters	High: all improvements contribute
Sensitivity to budget T	Moderate: efficiency term scales with T	Strong: values not comparable across different T
Plateau handling	Rewards early plateau at good value	Penalizes early plateau if further improvement is possible
Computational cost	$\mathcal{O}(T)$	$\mathcal{O}(T)$

ments GAP by rewarding sustained progress throughout the budget. Table 8 contrasts the two metrics along key evaluation dimensions.

We state two key properties of PA-GAP that justify its use as a trajectory-aware complement to GAP.

Lemma 1 (Trajectory monotonicity of PA-GAP). *Let $\{R_t\}_{t=1}^T$ and $\{\tilde{R}_t\}_{t=1}^T$ be two sequences of the best improvement ratios so far satisfying $R_t \geq \tilde{R}_t$ for all $t \in \{1, \dots, T\}$. Then $\text{PA-GAP}(\{R_t\}) \geq \text{PA-GAP}(\{\tilde{R}_t\})$, with equality if and only if $R_t = \tilde{R}_t$ for all t .*

Proof. Since the efficiency weights $w_t = \frac{T-(t-1)}{T} > 0$ for all t , the PA-GAP functional $\frac{1}{T} \sum_{t=1}^T R_t \cdot w_t$ is a positively weighted average of the improvement ratios. The conclusion follows from the strict positivity of the weights. $\square$

This property ensures that if one trajectory Pareto-dominates another at every trial, PA-GAP reflects this dominance. GAP does not satisfy this property in general: a trajectory with uniformly higher R_t can receive a lower GAP score if its final best is reached later.

Lemma 2 (Budget dependence of PA-GAP). *The maximum attainable PA-GAP value is $\frac{T+1}{2T}$, which is strictly decreasing in T and approaches $\frac{1}{2}$ as $T \rightarrow \infty$. Consequently, raw PA-GAP values are comparable only within the same budget T .*

Proof. The maximum is achieved when $R_t = 1$ for all t , giving $\text{PA-GAP} = \frac{1}{T} \sum_{t=1}^T \frac{T-(t-1)}{T} = \frac{T+1}{2T}$. Since $\frac{d}{dT} \left(\frac{T+1}{2T} \right) = -\frac{1}{2T^2} < 0$, the upper bound is strictly decreasing in T . $\square$

Practical recommendation. We recommend reporting both metrics alongside trajectory plots: GAP identifies methods that quickly find a strong final solution, while PA-GAP evaluates whether progress is sustained throughout the budget.

A.2 Implementation protocol

Reproducibility philosophy. All methods run with released default hyperparameters unless noted in Table 9. These are wrapper-default reproducibility results, not per-dataset tuned head-to-head comparisons. This choice prioritizes reproducibility and isolates the behavior of public implementations, but it can understate what a carefully tuned method might achieve. Appendix B.2 illustrates this sensitivity for MCBO: the best UCB β varies across datasets, budgets, and metrics, so the main table keeps the wrapper setting $\beta = 10$ and reports the sweep separately.

Software environment. All experiments are conducted in Python 3.10 using PyTorch 2.0 and GPyTorch 1.11 as the primary GP backend for methods that depend on Gaussian process surrogates. Each

Table 9: Implementation settings used in the benchmark. Method-specific defaults are preserved unless a wrapper override is needed for shared evaluation.

Method	Benchmark-controlled settings	Method-specific notes
BO	Trial budget, seed, initial design, objective direction, and intervention domain fixed by wrapper; no graph input used.	Graph-free GP-EI baseline on the joint intervention vector: GPy GPRRegression with RBF kernel (lengthscale = 1, variance = 1, ARD disabled), fixed observation noise 10^{-10} , emukit Expected Improvement, and GradientAcquisitionOptimizer. All methods, including BO, use the same fixed initial interventional design per seed and the same intervention parameterization and optimization direction.
MCBO	Trial budget and seed fixed by wrapper; noise scale set to 0; $\beta = 10$.	Uses GP networks with MC acquisition. Acquisition optimization follows the released implementation defaults, including restart and raw-sample rules.
ACBO	Trial budget and seed fixed by wrapper; noise scale set to 0; $\beta = 10$; default discrete action cardinality 4.	Used only for soft-intervention function-network tasks. Runs CBO-MW with multiplicative-weights updates.
cCBO	Trial budget fixed by wrapper; dataset configs determine observational samples, intervention sets, and constraints.	Uses the single-task cCBO setting in our benchmark. Protein uses an internal offset so that the exported trajectory contains the intended number of evaluated trials.
DCBO	Trial budget and seed fixed by wrapper; dataset YAML files specify the SCM and intervention domain.	Benchmark uses a single-time setting to make DCBO comparable with static hard-intervention tasks.
CEO	Trial budget and seed fixed by wrapper; small candidate graph and anchor-point settings follow the benchmark runner.	Uses the benchmark-native CEO runner by default. The original vendored CEO implementation can be enabled separately.
CoCaBO	Trial budget, seed, and objective direction fixed by wrapper.	Uses UCB over active policy scopes and HEBO-style inner optimization. Benchmark configs do not include additional contextual variables unless specified.
HCBO	Seed and trial settings fixed by wrapper when supported.	Used for compatible hard-intervention settings. Acquisition and initialization constants follow the released HCBO-style implementation.

CBO method is run in its own isolated virtual environment to prevent dependency conflicts. The benchmark evaluation pipeline—including trajectory export, GAP/PA-GAP computation, and figure generation—is implemented in a shared codebase that is independent of any individual method’s software stack. This separation ensures that scoring differences reflect algorithmic behavior rather than library version effects. The exact package versions for each method are documented in the benchmark repository.

Computational resources. All experiments were conducted on a single-GPU CUDA workstation equipped with one NVIDIA A100 GPU with 40 GiB of VRAM, Intel Xeon Platinum 8481C CPUs, and 754 GiB of system memory. Most methods complete 100 trials on a single dataset within 5–30 minutes. The main exceptions are CEO, which incurs additional overhead from graph posterior sampling, and MCBO, which propagates uncertainty through multiple mechanism-level GPs. GPU acceleration is used when supported by the corresponding implementation; otherwise, experiments are executed on the CPU.

Wrapper design. Wrapper-level overrides are limited to shared experimental controls: trial budget, random seed, dataset mapping, objective direction (minimization or maximization), and trajectory export format. Method-internal choices, including the selection of the GP kernel, the acquisition optimization strategy, the restart schedules, the scope enumeration, and the internal noise models, are left at their released defaults.

Each wrapper converts the method’s internal output into the common trajectory format described below. Wrappers do not modify the optimization loop itself; they only control initialization and output extraction.

Trajectory normalization. To make results comparable across heterogeneous codebases, every method’s output is converted to a common trajectory format with two fields: trial number and best-so-far objective value. For minimization tasks, this value is the running minimum observed up to the trial t ; for maximization tasks, the running maximum. The trajectory includes an initial best-of-so-far value (before any interventional trial), followed by one entry per interventional trial. Thus, a budget of B interventions produces a trajectory of $B + 1$ entries. GAP and PA-GAP are then computed from these exported trajectories using a single evaluation script shared by all methods, ensuring that scoring differences reflect algorithmic behavior rather than implementation artifacts. This design also makes it easy to add new methods or metrics without modifying existing implementations.

Zero-score diagnosis. A zero or near-zero GAP or PA-GAP value in the main tables means that the method did not improve over initialization under the shared scoring protocol; it does not mean that the run was silently removed (genuine failures such as missing outputs or non-finite trajectories are logged separately under the failure-handling protocol). In principle, a zero score could arise from (a) a poorly calibrated observational prior, (b) a loose or unreliable reference optimum, or (c) graph-induced scope selection that excludes useful interventions. In the current logs, the dominant cause is failure to improve over the initial design; in a smaller number of cases, the zero score reflects method–task incompatibility, unsupported or empty candidate scopes (cause c), or poor calibration of the observational prior (cause a). We found no case attributable to an unreliable reference optimum (cause b): the protocol in Appendix A.3 computes references offline on the ground-truth SCM, and reference error would rescale all methods’ scores on a dataset rather than zero out a single method. Two named cases illustrate the distinction. cCBO on Protein-reconstructed completes all 20 seeds and returns valid, finite trajectories, but its best-so-far value remains flat over the entire budget under the reconstructed SCM and its wide Sachs-derived intervention domains, so GAP and PA-GAP are zero by definition; the runs do not crash, and the discrepancy with the original cCBO Protein results is consistent with our task being a transparent reconstruction rather than the unreleased fitted SCM used in the original experiments (Appendix A.6.2). MCBO’s zero entries at $T = 20$ and $T = 50$ on several datasets reflect delayed improvement: the same runs reach non-zero scores at $T = 100$, consistent with a method that requires many trials before improving on the initial design.

A.3 Reference optima

The reference optimum y^* used by GAP and PA-GAP is computed offline from the SCM benchmark, function generator or oracle interventional data, as detailed in Table 10. It is used *only* for scoring and is never provided to any optimizer during a run. At evaluation time, all methods on a given data set are compared against the same fixed reference value using the same task direction.

Reliability categories. The reference optima fall into two reliability classes. For datasets with analytically known optima (Dropwave, Alpine2, Ackley, Rosenbrock), the reference is exact and introduces no scoring uncertainty. For SCM-based datasets, the reference is obtained by dense grid search or large-sample uniform random search over the admissible intervention domain. In these cases, y^* represents the *best known achievable outcome*, not necessarily the true global optimum. This distinction is important for datasets with complex, potentially multimodal objective landscapes: if the reference underestimates the true optimum, both GAP and PA-GAP will overestimate the improvement ratios R_t , potentially making all methods appear closer to optimal than they actually are. In the opposite direction, an overestimated reference (which cannot occur by construction in our protocol, since we use oracle SCM access) would inflate the denominator and deflate all scores. To mitigate this risk, we use high-budget offline search (10,000 samples for Ecology and Epidemiology, 200×200 grids for Protein-reconstructed, and 1,000-point grids for Synthetic-2) and cross-validate against known interventional datasets where available.

Impact on metric comparisons. Because both GAP and PA-GAP normalize the improvement relative to $y^* - y(\mathbf{x}_{\text{init}})$, the quality of y^* affects the absolute metric values, but does not affect *the relative rankings*

between methods in the same dataset, provided that all methods are scored against the same reference. Cross data set comparisons of absolute GAP or PA-GAP values should therefore be made with caution, as differences in reference quality across datasets can introduce systematic biases.

Table 10: Reference optima used for GAP and PA-GAP scoring. Analytic optima are exact; grid/random-search references represent the best known achievable value.

Dataset	Task	y^*	Offline computation protocol
Dropwave	Max	1.0000	Analytic optimum from the known function generator.
Alpine2	Max	400.0000	Analytic optimum from the known function generator.
Ackley	Max	0.0000	Analytic optimum of the negated Ackley objective.
Rosenbrock	Max	0.0000	Analytic optimum of the negated Rosenbrock objective.
Chain-soft	Min	-3.3936	Reference value from the Chain-soft function generator.
Protein-recon.	Min	-240.0234	Reconstructed linear SEM; 200×200 grid over PKC and PKA; minimum Erk.
Synthetic-2	Min	-3.9690	Best offline oracle outcome; 1000-point grids on X and Z .
Synthetic	Min	-4.4474	Best value across BO and CBO oracle interventional datasets.
ToyGraph	Min	-2.5718	Best value across BO and CBO oracle interventional datasets.
Chain-hard	Min	-15.8060	Best value across BO and CBO oracle interventional datasets.
Ecology	Max	9.2652	10,000-sample uniform search over the manipulative-variable domain.
Epidemiology	Min	-3.3906	10,000-sample uniform search over the manipulative-variable domain.
Healthcare	Min	4.0401	Reference value from the distributed real-data benchmark artifact.

A.4 Consolidated dataset summary

Table 11 records benchmark provenance, including whether a task is inherited, reconstructed, or newly added, and Table 12 provides the code-authoritative structural and domain summary. Together these tables clarify that the benchmark is a standardized harness over heterogeneous prior tasks rather than a claim that every dataset is newly introduced here.

Table 12 provides a unified overview of all 13 benchmark datasets, consolidating key structural and experimental properties in a single reference. This table complements the SCM specifications per-dataset in Appendices A.5–A.7 and the usage-coverage table (Table 3) in the main text. The column “#Nodes” counts endogenous (observed) variables only; latent variables are listed separately where applicable. The column “Intervention domain” specifies the admissible range per manipulable variable as used in the benchmark.

Table 11: Benchmark provenance and domain reconciliation. “Current domain” refers to the code-authoritative domain used in our benchmark.

Dataset	Source / provenance	Current domain note	Modification status
ToyGraph	Aglietti et al. (2020)	$X \in [-5, 5]$, $Z \in [-5, 20]$	Inherited and standardized
Synthetic	CBO-style synthetic SCM	B, D, E code domains	Newly added SCM
Synthetic-2	cCBO-style synthetic SCM	$X \in [-3, 2]$, $Z \in [-1, 1]$	Newly added / reconstructed SCM
Chain-hard	Chain benchmark from fCBO-style SCM	$W, Z \in [-1, 1]$	Hard-intervention variant
Ecology	Aglietti et al. (2020) / ecological SCM	Code feature-parameter ranges	Inherited and repackaged
Protein-reconstructed	Sachs protein data / cCBO graph	Sachs-derived intervention ranges	Reconstructed fitted SCM
Healthcare	Aglietti et al. (2020) PSA example	Aspirin, Statin in $[0, 1]$	Inherited and repackaged
Epidemiology	CEO/DCBO epidemiology SCM	$L \in [0.479, 801.787]$, $B \in [-0.994, 0.999]$	Reconciled to code-authoritative intervention variables
Ackley / Rosenbrock / Dropwave / Alpine2	Function-network BO (Astudillo & Frazier, 2021); used by MCBO/ACBO	Standard function-network domains after wrapper normalization	Inherited function-network tasks
Chain-soft	Gultchin et al. (2023)	Policy coefficients in $[-0.27, 0.27]$, $W \in [-1, 1]$	Inherited policy-intervention task

Table 12: Consolidated summary of all thirteen benchmark datasets. For each dataset, we list the category, number of observed nodes, manipulable variables, intervention type, admissible domain per variable, optimization direction, and source reference. Latent variables (if any) are noted in parentheses.

Dataset	Category	#Nodes	Manipulable	Type	Domain	Task	Source
ToyGraph	Synthetic (hard)	3	X, Z	Hard	$X \in [-5, 5]; Z \in [-5, 20]$	Min Y	Aglietti et al. (2020)
Synthetic	Synthetic (hard)	7 (+2 latent)	B, D, E	Hard	$B \in [-5, 4]; D \in [-5, 5]; E \in [-6, 3]$	Min Y	Aglietti et al. (2020)
Synthetic-2	Synthetic (hard)	3	X, Z	Hard	$X \in [-3, 2]; Z \in [-1, 1]$	Min Y	Aglietti et al. (2023)
Chain-hard	Synthetic (hard)	4	W, Z	Hard	$[-1, 1]$ each	Min Y	Gultchin et al. (2023)
Ecology	Real/fitted (hard)	11	N, O, C, T, D	Hard	$N \in [-2, 5]; O \in [2, 4]; C \in [0, 1]; T \in [2200, 2500]; D \in [1950, 2100]$	Max Y	Aglietti et al. (2020)
Protein-recon.	Real/fitted (hard)	8	PKC, PKA, Mek, Akt	Hard	PKC $\in [0.5, 106.5];$ PKA $\in [1.45, 4491.5];$ Mek $\in [0.5, 389.5];$ Akt $\in [1.2, 3555.5]$	Min Erk	Aglietti et al. (2023)
Healthcare	Real/fitted (hard)	6	Aspirin, Statin	Hard	$[0, 1]$ each	Min PSA	Aglietti et al. (2020)
Epidemiology	Real/fitted (hard)	5	L, B	Hard	$L \in [0.479, 801.787]; B \in [-0.994, 0.999]$	Min Y	Branchini et al. (2023)
Ackley	Synthetic (soft)	9	$a_0, \dots, a_5$	Soft (function network)	$[-2, 2]$ each	Max Y	Astudillo & Frazier (2021); Sussex et al. (2023)
Rosenbrock	Synthetic (soft)	$D+D-1$	$a_0, \dots, a_{D-1}$	Soft (function network)	$[-2, 2]$ each	Max Y	Astudillo & Frazier (2021); Sussex et al. (2023)
Dropwave	Synthetic (soft)	4	a_0, a_1	Soft (function network)	$[-5.12, 5.12]$ each	Max Y	Astudillo & Frazier (2021); Sussex et al. (2023)
Alpine2	Synthetic (soft)	12	$a_0, \dots, a_5$	Soft (function network)	$[0, 10]$ each	Max Y	Astudillo & Frazier (2021); Sussex et al. (2023)
Chain-soft	Synthetic (soft)	4	W (hard), Z (policy)	Mixed hard/soft	$W \in [-1, 1]; Z = \pi(X)$ with coeffs. $\in [-0.27, 0.27]$	Min Y	Gultchin et al. (2023)

A.5 Synthetic dataset SCMs

This section provides the full structural equations, noise distributions, intervention domains, and causal graph diagrams for the four synthetic hard-intervention datasets used in our benchmark. For each data set, we follow a standardized format: (i) a brief description including the source reference, the number of nodes, and the direction of the task; (ii) the exogenous variables and their distributions; (iii) the endogenous structural equations; (iv) the noise model with explicit distributional parameters; (v) the intervention domain; and (vi) a figure of the causal graph with a consistent color scheme (green = manipulable, red = target, light-gray = non-manipulable, dark-gray = latent).

A.5.1 ToyGraph

A minimal three-node chain-structured SCM introduced by Aglietti et al. (2020), consisting of variables X , Z , and Y . Both X and Z are manipulable (intervenable); Y is the target. The task is to minimize Y . Figure 7 shows the causal graph.

Exogenous variable.

$$X \sim P_X \quad (\text{input variable, exogenous})$$

Endogenous variables.

$$Z = e^{-X} + \varepsilon_Z,$$

$$Y = \cos Z - e^{-Z/20} + \varepsilon_Y.$$

Noise model. All noise terms $\varepsilon_Z, \varepsilon_Y$ are mutually independent. In the original CBO implementation, these are drawn from $\mathcal{N}(0, \sigma^2)$ with a small variance σ^2 .

Intervention domain. Hard interventions on X and Z use dataset-specific domains:

$$X \in [-5, 5], \quad Z \in [-5, 20].$$

The admissible scopes are $\{X\}$, $\{Z\}$, and $\{X, Z\}$.




Figure 7: ToyGraph SCM. Green nodes represent manipulable variables; the red node is the target variable Y .

A.5.2 Synthetic

A seven-node SCM with two latent variables, introduced by Aglietti et al. (2020). The observed nodes are A, B, C, D, E, F, Y , with latent confounders U_1 and U_2 . Among the observed nodes, B, D , and E are manipulable; A, C , and F are not manipulable; and Y is the target. The task is to minimize Y . Figure 8 shows the causal graph.

Latent variables.

$$U_1 = \varepsilon_{YA} \sim \mathcal{N}(0, 1),$$

$$U_2 = \varepsilon_{YB} \sim \mathcal{N}(0, 1).$$

Exogenous variable.

$$F = \varepsilon_F \sim \mathcal{N}(0, 1).$$

Endogenous variables.

$$A = F^2 + U_1 + \varepsilon_A,$$

$$B = U_2 + \varepsilon_B,$$

$$C = e^{-B} + \varepsilon_C,$$

$$D = \frac{e^{-C}}{10} + \varepsilon_D,$$

$$E = \cos A + \frac{C}{10} + \varepsilon_E,$$

$$Y = \cos D + \sin E + U_1 + U_2 \varepsilon_Y.$$

Noise model. Each $\varepsilon_\bullet$ is an independent noise term. In the original CBO implementation, the endogenous noise terms are drawn from $\mathcal{N}(0, \sigma^2)$ with small variance σ^2 .

Intervention domain. Hard interventions on B , D , and E use dataset-specific domains:

$$B \in [-5, 4], \quad D \in [-5, 5], \quad E \in [-6, 3].$$

The admissible scopes include all non-empty subsets of $\{B, D, E\}$ considered by the benchmark intervention-set construction.




Figure 8: Synthetic SCM. Light-gray nodes are non-manipulable variables; dark-gray nodes are latent confounders; green nodes are manipulable variables; the red node is the target variable Y .

A.5.3 Synthetic-2

A three-node SCM introduced in the cCBO line of work (Aglietti et al., 2023), consisting of variables X , Z , and Y . Both X and Z are manipulable; Y is the target. The task is to minimize Y . The structure is similar to ToyGraph but uses explicitly specified Gaussian noise. Figure 9 shows the causal graph.

Exogenous variable.

$$X = \varepsilon_X, \quad \varepsilon_X \sim \mathcal{N}(0, 1).$$

Endogenous variables.

$$Z = e^{-X} + \varepsilon_Z,$$

$$Y = \cos Z - e^{-Z/20} + \varepsilon_Y.$$

Noise model. $\varepsilon_X \sim \mathcal{N}(0, 1)$, $\varepsilon_Z \sim \mathcal{N}(0, 1)$, and $\varepsilon_Y \sim \mathcal{N}(0, 1)$. All noise terms are mutually independent.

Intervention domain. Hard interventions on X and Z use $X \in [-3, 2]$ and $Z \in [-1, 1]$. The admissible scopes are $\{X\}$, $\{Z\}$, and $\{X, Z\}$.




Figure 9: Synthetic-2 SCM. Green nodes represent manipulable variables; the red node is the target variable Y .

A.5.4 Chain-hard

A four-node SCM introduced by Gultchin et al. (2023) as the hard-intervention version of the Chain benchmark. The observed variables are X , W , Z , and Y . Among them, W and Z are manipulable; X is a non-manipulable context variable; and Y is the target. The task is to minimize Y . Figure 10 illustrates the causal graph.

Structural equations.

$$\begin{aligned}
 X &= U_X, \\
 W &= U_W, \\
 Z &= -0.5X + U_Z, \\
 Y &= -W - 3ZX + U_Y.
 \end{aligned}$$

Noise model. $U_X, U_W, U_Z, U_Y \sim \mathcal{N}(0, 1)$, mutually independent.

Intervention domain. Hard interventions in W and Z are each restricted to $[-1, 1]$. A hard intervention replaces the corresponding structural equation with a fixed constant. The admissible scopes are $\{W\}$, $\{Z\}$, and $\{W, Z\}$.




Figure 10: Chain-hard SCM. Light-gray nodes are non-manipulable variables; green nodes are manipulable variables; the red node is the target variable Y .

A.6 Real dataset SCMs

This section provides the full structural equations, noise distributions, fitting details, and intervention domains for the four real or fitted-SCM hard-intervention datasets. Unlike the synthetic datasets above, these

SCMs are either fitted to real observational data or derived from domain-specific causal models. We follow the same standardized format as in the Appendix A.5. For protein reconstruction, we additionally document the fitting procedure to ensure transparent reproducibility.

A.6.1 Ecology

An 11-node SCM based on the Bermuda reef calcification model of Courtney et al. (2017), introduced in CBO (Aglietti et al., 2020). The manipulable variables are Nut, Chl α , TA, DIC, and Ω_A . The non-manipulable variables are Tem, Sal, P_{co_2} , Light, and pHsw. The target is NEC (net ecosystem calcification). The task is to *maximize* NEC. Figure 11 shows the causal graph.

Exogenous variables.

$$\begin{aligned} \text{Tem} &= U_{\text{Tem}} \sim \mathcal{N}(24.184130, 3.220405^2), \\ \text{Sal} &= U_{\text{Sal}} \sim \mathcal{N}(36.591624, 0.149197^2), \\ \text{Nut} &= U_{\text{Nut}} \sim \mathcal{N}(0.492065, 1.592408^2), \\ \text{TA} &= U_{\text{TA}} \sim \mathcal{N}(2357.893696, 27.609355^2). \end{aligned}$$

Endogenous variables.

$$\begin{aligned} P_{co_2} &= 18.798174 + 15.797384 \text{Tem} + U_{P_{co_2}}, \\ \text{Chl}\alpha &= 0.373420 - 0.002400 \text{Nut} + U_{\text{Chl}\alpha}, \\ \text{Light} &= 6665.081996 - 10737.462582 \text{Chl}\alpha + U_{\text{Light}}, \\ \text{pHsw} &= 8.427706 - 0.000966 P_{co_2} + U_{\text{pHsw}}, \\ \text{DIC} &= 2131.672107 - 0.216560 P_{co_2} + U_{\text{DIC}}, \\ \Omega_A &= 3.245248 + 0.094332 \text{Tem} + 0.006754 \text{Sal} - 0.005737 P_{co_2} + U_{\Omega_A}, \\ \text{NEC} &= 211.422555 - 0.000030 \text{Light} + 0.016680 \text{Nut} - 25.719277 \text{pHsw} - 0.500403 \Omega_A + U_{\text{NEC}}. \end{aligned}$$

Noise model. $U_{P_{co_2}} \sim \mathcal{N}(0, 28.896639^2)$, $U_{\text{Chl}\alpha} \sim \mathcal{N}(0, 0.039295^2)$, $U_{\text{Light}} \sim \mathcal{N}(0, 1546.913034^2)$, $U_{\text{pHsw}} \sim \mathcal{N}(0, 0.005258^2)$, $U_{\text{DIC}} \sim \mathcal{N}(0, 18.412100^2)$, $U_{\Omega_A} \sim \mathcal{N}(0, 0.036958^2)$, and $U_{\text{NEC}} \sim \mathcal{N}(0, 1.073340^2)$. All noise terms are mutually independent.

Intervention domain. Hard interventions in Nut, Chl α , TA, DIC, and Ω_A are applied in domain-specific ranges corresponding to the natural variability reported in Courtney et al. (2017). The benchmark uses the intervention domains specified in the public CBO repository.

A.6.2 Protein-reconstructed

An eight-node SCM reconstructed from protein-signaling data from Sachs et al. (2005), following the DAG used in cCBO (Aglietti et al., 2023). The manipulable variables are PKC, PKA, Mek, and Akt. The non-manipulable variables are Raf, P38, and Jnk. The target is Erk (extracellular signal-regulated kinase). The task is to minimize Erk. Figure 12 shows the causal graph.

Fitting procedure. This is a transparent reconstruction, not an exact reproduction of the unreleased fitted SCM from the original cCBO experiments. We retain the connected subgraph used in cCBO and fit *linear* structural mechanisms using ordinary least squares on 852 observational samples from Sachs et al. (2005). The parent sets are determined by the cCBO DAG. The linear model class was chosen for reproducibility and interpretability; nonlinear alternatives (e.g. kernel regression) were not explored.




Figure 11: Ecology SCM. Light-gray nodes are non-manipulable variables; green nodes are manipulable variables; the red node is the target variable NEC.

Exogenous variable.

$$\text{PKC} = U_{\text{PKC}} \sim \text{Uniform}(1, 106).$$

Endogenous variables.

$$\text{PKA} = 554.390731 + 0.841153 \text{PKC} + U_{\text{PKA}},$$

$$\text{Raf} = 62.199046 - 0.177745 \text{PKC} - 0.000379 \text{PKA} + U_{\text{Raf}},$$

$$\text{Mek} = -1.090275 + 0.039662 \text{PKC} - 0.000652 \text{PKA} + 0.520845 \text{Raf} + U_{\text{Mek}},$$

$$\text{P38} = 15.144328 + 1.234783 \text{PKC} + 0.000591 \text{PKA} + U_{\text{P38}},$$

$$\text{Jnk} = 52.953603 - 0.764801 \text{PKC} - 0.005306 \text{PKA} + U_{\text{Jnk}},$$

$$\text{Akt} = -31.110747 + 0.128905 \text{PKA} + U_{\text{Akt}},$$

$$\text{Erk} = -23.248743 + 0.081707 \text{PKA} - 0.029886 \text{Mek} + U_{\text{Erk}}.$$

Noise model. $U_{\text{PKA}} \sim \mathcal{N}(0, 427.437996^2)$, $U_{\text{Raf}} \sim \mathcal{N}(0, 41.768782^2)$, $U_{\text{Mek}} \sim \mathcal{N}(0, 16.695865^2)$, $U_{\text{P38}} \sim \mathcal{N}(0, 13.111164^2)$, $U_{\text{Jnk}} \sim \mathcal{N}(0, 42.074715^2)$, $U_{\text{Akt}} \sim \mathcal{N}(0, 113.958504^2)$, and $U_{\text{Erk}} \sim \mathcal{N}(0, 82.760198^2)$. All noise terms are mutually independent.

Intervention domain and constraints. Hard interventions in PKC and PKA are drawn from the range $[0.5, 106.5]$ and $[1.45, 4491.5]$, respectively. The interventions in Mek and Akt use the domains observed in the Sachs data, which are $[0.5, 389.5]$ and $[1.2, 3555.5]$. Biological plausibility constraints on PKC and PKA are enforced by clamping the proposed values to their observed ranges.

A.6.3 Healthcare

A six-node SCM derived from a real-world clinical setting (Thompson, 2019), involving Age, BMI, Aspirin, Statin, Cancer, and PSA. The manipulable variables are Aspirin and Statin. Non-manipulable variables are Age, BMI, and Cancer. The target is PSA (prostate-specific antigen). The task is to minimize PSA. Figure 13 shows the causal graph.

Exogenous variable.

$$\text{Age} \sim \text{Uniform}(55, 75).$$




Figure 12: Protein-signaling SCM (reconstructed). Light-gray nodes are non-manipulable variables; green nodes are manipulable variables; the red node is the target variable Erk.

Endogenous variables.

$$\text{BMI} = 27.0 - 0.01 \text{ Age} + \varepsilon_{\text{BMI}},$$

$$\text{Aspirin} = \sigma(-8.0 + 0.10 \text{ Age} + 0.03 \text{ BMI}) + \varepsilon_{\text{Aspirin}},$$

$$\text{Statin} = \sigma(-13.0 + 0.10 \text{ Age} + 0.20 \text{ BMI}) + \varepsilon_{\text{Statin}},$$

$$\text{Cancer} = \sigma(2.2 - 0.05 \text{ Age} + 0.01 \text{ BMI} - 0.04 \text{ Statin} + 0.02 \text{ Aspirin}) + \varepsilon_{\text{Cancer}},$$

$$\text{PSA} = 6.8 + 0.04 \text{ Age} - 0.15 \text{ BMI} - 0.60 \text{ Statin} + 0.55 \text{ Aspirin} + 1.00 \text{ Cancer} + \varepsilon_{\text{PSA}},$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$ denotes the sigmoid function.

Noise model. $\varepsilon_{\text{BMI}} \sim \mathcal{N}(0, 0.7^2)$, $\varepsilon_{\text{PSA}} \sim \mathcal{N}(0, 0.4^2)$. Each remaining $\varepsilon_{\bullet}$ is an independent noise term with small variance.

Intervention domain. Although the original data template specifies Aspirin and Statin as binary variables, existing CBO implementations commonly relax them to continuous values in $(0, 1)$ for BO-style optimization. We follow this convention to remain comparable to the prior work. The admissible scopes are $\{\text{Aspirin}\}$, $\{\text{Statin}\}$, and $\{\text{Aspirin}, \text{Statin}\}$.




Figure 13: Healthcare SCM. Light-gray nodes are non-manipulable variables; green nodes are manipulable variables; the red node is the target variable PSA.

A.6.4 Epidemiology

A five-node SCM introduced in CEO (Branchini et al., 2023), modeling an HIV treatment scenario based on Havercroft & Didelez (2012). The observed variables are B , T , L , R , and Y . In the code-authoritative

benchmark configuration, the manipulable variables are L and B ; T and R are non-intervention SCM variables. The target is Y (viral load). The task is to minimize Y . Figure 14 shows the causal graph.

Exogenous variables.

$$B \sim \text{Uniform}(-1, 1),$$

$$T \sim \text{Uniform}(4, 8).$$

Endogenous variables.

$$L = \exp(0.5T + U), \quad U \sim \mathcal{N}(0, 1),$$

$$R = 4 + LT,$$

$$Y = 0.5 + \cos(4T) + \sin(-L + 2R) + B + \varepsilon, \quad \varepsilon \sim \mathcal{N}(0, 1).$$

Noise model. $U \sim \mathcal{N}(0, 1)$ and $\varepsilon \sim \mathcal{N}(0, 1)$. All noise terms are mutually independent.

Intervention domain. Hard interventions in L are drawn from $[0.479, 801.787]$ and in B from $[-0.994, 0.999]$. The admissible scopes are $\{L\}$, $\{B\}$, and $\{L, B\}$. Since CEO uses a modified HIV SCM, more specific clinical interpretations of B and L are not provided in the original article.




Figure 14: Epidemiology SCM. Light-gray nodes are non-manipulable variables; green nodes are manipulable variables; the red node is the target variable Y (viral load).

A.7 Soft-intervention datasets

This section details the five soft-intervention datasets used in the benchmark. These datasets fall into two categories:

1. **Function network benchmarks** (Ackley, Rosenbrock, Dropwave, Alpine2): Classical global optimization benchmarks reformulated as deterministic function networks, where each action variable propagates through intermediate nodes to the target. Interventions correspond to setting action variables, and the network structure defines soft causal relationships. These are *maximization* tasks.
2. **SCM-based policy benchmark** (Chain-soft): A full structural causal model in which a soft intervention replaces the structural equation of one variable with a context-dependent policy. This is a *minimization* task.

For each data set, we specify the network architecture, all structural equations, the action domains, and the target computation. Figures use a consistent color scheme: blue = intermediate variables, orange squares = action (intervention) variables, red = target.




Figure 15: Ackley function network ($D = 6$). Blue nodes indicate intermediate variables; orange squares represent action (intervention) variables; the red node is the target variable Y .

A.7.1 Ackley

A multimodal, non-separable function-network benchmark with six action variables $a_0, \dots, a_5$, two intermediate nodes X_0 and X_1 , and target node Y . The task is to *maximize* Y . Following the standard configuration in the BO function-network (Sussex et al., 2023; 2024), we use the domain $a_i \in [-2, 2]$ for $i = 0, \dots, 5$, corresponding to $D = 6$. Figure 15 shows the function network.

Action variables. $a_i \in [-2, 2]$, $i = 0, \dots, 5$.

Intermediate variables.

$$X_0 = \frac{1}{D} \sum_{i=0}^5 a_i^2,$$

$$X_1 = \frac{1}{D} \sum_{i=0}^5 \cos(2\pi a_i),$$

where $D = 6$.

Target variable.

$$Y = 20 \exp(-0.2\sqrt{X_0}) + \exp(X_1) - 20 - e.$$

Interpretation. X_0 aggregates the average squared magnitude of the action variables, while X_1 aggregates the average cosine term. The target Y combines these two intermediate quantities to recover the Ackley objective in the form of maximum reward. The observational (unintervened) setting corresponds to $a_i = 0$ for all i .

A.7.2 Rosenbrock

A unimodal, non-separable function-network benchmark with a narrow curved valley. The network has action variables $a_0, \dots, a_{D-1}$, intermediate nodes $X_0, \dots, X_{D-2}$, and target node Y . The task is to *maximize* Y . We use the domain $a_i \in [-2, 2]$ and consider $D \in \{3, 5, 7\}$. Figure 16 shows the chain-structured function network.

Action variables. $a_i \in [-2, 2]$, $i = 0, \dots, D - 1$.




Figure 16: Rosenbrock function network (shown for $D = 3$). Blue nodes indicate intermediate variables; orange squares represent action (intervention) variables; the red node is the target variable Y .

Intermediate variables. The first intermediate node computes the first Rosenbrock term:

$$X_0 = -100(a_1 - a_0^2)^2 - (1 - a_0)^2.$$

Each subsequent node recursively accumulates the next pairwise contribution:

$$X_k = -100(a_{k+1} - a_k^2)^2 - (1 - a_k)^2 + X_{k-1}, \quad k = 1, \dots, D - 2.$$

Target variable.

$$Y = X_{D-2}.$$

Interpretation. The Rosenbrock function network forms a causal chain in which the first node computes the initial Rosenbrock component, and each later node adds one additional pairwise contribution to the accumulated objective. The target Y recovers the Rosenbrock objective in the form of maximum reward. The observational setting corresponds to $a_i = 0$ for all i .

A.7.3 Dropwave




Figure 17: Dropwave function network. Blue nodes indicate intermediate variables; orange squares represent action (intervention) variables; the red node is the target variable Y .

A highly multimodal two-dimensional function-network benchmark with two action variables a_0 and a_1 , an intermediate node X_0 and a target node Y . The task is to *maximize* Y . Following the BO setup of the standard function-network (Sussex et al., 2023; 2024), we use the domain $a_0, a_1 \in [-5.12, 5.12]$. Figure 17 shows the function network.

Action variables. $a_0, a_1 \in [-5.12, 5.12]$.

Intermediate variable.

$$X_0 = \sqrt{a_0^2 + a_1^2} + \epsilon_0.$$

Target variable.

$$Y = \frac{1 + \cos(12X_0)}{2 + 0.5X_0^2} + \epsilon_Y.$$

Interpretation. The intermediate node X_0 computes the radial distance from the origin induced by the two action variables. The target Y applies the oscillatory Dropwave transformation to this radial value. Following MCBO and related function-network BO settings, this benchmark can also be evaluated in a noisy setting by choosing non-zero ϵ_0 and ϵ_Y .

A.7.4 Alpine2




Figure 18: Alpine2 function network ($K = 6$). Blue nodes indicate intermediate variables; orange squares represent action (intervention) variables; the red node is the target variable Y .

A multimodal, separable chain-structured function-network benchmark with action variables $a_0, \dots, a_5$, intermediate nodes $X_0, \dots, X_4$, and target node Y . The task is to *maximize* Y . We use the setting $K = 6$ with the action domain $a_i \in [0, 10]$ for $i = 0, \dots, 5$. Figure 18 shows the function network.

Action variables. $a_i \in [0, 10]$, $i = 0, \dots, 5$.

Intermediate variables. The first intermediate node is defined as

$$X_0 = -\sqrt{a_0} \sin(a_0).$$

Each subsequent node combines its own action variable with the output of the previous node through the Alpine2 multiplicative structure:

$$X_k = \sqrt{a_k} \sin(a_k) X_{k-1}, \quad k = 1, \dots, 4.$$

Target variable.

$$Y = \sqrt{a_5} \sin(a_5) X_4.$$

Interpretation. Alpine2 forms a causal chain in which each stage combines its own decision variable with the output of the previous stage through a multiplicative structure. The analytical global maximum is $\prod_{i=0}^5 \sqrt{a_i^*} \sin(a_i^*) = 2.808^6 \approx 400$ at $a_i^* \approx 7.917$ for all i . This subsection corresponds to the $K = 6$ configuration used in our experiments.

A.7.5 Chain-soft

The soft-intervention Chain benchmark uses the same base SCM as Chain-hard (Appendix A.5.4), with observed nodes X, W, Z, Y . Here, X is a non-manipulable context variable, W is hard-intervenable, Z is *soft*-intervenable via a context-dependent policy, and Y is the target. The task is to *minimize* Y . Figure 19 illustrates the SCM with the intervention policy.




Figure 19: Chain SCM under soft intervention. Light-gray nodes are non-manipulable variables; green nodes indicate hard-intervenable variables; blue nodes indicate intermediate variables; the orange rectangle denotes the intervention policy π_0 ; the red node is the target variable Y .

Exogenous variables.

$$X = U_X \sim \mathcal{N}(0, 1),$$

$$W = U_W \sim \mathcal{N}(0, 1).$$

Endogenous variables (unintervened).

$$Z = -0.5X + U_Z,$$

$$Y = -W - 3ZX + U_Y,$$

where $U_Z \sim \mathcal{N}(0, 1)$ and $U_Y \sim \mathcal{N}(0, 1)$, all mutually independent.

Intervention policy. Under soft intervention, the structural mechanism of Z is replaced by a context-dependent policy π_0 :

$$Z := \pi_0(X),$$

so that the intervened SCM becomes:

$$X = U_X,$$

$$W = U_W \quad \text{or} \quad W := w, \quad w \in [-1, 1],$$

$$Z = \pi_0(X),$$

$$Y = -W - 3ZX + U_Y.$$

Intervention domain and policy parameterization. The admissible scopes include a purely functional intervention in Z (through the policy π_0) and a mixed intervention combining the functional intervention in Z with a hard intervention in W over $[-1, 1]$. Following Gultchin et al. (2023), each functional intervention is represented using samples `GridSize` = 10 and $N_\alpha = N_\beta = 10$ for the context variable, with coefficients α_i and β_j uniformly sampled from $[-0.27, 0.27]$. This parameterization keeps the induced intervention values comparable to the hard-intervention range $[-1, 1]$.

Interpretation. Because the target $Y = -W - 3ZX + U_Y$ depends on the interaction between Z and the context X , the optimal intervention in Z generally varies with X . This makes Chain-soft a key benchmark for evaluating context-adaptive intervention strategies. In the figure, the dashed arrows $X \dashrightarrow \pi_0 \dashrightarrow Z$ indicate that the policy takes X as input and determines the intervened value of Z .




Figure 20: Perturbation taxonomy on the ToyGraph SCM ($X \rightarrow Z \rightarrow Y$). From left to right: the true benchmark graph, edge deletion ($X \rightarrow Z$ removed), edge addition (a spurious $X \rightarrow Y$ inserted), and edge reversal ($X \rightarrow Z$ flipped to $Z \rightarrow X$). Green nodes are manipulable and the red node is the target Y ; kept edges are solid black, the spurious added edge is solid blue, and the deleted edge is dashed red. Interventional outcomes and the reference optimum y^* are always computed from the true SCM (leftmost panel); only the graph handed to the graph-using optimizers is perturbed.

B Additional Results

B.1 Robustness Stress Tests

B.1.1 Graph-misspecification stress test

To go beyond a limitation statement, we added a graph-misspecification stress test. In this experiment, the true environment is unchanged: interventional outcomes and reference optima are always computed from the true benchmark SCM. Only the graph supplied to graph-using optimizers is perturbed. Thus, the experiment tests sensitivity to incorrect causal information without changing the true intervention task. We consider three perturbation types: edge addition, edge deletion, and edge reversal. The results are shown in Tables 14, 15, and 16. Each perturbation setting is evaluated over 20 random seeds, matching the seed count used in the main benchmark. We present these stress tests as robustness diagnostics under controlled graph misspecification, not as a complete deployment-robustness benchmark.

To make the perturbations concrete and reproducible, we visualize them directly on the benchmark graphs. Figure 20 illustrates the three perturbation types on the canonical ToyGraph chain $X \rightarrow Z \rightarrow Y$, and Figures 21 and 22 show, for every other benchmark graph, the specific edge that is added, deleted, or reversed and, where a hidden-confounder variant is defined, the node that is removed from the learner’s view. All panels share the visual encoding of Figure 20, with a hidden node drawn as an open dashed circle. In every panel exactly one edge (or, for the hidden-confounder variant, one node) is perturbed relative to the true SCM, so the schematics also document the minimal, single-perturbation nature of each stress test.

The stress test confirms the concern that motivated it: method rankings change substantially under graph perturbation. For example, under edge addition at 100 trials, the GAP winner varies across datasets: CoCaBO is best on ToyGraph, BO is best on Synthetic, DCBO is best on Synthetic-2, Ecology, and Epidemiology, CBO is best on Chain-hard and Healthcare, and MCBO is best on Protein-reconstructed (Table 14). Under edge deletion and edge reversal, the winners change again (Tables 15 and 16). PA-GAP also produces different rankings from GAP because it rewards full-trajectory progress rather than only final discovery time. These results support a more cautious conclusion: causal structure can help, but current CBO pipelines are sensitive to how that structure is specified and operationalized. This particular stress test addresses learner-side graph misspecification; hidden confounding is a separate dimension, which we probe with a companion omitted-variable experiment (Table 17, below). We treat comprehensive robustness to unobserved confounding as still open rather than resolved by these pilots.

Some misspecified-graph runs outperform the corresponding correct-graph run. We do not interpret this as evidence that the misspecified graph is causally preferable. Since outcomes and y^* are still computed from the true SCM, such improvements reflect finite-budget search effects. A perturbed graph can weaken an overconfident observational prior, change candidate intervention scopes, alter exploration, or accidentally

Table 13: Degradation summary for the graph-misspecification stress tests. Each cell compares a graph-using method’s mean score under the perturbed graph against the corresponding correct-graph mean score in the main benchmark tables, over 6 methods $\times$ 8 datasets $\times$ 2 metrics $\times$ 3 budgets ($N=288$ per perturbation type).

Perturbation	Degrade	Improve	Unchanged
Edge addition	150 (52%)	78 (27%)	60 (21%)
Edge deletion	153 (53%)	76 (26%)	59 (20%)
Edge reversal	156 (54%)	75 (26%)	57 (20%)

induce a more favorable acquisition trajectory. We therefore report these cases as evidence that CBO is sensitive to causal inductive bias, not as evidence that misspecification is beneficial.

To make the comparison interpretable, we compare each perturbed-graph score against the corresponding correct-graph score obtained with the same dataset, method, metric, and budget. Table 13 reports the resulting compact degradation summary: for each perturbation type, the fraction of graph-using method cells (6 methods $\times$ 8 datasets $\times$ 2 metrics $\times$ 3 budgets, $N=288$) whose mean score degrades, improves, or is unchanged relative to the correct-graph score in the main benchmark tables. Across all three perturbation types, roughly half of the cells degrade (52–54%), about a quarter improve (26–27%), and about a fifth are unchanged (20–21%); CBO is the most consistently affected pipeline (34–37 of 48 cells degrade per perturbation), while HCBO improves nearly as often as it degrades. Because the stress-test pilots were executed as separate runs (visible in the shared BO control, whose scores differ slightly from the main-table BO scores), these fractions should be read as coarse sensitivity indicators rather than exact paired effect sizes. This summary prevents individual improvements under misspecification from being misread as evidence that the perturbed graph is causally better: improvements are a minority pattern and are consistent with the finite-budget search effects described above.

For methods not directly affected by the supplied fixed graph in this pilot, the perturbation does not change the optimizer. The non-causal BO baseline never reads the causal graph, and the CEO runner used in this pilot does not consume a single fixed supplied graph in the same way as the known-graph methods. We therefore treat BO and CEO as isolation controls and report a single shared baseline for each dataset–budget–metric setting, identical across the edge-addition, edge-deletion, and edge-reversal tables. Thus, cross-table differences are confined to the graph-using methods (CBO, cCBO, DCBO, MCBO, HCBO, and CoCaBO).

Table 14: Graph-misspecification stress test under edge addition (mean $\pm$ standard error over 20 seeds; illustrative pilot stress test; higher is better; best mean among methods under the same perturbation setting is in bold).

Trial limit	Metric	Dataset	BO	CBO	cCBO	DCBO	MCBO	HCBO	CoCaBO	CEO
100	GAP	ToyGraph	0.232 $\pm$.079	0.288 $\pm$.071	0.567 $\pm$.001	0.569 $\pm$.055	0.558 $\pm$.286	0.509 $\pm$.049	0.682$\pm$.109	0.459 $\pm$.084
		Synthetic	0.651$\pm$.065	0.454 $\pm$.072	0.642 $\pm$.056	0.337 $\pm$.036	0.456 $\pm$.082	0.391 $\pm$.088	0.617 $\pm$.093	0.396 $\pm$.042
		Synthetic-2	0.001 $\pm$.004	0.281 $\pm$.125	0.289 $\pm$.128	0.578$\pm$.025	0.000 $\pm$.000	0.472 $\pm$.200	0.122 $\pm$.076	0.526 $\pm$.183
		Chain-hard	0.787 $\pm$.103	0.814$\pm$.128	0.618 $\pm$.151	0.333 $\pm$.333	0.284 $\pm$.041	0.095 $\pm$.062	0.786 $\pm$.107	0.347 $\pm$.046
		Ecology	0.160 $\pm$.128	0.018 $\pm$.018	0.426 $\pm$.068	0.519$\pm$.001	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.228 $\pm$.051
		Protein-reconstructed	0.331 $\pm$.125	0.363 $\pm$.139	0.000 $\pm$.000	0.000 $\pm$.000	0.615$\pm$.097	0.288 $\pm$.091	0.474 $\pm$.045	0.193 $\pm$.048
	PA-GAP	Healthcare	0.315 $\pm$.054	0.510$\pm$.035	0.509 $\pm$.083	0.341 $\pm$.093	0.357 $\pm$.029	0.234 $\pm$.054	0.257 $\pm$.129	0.483 $\pm$.086
		Epidemiology	0.506 $\pm$.056	0.563 $\pm$.038	0.000 $\pm$.000	0.599$\pm$.008	0.000 $\pm$.000	0.390 $\pm$.077	0.300 $\pm$.048	0.218 $\pm$.081
		ToyGraph	0.051 $\pm$.039	0.089 $\pm$.058	0.073 $\pm$.001	0.088 $\pm$.061	0.172 $\pm$.122	0.199 $\pm$.082	0.320$\pm$.032	0.203 $\pm$.099
		Synthetic	0.214 $\pm$.051	0.041 $\pm$.096	0.286 $\pm$.027	0.025 $\pm$.056	0.113 $\pm$.082	0.075 $\pm$.093	0.397$\pm$.033	0.253 $\pm$.054
		Synthetic-2	0.005 $\pm$.002	0.003 $\pm$.002	0.139 $\pm$.069	0.089 $\pm$.029	0.000 $\pm$.000	0.145 $\pm$.077	0.090 $\pm$.054	0.155$\pm$.057
		Chain-hard	0.497$\pm$.015	0.437 $\pm$.005	0.345 $\pm$.129	0.168 $\pm$.169	0.077 $\pm$.014	0.066 $\pm$.051	0.424 $\pm$.069	0.007 $\pm$.033
50	GAP	Ecology	0.031 $\pm$.011	0.002 $\pm$.002	0.174 $\pm$.038	0.178 $\pm$.001	0.000 $\pm$.000	0.210$\pm$.120	0.000 $\pm$.000	0.007 $\pm$.069
		Protein-reconstructed	0.153$\pm$.061	0.057 $\pm$.018	0.000 $\pm$.000	0.000 $\pm$.000	0.031 $\pm$.036	0.000 $\pm$.000	0.051 $\pm$.024	0.010 $\pm$.072
		Healthcare	0.120 $\pm$.021	0.077 $\pm$.015	0.202$\pm$.014	0.000 $\pm$.000	0.071 $\pm$.051	0.010 $\pm$.006	0.116 $\pm$.058	0.138 $\pm$.049
		Epidemiology	0.242 $\pm$.021	0.185 $\pm$.059	0.000 $\pm$.000	0.341$\pm$.002	0.000 $\pm$.000	0.138 $\pm$.080	0.122 $\pm$.078	0.001 $\pm$.042
	PA-GAP	ToyGraph	0.202 $\pm$.099	0.235 $\pm$.111	0.560 $\pm$.009	0.545 $\pm$.046	0.296 $\pm$.296	0.450 $\pm$.093	0.796$\pm$.041	0.489 $\pm$.073
		Synthetic	0.567 $\pm$.091	0.000 $\pm$.097	0.440 $\pm$.048	0.295 $\pm$.050	0.456 $\pm$.099	0.391 $\pm$.088	0.651$\pm$.028	0.254 $\pm$.057
		Synthetic-2	0.007 $\pm$.005	0.310 $\pm$.033	0.377 $\pm$.192	0.562$\pm$.017	0.000 $\pm$.000	0.428 $\pm$.217	0.122 $\pm$.076	0.484 $\pm$.222
		Chain-hard	0.686 $\pm$.086	0.889$\pm$.016	0.645 $\pm$.198	0.333 $\pm$.333	0.360 $\pm$.107	0.095 $\pm$.062	0.725 $\pm$.143	0.356 $\pm$.101
		Ecology	0.259 $\pm$.148	0.134 $\pm$.134	0.502$\pm$.105	0.357 $\pm$.001	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.101 $\pm$.031
		Protein-reconstructed	0.365 $\pm$.061	0.413 $\pm$.046	0.000 $\pm$.000	0.000 $\pm$.000	0.526 $\pm$.097	0.591$\pm$.040	0.392 $\pm$.120	0.264 $\pm$.028
20	GAP	Healthcare	0.439 $\pm$.121	0.432 $\pm$.073	0.604$\pm$.007	0.332 $\pm$.093	0.357 $\pm$.029	0.234 $\pm$.054	0.244 $\pm$.132	0.489 $\pm$.076
		Epidemiology	0.358 $\pm$.052	0.355 $\pm$.021	0.000 $\pm$.000	0.671$\pm$.005	0.000 $\pm$.000	0.390 $\pm$.077	0.246 $\pm$.079	0.060 $\pm$.050
	PA-GAP	ToyGraph	0.024 $\pm$.016	0.070 $\pm$.053	0.073 $\pm$.006	0.083 $\pm$.057	0.108 $\pm$.109	0.199 $\pm$.082	0.321$\pm$.034	0.213 $\pm$.099
		Synthetic	0.194 $\pm$.054	0.000 $\pm$.036	0.201 $\pm$.034	0.023 $\pm$.050	0.113 $\pm$.082	0.075 $\pm$.093	0.344$\pm$.043	0.246 $\pm$.077
		Synthetic-2	0.003 $\pm$.003	0.003 $\pm$.002	0.110 $\pm$.056	0.085 $\pm$.026	0.000 $\pm$.000	0.081 $\pm$.041	0.090 $\pm$.054	0.146$\pm$.056
		Chain-hard	0.481$\pm$.029	0.379 $\pm$.010	0.327 $\pm$.124	0.170 $\pm$.170	0.050 $\pm$.025	0.066 $\pm$.051	0.391 $\pm$.098	0.001 $\pm$.054
	PA-GAP	Ecology	0.017 $\pm$.010	0.002 $\pm$.002	0.166 $\pm$.040	0.175 $\pm$.001	0.000 $\pm$.000	0.208$\pm$.120	0.000 $\pm$.000	0.006 $\pm$.075
		Protein-reconstructed	0.137$\pm$.054	0.047 $\pm$.017	0.000 $\pm$.000	0.000 $\pm$.000	0.031 $\pm$.036	0.000 $\pm$.000	0.045 $\pm$.019	0.009 $\pm$.075
		Healthcare	0.105 $\pm$.017	0.069 $\pm$.014	0.178$\pm$.010	0.000 $\pm$.000	0.071 $\pm$.051	0.009 $\pm$.005	0.100 $\pm$.051	0.147 $\pm$.087
		Epidemiology	0.232 $\pm$.025	0.129 $\pm$.061	0.000 $\pm$.000	0.315$\pm$.003	0.000 $\pm$.000	0.127 $\pm$.073	0.085 $\pm$.069	0.008 $\pm$.002
	GAP	ToyGraph	0.272 $\pm$.097	0.167 $\pm$.064	0.538 $\pm$.000	0.472 $\pm$.034	0.239 $\pm$.239	0.450 $\pm$.093	0.763$\pm$.057	0.489 $\pm$.073
		Synthetic	0.390 $\pm$.138	0.000 $\pm$.000	0.475 $\pm$.057	0.295 $\pm$.050	0.456 $\pm$.099	0.391 $\pm$.088	0.590$\pm$.095	0.254 $\pm$.057
		Synthetic-2	0.003 $\pm$.002	0.185 $\pm$.125	0.515$\pm$.009	0.129 $\pm$.129	0.000 $\pm$.000	0.428 $\pm$.217	0.000 $\pm$.000	0.484 $\pm$.222
		Chain-hard	0.761$\pm$.124	0.718 $\pm$.039	0.486 $\pm$.119	0.333 $\pm$.333	0.380 $\pm$.129	0.000 $\pm$.000	0.731 $\pm$.228	0.356 $\pm$.101
		Ecology	0.145 $\pm$.145	0.080 $\pm$.080	0.422 $\pm$.159	0.442$\pm$.000	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.101 $\pm$.031
		Protein-reconstructed	0.370 $\pm$.073	0.194 $\pm$.091	0.000 $\pm$.000	0.000 $\pm$.000	0.526 $\pm$.097	0.556$\pm$.044	0.516 $\pm$.027	0.264 $\pm$.028
20	GAP	Healthcare	0.348 $\pm$.119	0.268 $\pm$.142	0.413 $\pm$.036	0.332 $\pm$.093	0.000 $\pm$.000	0.234 $\pm$.054	0.198 $\pm$.100	0.489$\pm$.076
		Epidemiology	0.694$\pm$.042	0.244 $\pm$.122	0.000 $\pm$.000	0.373 $\pm$.005	0.000 $\pm$.000	0.390 $\pm$.077	0.259 $\pm$.165	0.001 $\pm$.005
	PA-GAP	ToyGraph	0.006 $\pm$.004	0.042 $\pm$.040	0.070 $\pm$.000	0.069 $\pm$.047	0.044 $\pm$.044	0.199 $\pm$.082	0.324$\pm$.039	0.213 $\pm$.099
		Synthetic	0.156 $\pm$.059	0.000 $\pm$.000	0.154 $\pm$.036	0.023 $\pm$.050	0.113 $\pm$.082	0.000 $\pm$.000	0.244 $\pm$.054	0.246$\pm$.077
		Synthetic-2	0.004 $\pm$.001	0.057 $\pm$.029	0.075 $\pm$.017	0.001 $\pm$.001	0.000 $\pm$.000	0.081 $\pm$.041	0.000 $\pm$.000	0.146$\pm$.056
		Chain-hard	0.463$\pm$.062	0.228 $\pm$.019	0.275 $\pm$.112	0.175 $\pm$.175	0.036 $\pm$.022	0.000 $\pm$.000	0.378 $\pm$.108	0.001 $\pm$.054
	PA-GAP	Ecology	0.011 $\pm$.011	0.001 $\pm$.001	0.150 $\pm$.047	0.168 $\pm$.001	0.000 $\pm$.000	0.208$\pm$.120	0.000 $\pm$.000	0.006 $\pm$.075
		Protein-reconstructed	0.113 $\pm$.049	0.025 $\pm$.013	0.000 $\pm$.000	0.193$\pm$.097	0.031 $\pm$.036	0.000 $\pm$.000	0.045 $\pm$.020	0.009 $\pm$.075
		Healthcare	0.091 $\pm$.015	0.056 $\pm$.011	0.119 $\pm$.005	0.000 $\pm$.000	0.000 $\pm$.000	0.009 $\pm$.005	0.074 $\pm$.040	0.147$\pm$.087
		Epidemiology	0.228 $\pm$.028	0.087 $\pm$.045	0.000 $\pm$.000	0.279$\pm$.005	0.000 $\pm$.000	0.127 $\pm$.073	0.063 $\pm$.060	0.007 $\pm$.001

Table 15: Graph-misspecification stress test under edge deletion (mean $\pm$ standard error over 20 seeds; illustrative pilot stress test; higher is better; best mean among methods under the same perturbation setting is in bold).

Trial limit	Metric	Dataset	BO	CBO	cCBO	DCBO	MCBO	HCBO	CoCaBO	CEO
100	GAP	ToyGraph	0.232 $\pm$.079	0.392 $\pm$.009	0.826$\pm$.001	0.564 $\pm$.060	0.394 $\pm$.200	0.234 $\pm$.054	0.335 $\pm$.096	0.459 $\pm$.084
		Synthetic	0.651 $\pm$.065	0.355 $\pm$.184	0.642 $\pm$.056	0.337 $\pm$.036	0.456 $\pm$.082	0.391 $\pm$.088	0.663$\pm$.020	0.396 $\pm$.042
		Synthetic-2	0.001 $\pm$.004	0.281 $\pm$.125	0.306 $\pm$.145	0.000 $\pm$.004	0.000 $\pm$.000	0.657$\pm$.139	0.122 $\pm$.076	0.526 $\pm$.183
		Chain-hard	0.787$\pm$.103	0.700 $\pm$.128	0.767 $\pm$.060	0.448 $\pm$.300	0.428 $\pm$.085	0.095 $\pm$.062	0.688 $\pm$.136	0.347 $\pm$.046
		Ecology	0.160 $\pm$.128	0.012 $\pm$.012	0.426$\pm$.068	0.255 $\pm$.008	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.228 $\pm$.051
		Protein-reconstructed	0.331 $\pm$.125	0.363 $\pm$.139	0.000 $\pm$.000	0.000 $\pm$.000	0.615$\pm$.097	0.288 $\pm$.091	0.578 $\pm$.082	0.193 $\pm$.048
	PA-GAP	Healthcare	0.315 $\pm$.054	0.510$\pm$.035	0.445 $\pm$.100	0.341 $\pm$.093	0.357 $\pm$.029	0.234 $\pm$.054	0.431 $\pm$.216	0.483 $\pm$.086
		Epidemiology	0.506 $\pm$.056	0.563 $\pm$.038	0.000 $\pm$.000	0.661$\pm$.092	0.000 $\pm$.000	0.390 $\pm$.077	0.572 $\pm$.134	0.218 $\pm$.081
		ToyGraph	0.051 $\pm$.039	0.104 $\pm$.054	0.355$\pm$.001	0.095 $\pm$.056	0.248 $\pm$.129	0.199 $\pm$.082	0.103 $\pm$.051	0.203 $\pm$.099
		Synthetic	0.214 $\pm$.051	0.087 $\pm$.044	0.286 $\pm$.027	0.025 $\pm$.056	0.113 $\pm$.082	0.075 $\pm$.093	0.396$\pm$.039	0.253 $\pm$.054
		Synthetic-2	0.005 $\pm$.002	0.003 $\pm$.002	0.133 $\pm$.069	0.000 $\pm$.006	0.000 $\pm$.000	0.178$\pm$.131	0.090 $\pm$.054	0.155 $\pm$.057
		Chain-hard	0.497$\pm$.015	0.437 $\pm$.005	0.414 $\pm$.030	0.221 $\pm$.147	0.085 $\pm$.029	0.066 $\pm$.051	0.388 $\pm$.098	0.007 $\pm$.033
50	GAP	Ecology	0.031 $\pm$.011	0.000 $\pm$.005	0.174 $\pm$.038	0.041 $\pm$.008	0.000 $\pm$.000	0.210$\pm$.120	0.000 $\pm$.000	0.007 $\pm$.069
		Protein-reconstructed	0.153$\pm$.061	0.057 $\pm$.018	0.000 $\pm$.000	0.000 $\pm$.000	0.031 $\pm$.036	0.000 $\pm$.000	0.094 $\pm$.077	0.010 $\pm$.072
		Healthcare	0.120 $\pm$.021	0.077 $\pm$.015	0.158$\pm$.052	0.000 $\pm$.000	0.071 $\pm$.051	0.010 $\pm$.006	0.134 $\pm$.068	0.138 $\pm$.049
		Epidemiology	0.242 $\pm$.021	0.185 $\pm$.059	0.000 $\pm$.000	0.273$\pm$.038	0.000 $\pm$.000	0.138 $\pm$.080	0.227 $\pm$.087	0.001 $\pm$.042
	PA-GAP	ToyGraph	0.202 $\pm$.099	0.363 $\pm$.141	0.780$\pm$.009	0.525 $\pm$.062	0.569 $\pm$.288	0.450 $\pm$.093	0.195 $\pm$.059	0.489 $\pm$.073
		Synthetic	0.567 $\pm$.091	0.271 $\pm$.177	0.440 $\pm$.048	0.295 $\pm$.050	0.456 $\pm$.099	0.391 $\pm$.088	0.688$\pm$.082	0.254 $\pm$.057
		Synthetic-2	0.007 $\pm$.005	0.310 $\pm$.033	0.229 $\pm$.118	0.000 $\pm$.002	0.000 $\pm$.000	0.467 $\pm$.251	0.122 $\pm$.076	0.484$\pm$.222
		Chain-hard	0.686 $\pm$.086	0.889$\pm$.016	0.741 $\pm$.151	0.448 $\pm$.300	0.346 $\pm$.109	0.095 $\pm$.062	0.601 $\pm$.165	0.356 $\pm$.101
		Ecology	0.259 $\pm$.148	0.000 $\pm$.005	0.502$\pm$.105	0.158 $\pm$.008	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.101 $\pm$.031
		Protein-reconstructed	0.365 $\pm$.061	0.413 $\pm$.046	0.000 $\pm$.000	0.000 $\pm$.000	0.526 $\pm$.097	0.591$\pm$.040	0.562 $\pm$.088	0.264 $\pm$.028
20	GAP	Healthcare	0.439 $\pm$.121	0.432 $\pm$.073	0.516$\pm$.102	0.332 $\pm$.093	0.357 $\pm$.029	0.234 $\pm$.054	0.380 $\pm$.197	0.489 $\pm$.076
		Epidemiology	0.358 $\pm$.052	0.355 $\pm$.021	0.000 $\pm$.000	0.687$\pm$.032	0.000 $\pm$.000	0.390 $\pm$.077	0.482 $\pm$.140	0.060 $\pm$.050
	PA-GAP	ToyGraph	0.024 $\pm$.016	0.093 $\pm$.045	0.338$\pm$.006	0.088 $\pm$.054	0.180 $\pm$.110	0.199 $\pm$.082	0.074 $\pm$.036	0.213 $\pm$.099
		Synthetic	0.194 $\pm$.054	0.076 $\pm$.038	0.201 $\pm$.034	0.023 $\pm$.050	0.113 $\pm$.082	0.075 $\pm$.093	0.346$\pm$.055	0.246 $\pm$.077
		Synthetic-2	0.003 $\pm$.003	0.003 $\pm$.002	0.104 $\pm$.058	0.000 $\pm$.005	0.000 $\pm$.000	0.158$\pm$.137	0.090 $\pm$.054	0.146 $\pm$.056
		Chain-hard	0.481$\pm$.029	0.379 $\pm$.010	0.387 $\pm$.030	0.222 $\pm$.148	0.062 $\pm$.030	0.066 $\pm$.051	0.380 $\pm$.099	0.001 $\pm$.054
	PA-GAP	Ecology	0.017 $\pm$.010	0.000 $\pm$.002	0.166 $\pm$.040	0.031 $\pm$.010	0.000 $\pm$.000	0.210$\pm$.120	0.000 $\pm$.000	0.006 $\pm$.075
		Protein-reconstructed	0.137$\pm$.054	0.047 $\pm$.017	0.000 $\pm$.000	0.000 $\pm$.000	0.031 $\pm$.036	0.000 $\pm$.000	0.095 $\pm$.078	0.009 $\pm$.075
		Healthcare	0.105 $\pm$.017	0.069 $\pm$.014	0.141 $\pm$.044	0.000 $\pm$.000	0.071 $\pm$.051	0.009 $\pm$.005	0.122 $\pm$.061	0.147$\pm$.087
		Epidemiology	0.232 $\pm$.025	0.129 $\pm$.061	0.000 $\pm$.000	0.258$\pm$.033	0.000 $\pm$.000	0.127 $\pm$.073	0.196 $\pm$.085	0.008 $\pm$.002
20	GAP	ToyGraph	0.272 $\pm$.097	0.375 $\pm$.145	0.639$\pm$.000	0.405 $\pm$.074	0.274 $\pm$.274	0.450 $\pm$.093	0.097 $\pm$.039	0.489 $\pm$.073
		Synthetic	0.390 $\pm$.138	0.276 $\pm$.163	0.475 $\pm$.057	0.295 $\pm$.050	0.456 $\pm$.099	0.391 $\pm$.088	0.511$\pm$.058	0.254 $\pm$.057
		Synthetic-2	0.003 $\pm$.002	0.245 $\pm$.125	0.000 $\pm$.000	0.507$\pm$.026	0.000 $\pm$.000	0.467 $\pm$.251	0.000 $\pm$.000	0.484 $\pm$.222
		Chain-hard	0.761 $\pm$.124	0.718 $\pm$.039	0.778$\pm$.019	0.318 $\pm$.318	0.250 $\pm$.077	0.000 $\pm$.000	0.671 $\pm$.202	0.356 $\pm$.101
		Ecology	0.145 $\pm$.145	0.000 $\pm$.000	0.422$\pm$.159	0.221 $\pm$.147	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.101 $\pm$.031
		Protein-reconstructed	0.370 $\pm$.073	0.194 $\pm$.091	0.000 $\pm$.000	0.000 $\pm$.000	0.526 $\pm$.097	0.556$\pm$.044	0.515 $\pm$.109	0.264 $\pm$.028
	PA-GAP	Healthcare	0.348 $\pm$.119	0.268 $\pm$.142	0.413 $\pm$.036	0.332 $\pm$.093	0.000 $\pm$.000	0.234 $\pm$.054	0.317 $\pm$.180	0.489$\pm$.076
		Epidemiology	0.694$\pm$.042	0.244 $\pm$.122	0.000 $\pm$.000	0.514 $\pm$.079	0.000 $\pm$.000	0.390 $\pm$.077	0.483 $\pm$.102	0.001 $\pm$.005
		ToyGraph	0.006 $\pm$.004	0.066 $\pm$.025	0.296$\pm$.000	0.068 $\pm$.047	0.076 $\pm$.076	0.199 $\pm$.082	0.047 $\pm$.034	0.213 $\pm$.099
		Synthetic	0.156 $\pm$.059	0.063 $\pm$.034	0.154 $\pm$.036	0.023 $\pm$.050	0.113 $\pm$.082	0.000 $\pm$.000	0.286$\pm$.062	0.246 $\pm$.077
		Synthetic-2	0.004 $\pm$.001	0.068 $\pm$.037	0.000 $\pm$.000	0.158$\pm$.017	0.000 $\pm$.000	0.158$\pm$.137	0.000 $\pm$.000	0.146 $\pm$.056
		Chain-hard	0.463$\pm$.062	0.228 $\pm$.019	0.312 $\pm$.031	0.160 $\pm$.160	0.029 $\pm$.017	0.000 $\pm$.000	0.368 $\pm$.102	0.001 $\pm$.054
20	PA-GAP	Ecology	0.011 $\pm$.011	0.000 $\pm$.000	0.150 $\pm$.047	0.017 $\pm$.016	0.000 $\pm$.000	0.210$\pm$.120	0.000 $\pm$.000	0.006 $\pm$.075
		Protein-reconstructed	0.113 $\pm$.049	0.025 $\pm$.013	0.000 $\pm$.000	0.193$\pm$.097	0.031 $\pm$.036	0.000 $\pm$.000	0.096 $\pm$.080	0.009 $\pm$.075
		Healthcare	0.091 $\pm$.015	0.056 $\pm$.011	0.119 $\pm$.005	0.000 $\pm$.000	0.000 $\pm$.000	0.009 $\pm$.005	0.102 $\pm$.052	0.147$\pm$.087
		Epidemiology	0.228 $\pm$.028	0.087 $\pm$.045	0.000 $\pm$.000	0.241$\pm$.013	0.000 $\pm$.000	0.127 $\pm$.073	0.143 $\pm$.067	0.007 $\pm$.001

Table 16: Graph-misspecification stress test under edge reversal (mean $\pm$ standard error over 20 seeds; illustrative pilot stress test; higher is better; best mean among methods under the same perturbation setting is in bold).

Trial limit	Metric	Dataset	BO	CBO	cCBO	DCBO	MCBO	HCBO	CoCaBO	CEO
100	GAP	ToyGraph	0.232 $\pm$.079	0.333 $\pm$.118	0.826$\pm$.001	0.564 $\pm$.060	0.513 $\pm$.257	0.358 $\pm$.054	0.279 $\pm$.074	0.459 $\pm$.084
		Synthetic	0.651$\pm$.065	0.340 $\pm$.084	0.548 $\pm$.015	0.337 $\pm$.036	0.456 $\pm$.082	0.391 $\pm$.088	0.589 $\pm$.042	0.396 $\pm$.042
		Synthetic-2	0.001 $\pm$.004	0.281 $\pm$.125	0.335 $\pm$.114	0.000 $\pm$.004	0.000 $\pm$.000	0.757$\pm$.012	0.122 $\pm$.076	0.526 $\pm$.183
		Chain-hard	0.787 $\pm$.103	0.874$\pm$.051	0.632 $\pm$.118	0.295 $\pm$.050	0.409 $\pm$.093	0.095 $\pm$.062	0.780 $\pm$.071	0.347 $\pm$.046
		Ecology	0.160 $\pm$.128	0.012 $\pm$.012	0.426$\pm$.068	0.255 $\pm$.008	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.228 $\pm$.051
		Protein-reconstructed	0.331 $\pm$.125	0.363 $\pm$.139	0.000 $\pm$.000	0.000 $\pm$.000	0.615$\pm$.097	0.288 $\pm$.091	0.423 $\pm$.096	0.193 $\pm$.048
		Healthcare	0.315 $\pm$.054	0.561$\pm$.004	0.509 $\pm$.083	0.341 $\pm$.093	0.357 $\pm$.029	0.234 $\pm$.054	0.432 $\pm$.104	0.483 $\pm$.086
		Epidemiology	0.506 $\pm$.056	0.533$\pm$.134	0.000 $\pm$.000	0.456 $\pm$.029	0.000 $\pm$.000	0.390 $\pm$.077	0.464 $\pm$.200	0.218 $\pm$.081
	PA-GAP	ToyGraph	0.051 $\pm$.039	0.102 $\pm$.054	0.355$\pm$.001	0.095 $\pm$.056	0.271 $\pm$.141	0.199 $\pm$.082	0.139 $\pm$.079	0.203 $\pm$.099
		Synthetic	0.214 $\pm$.051	0.141 $\pm$.036	0.257 $\pm$.045	0.025 $\pm$.056	0.113 $\pm$.082	0.075 $\pm$.093	0.376$\pm$.027	0.253 $\pm$.054
		Synthetic-2	0.005 $\pm$.002	0.003 $\pm$.002	0.139 $\pm$.070	0.000 $\pm$.006	0.000 $\pm$.000	0.213$\pm$.043	0.090 $\pm$.054	0.155 $\pm$.057
		Chain-hard	0.497$\pm$.015	0.434 $\pm$.006	0.325 $\pm$.099	0.087 $\pm$.050	0.087 $\pm$.001	0.066 $\pm$.051	0.454 $\pm$.043	0.007 $\pm$.033
		Ecology	0.031 $\pm$.011	0.000 $\pm$.005	0.174 $\pm$.038	0.041 $\pm$.008	0.000 $\pm$.000	0.210$\pm$.120	0.000 $\pm$.000	0.007 $\pm$.069
		Protein-reconstructed	0.153$\pm$.061	0.057 $\pm$.018	0.000 $\pm$.000	0.000 $\pm$.000	0.031 $\pm$.036	0.000 $\pm$.000	0.018 $\pm$.009	0.010 $\pm$.072
		Healthcare	0.120 $\pm$.021	0.090 $\pm$.014	0.202$\pm$.014	0.000 $\pm$.000	0.071 $\pm$.051	0.010 $\pm$.006	0.196 $\pm$.010	0.138 $\pm$.049
		Epidemiology	0.242 $\pm$.021	0.173 $\pm$.077	0.000 $\pm$.000	0.065 $\pm$.074	0.000 $\pm$.000	0.138 $\pm$.080	0.251$\pm$.094	0.001 $\pm$.042
50	GAP	ToyGraph	0.202 $\pm$.099	0.477 $\pm$.097	0.780$\pm$.009	0.525 $\pm$.062	0.411 $\pm$.206	0.450 $\pm$.093	0.276 $\pm$.075	0.489 $\pm$.073
		Synthetic	0.567 $\pm$.091	0.534 $\pm$.087	0.425 $\pm$.086	0.295 $\pm$.050	0.456 $\pm$.099	0.391 $\pm$.088	0.569$\pm$.118	0.254 $\pm$.057
		Synthetic-2	0.007 $\pm$.005	0.310 $\pm$.033	0.243 $\pm$.127	0.000 $\pm$.002	0.000 $\pm$.000	0.536$\pm$.099	0.122 $\pm$.076	0.484 $\pm$.222
		Chain-hard	0.686 $\pm$.086	0.747$\pm$.102	0.411 $\pm$.128	0.295 $\pm$.050	0.383 $\pm$.096	0.095 $\pm$.062	0.744 $\pm$.062	0.356 $\pm$.101
		Ecology	0.259 $\pm$.148	0.000 $\pm$.005	0.502$\pm$.105	0.158 $\pm$.008	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.101 $\pm$.031
		Protein-reconstructed	0.365 $\pm$.061	0.413 $\pm$.046	0.000 $\pm$.000	0.000 $\pm$.000	0.526 $\pm$.097	0.591$\pm$.040	0.342 $\pm$.170	0.264 $\pm$.028
		Healthcare	0.439 $\pm$.121	0.522 $\pm$.020	0.604$\pm$.007	0.332 $\pm$.093	0.357 $\pm$.029	0.234 $\pm$.054	0.342 $\pm$.077	0.489 $\pm$.076
		Epidemiology	0.358 $\pm$.052	0.450 $\pm$.170	0.000 $\pm$.000	0.456 $\pm$.039	0.000 $\pm$.000	0.390 $\pm$.077	0.624$\pm$.150	0.060 $\pm$.050
	PA-GAP	ToyGraph	0.024 $\pm$.016	0.089 $\pm$.047	0.338$\pm$.006	0.088 $\pm$.054	0.225 $\pm$.134	0.199 $\pm$.082	0.131 $\pm$.073	0.213 $\pm$.099
		Synthetic	0.194 $\pm$.054	0.123 $\pm$.034	0.214 $\pm$.051	0.023 $\pm$.050	0.113 $\pm$.082	0.075 $\pm$.093	0.328$\pm$.036	0.246 $\pm$.077
		Synthetic-2	0.003 $\pm$.003	0.003 $\pm$.002	0.110 $\pm$.059	0.000 $\pm$.005	0.000 $\pm$.000	0.104 $\pm$.055	0.090 $\pm$.054	0.146$\pm$.056
		Chain-hard	0.481$\pm$.029	0.373 $\pm$.012	0.298 $\pm$.088	0.087 $\pm$.050	0.060 $\pm$.010	0.066 $\pm$.051	0.445 $\pm$.051	0.001 $\pm$.054
		Ecology	0.017 $\pm$.010	0.000 $\pm$.002	0.166 $\pm$.040	0.031 $\pm$.010	0.000 $\pm$.000	0.210$\pm$.120	0.000 $\pm$.000	0.006 $\pm$.075
		Protein-reconstructed	0.137$\pm$.054	0.047 $\pm$.017	0.000 $\pm$.000	0.000 $\pm$.000	0.031 $\pm$.036	0.000 $\pm$.000	0.018 $\pm$.009	0.009 $\pm$.075
		Healthcare	0.105 $\pm$.017	0.082 $\pm$.010	0.178 $\pm$.010	0.000 $\pm$.000	0.071 $\pm$.051	0.010 $\pm$.006	0.180$\pm$.008	0.147 $\pm$.087
		Epidemiology	0.232$\pm$.025	0.150 $\pm$.083	0.000 $\pm$.000	0.065 $\pm$.098	0.000 $\pm$.000	0.138 $\pm$.080	0.225 $\pm$.093	0.008 $\pm$.002
20	GAP	ToyGraph	0.272 $\pm$.097	0.264 $\pm$.152	0.639$\pm$.000	0.405 $\pm$.074	0.496 $\pm$.275	0.450 $\pm$.093	0.167 $\pm$.073	0.489 $\pm$.073
		Synthetic	0.390 $\pm$.138	0.350 $\pm$.208	0.518$\pm$.156	0.295 $\pm$.050	0.456 $\pm$.099	0.391 $\pm$.088	0.506 $\pm$.067	0.254 $\pm$.057
		Synthetic-2	0.003 $\pm$.002	0.313 $\pm$.163	0.000 $\pm$.000	0.507 $\pm$.026	0.000 $\pm$.000	0.536$\pm$.099	0.000 $\pm$.000	0.484 $\pm$.222
		Chain-hard	0.761 $\pm$.124	0.650 $\pm$.090	0.691 $\pm$.051	0.866 $\pm$.021	0.249 $\pm$.090	0.000 $\pm$.000	0.905$\pm$.066	0.356 $\pm$.101
		Ecology	0.145 $\pm$.145	0.000 $\pm$.000	0.422$\pm$.159	0.221 $\pm$.147	0.000 $\pm$.000	0.234 $\pm$.102	0.000 $\pm$.000	0.101 $\pm$.031
		Protein-reconstructed	0.370 $\pm$.073	0.194 $\pm$.091	0.000 $\pm$.000	0.000 $\pm$.000	0.526 $\pm$.097	0.556$\pm$.044	0.326 $\pm$.163	0.264 $\pm$.028
		Healthcare	0.348 $\pm$.119	0.401 $\pm$.079	0.413 $\pm$.036	0.332 $\pm$.093	0.000 $\pm$.000	0.234 $\pm$.054	0.407 $\pm$.108	0.489$\pm$.076
		Epidemiology	0.694$\pm$.042	0.378 $\pm$.226	0.000 $\pm$.000	0.105 $\pm$.053	0.000 $\pm$.000	0.390 $\pm$.077	0.541 $\pm$.079	0.001 $\pm$.005
	PA-GAP	ToyGraph	0.006 $\pm$.004	0.056 $\pm$.028	0.296$\pm$.000	0.068 $\pm$.047	0.161 $\pm$.157	0.199 $\pm$.082	0.113 $\pm$.059	0.213 $\pm$.099
		Synthetic	0.156 $\pm$.059	0.088 $\pm$.038	0.174 $\pm$.061	0.023 $\pm$.050	0.113 $\pm$.082	0.000 $\pm$.000	0.259$\pm$.047	0.246 $\pm$.077
		Synthetic-2	0.004 $\pm$.001	0.065 $\pm$.036	0.000 $\pm$.000	0.158$\pm$.017	0.000 $\pm$.000	0.104 $\pm$.055	0.000 $\pm$.000	0.146 $\pm$.056
		Chain-hard	0.463$\pm$.062	0.218 $\pm$.023	0.231 $\pm$.057	0.414 $\pm$.021	0.024 $\pm$.020	0.000 $\pm$.000	0.426 $\pm$.069	0.001 $\pm$.054
		Ecology	0.011 $\pm$.011	0.000 $\pm$.000	0.150 $\pm$.047	0.017 $\pm$.016	0.000 $\pm$.000	0.210$\pm$.120	0.000 $\pm$.000	0.006 $\pm$.075
		Protein-reconstructed	0.113 $\pm$.049	0.025 $\pm$.013	0.000 $\pm$.000	0.193$\pm$.097	0.031 $\pm$.036	0.000 $\pm$.000	0.017 $\pm$.009	0.009 $\pm$.075
		Healthcare	0.091 $\pm$.015	0.064 $\pm$.006	0.119 $\pm$.005	0.000 $\pm$.000	0.000 $\pm$.000	0.010 $\pm$.006	0.163$\pm$.005	0.147 $\pm$.087
		Epidemiology	0.228$\pm$.028	0.139 $\pm$.087	0.000 $\pm$.000	0.003 $\pm$.001	0.000 $\pm$.000	0.138 $\pm$.080	0.176 $\pm$.082	0.007 $\pm$.001




Figure 21: Graph-misspecification schematics for the synthetic SCMs (Synthetic, Synthetic-2, and Chain-hard). Visual encoding follows Figure 20: each row shows the true SCM followed by learner-belief graphs under the edge-deletion, edge-addition, and edge-reversal perturbations named in the panel titles. Only the learner’s supplied graph is perturbed; the true SCM and y^* are unchanged.




Figure 22: Graph-misspecification schematics for Protein-reconstructed and the domain-inspired benchmark SCMs (Ecology, Healthcare, and Epidemiology). Layout and visual encoding follow Figure 21. Each row shows the true SCM followed by learner-belief graphs under the misspecification perturbations shown in the panel titles. Protein-reconstructed uses only the three edge perturbations: $\text{PKC} \rightarrow \text{PKA}$ deletion, $\text{PKC} \rightarrow \text{Erk}$ addition, and $\text{PKC} \rightarrow \text{PKA}$ reversal to $\text{PKA} \rightarrow \text{PKC}$. For Ecology, the perturbations are $\text{TA} \rightarrow \Omega_A$ deletion, $\text{Sal} \rightarrow \text{NEC}$ addition, $\text{TA} \rightarrow \Omega_A$ reversal to $\Omega_A \rightarrow \text{TA}$, and hiding Tem. For Healthcare, the perturbations are $\text{BMI} \rightarrow \text{aspirin}$ deletion, $\text{aspirin} \rightarrow \text{statin}$ addition, $\text{BMI} \rightarrow \text{aspirin}$ reversal to $\text{aspirin} \rightarrow \text{BMI}$, and hiding BMI. For Epidemiology, the perturbations are $T \rightarrow L$ deletion, $B \rightarrow R$ addition, $T \rightarrow L$ reversal to $L \rightarrow T$, and hiding T . The Ecology schematic is illustrative and uses representative perturbations rather than a run-pilot setting. The Healthcare hide-BMI and Epidemiology hide- T panels are the omitted common causes used in the omitted-variable stress test of Table 17. Only the learner’s supplied graph is perturbed; the true SCM and y^* are unchanged.

B.1.2 Hidden confounding (omitted-variable stress test).

Hidden confounding is a distinct assumption violation, and the graph-perturbation study above does not address it. We therefore added a companion omitted-variable stress test (Table 17). Here the learner’s graph and observational view omit a genuine common cause (BMI in Healthcare, X in Chain-hard, and T in Epidemiology) while the true SCM and reference optimum y^* are held fixed, so the comparison isolates the effect of an unobserved confounder rather than changing the task. This omitted-variable stress test is evaluated over 20 random seeds. We currently include BO as a graph-free control and CBO/CoCaBO as belief-graph methods.

The effect is method- and dataset-dependent. Omitting the confounder degrades CBO on Chain-hard, whereas on Healthcare and some Epidemiology settings the omitted-variable runs of CBO and CoCaBO sometimes match or slightly exceed their full-information counterparts. We do not interpret these numerical increases as evidence that confounding is beneficial or harmless. Rather, they reflect finite-budget search effects: a coarser or misspecified observational view can weaken an overconfident causal prior, change candidate scopes, or accidentally induce a more favorable acquisition trajectory. The graph-free BO baseline is unchanged by construction, since it does not use the omitted causal structure. Overall, the omitted-variable stress test reinforces the same message as the graph-perturbation study: current CBO pipelines are sensitive to the causal information they are given, and robustness to hidden confounding remains an open evaluation dimension.

Table 17: Hidden-confounding / omitted-variable stress test (mean $\pm$ standard error over 20 seeds; higher is better; best mean per row in bold). The learner’s graph and observational view omit a genuine common cause (Healthcare: BMI; Chain-hard: X ; Epidemiology: T); the true SCM and y^* are unchanged. BO is graph-free and uses the full variable set, so its Full and Hidden rows are identical by construction.

Trial limit	Metric	Dataset	Information	BO	CBO	CoCaBO
100	GAP	Healthcare	Full information	0.309 $\pm$.057	0.510$\pm$.035	0.226 $\pm$.118
			Hidden variable	0.309 $\pm$.057	0.543$\pm$.012	0.251 $\pm$.126
		Chain-hard	Full information	0.780 $\pm$.102	0.814$\pm$.128	0.761 $\pm$.129
			Hidden variable	0.780$\pm$.102	0.695 $\pm$.106	0.770 $\pm$.139
		Epidemiology	Full information	0.508 $\pm$.057	0.563$\pm$.038	0.424 $\pm$.178
			Hidden variable	0.508 $\pm$.057	0.546$\pm$.116	0.399 $\pm$.094
	PA-GAP	Healthcare	Full information	0.122 $\pm$.018	0.077 $\pm$.015	0.125$\pm$.062
			Hidden variable	0.122 $\pm$.018	0.080 $\pm$.006	0.139$\pm$.070
		Chain-hard	Full information	0.492$\pm$.013	0.437 $\pm$.005	0.486 $\pm$.008
			Hidden variable	0.492$\pm$.013	0.420 $\pm$.010	0.479 $\pm$.008
		Epidemiology	Full information	0.244$\pm$.023	0.185 $\pm$.059	0.163 $\pm$.080
			Hidden variable	0.244$\pm$.023	0.182 $\pm$.057	0.130 $\pm$.091
50	GAP	Healthcare	Full information	0.438$\pm$.124	0.432 $\pm$.073	0.225 $\pm$.128
			Hidden variable	0.438 $\pm$.124	0.502$\pm$.020	0.378 $\pm$.196
		Chain-hard	Full information	0.680 $\pm$.087	0.889$\pm$.016	0.815 $\pm$.146
			Hidden variable	0.680 $\pm$.087	0.811 $\pm$.054	0.873$\pm$.102
		Epidemiology	Full information	0.357 $\pm$.053	0.355 $\pm$.021	0.565$\pm$.076
			Hidden variable	0.357 $\pm$.053	0.590$\pm$.103	0.344 $\pm$.188
	PA-GAP	Healthcare	Full information	0.111 $\pm$.017	0.069 $\pm$.014	0.114$\pm$.057
			Hidden variable	0.111 $\pm$.017	0.076 $\pm$.005	0.131$\pm$.066
		Chain-hard	Full information	0.484$\pm$.026	0.379 $\pm$.010	0.474 $\pm$.016
			Hidden variable	0.484$\pm$.026	0.348 $\pm$.018	0.468 $\pm$.013
		Epidemiology	Full information	0.230$\pm$.022	0.129 $\pm$.061	0.141 $\pm$.077
			Hidden variable	0.230$\pm$.022	0.162 $\pm$.065	0.102 $\pm$.084
20	GAP	Healthcare	Full information	0.348$\pm$.120	0.268 $\pm$.142	0.328 $\pm$.180
			Hidden variable	0.348 $\pm$.120	0.376$\pm$.048	0.250 $\pm$.159
		Chain-hard	Full information	0.761 $\pm$.124	0.718 $\pm$.039	0.913$\pm$.047
			Hidden variable	0.761 $\pm$.124	0.615 $\pm$.090	0.938$\pm$.023
		Epidemiology	Full information	0.694$\pm$.042	0.244 $\pm$.122	0.389 $\pm$.071
			Hidden variable	0.694$\pm$.042	0.429 $\pm$.181	0.156 $\pm$.156
	PA-GAP	Healthcare	Full information	0.091 $\pm$.015	0.056 $\pm$.011	0.100$\pm$.051
			Hidden variable	0.091 $\pm$.015	0.068 $\pm$.003	0.112$\pm$.057
		Chain-hard	Full information	0.463$\pm$.062	0.228 $\pm$.020	0.449 $\pm$.038
			Hidden variable	0.463$\pm$.062	0.178 $\pm$.029	0.437 $\pm$.033
		Epidemiology	Full information	0.228$\pm$.028	0.087 $\pm$.045	0.104 $\pm$.069
			Hidden variable	0.228$\pm$.028	0.128 $\pm$.080	0.068 $\pm$.068

Table 18: MCBO GAP and PA-GAP vs. UCB β at $T \in \{100, 50, 20\}$, $\beta=10$ columns are the **main_hard** values (the parameter used in the paper); $\beta \in \{0.5, 1, 2, 5\}$ are 20-seed sweep re-runs (100-trial trajectory truncated to T). Higher is better; best score per row in bold.

Metric	T	Dataset	$\beta=0.5$	$\beta=1$	$\beta=2$	$\beta=5$	$\beta=10$
GAP	100	ToyGraph	0.748	0.381	0.746	0.357	0.547
		Synthetic	0.659	0.686	0.684	0.681	0.652
		Synthetic-2	0.000	0.000	0.000	0.000	0.000
		Chain-hard	0.516	0.554	0.435	0.423	0.502
		Ecology	0.000	0.000	0.000	0.000	0.000
		Protein-reconstructed	0.419	0.337	0.109	0.286	0.878
		Healthcare	0.485	0.524	0.545	0.519	0.510
		Epidemiology	0.000	0.000	0.000	0.000	0.000
	50	ToyGraph	0.545	0.576	0.492	0.253	0.513
		Synthetic	0.605	0.606	0.651	0.587	0.630
		Synthetic-2	0.000	0.000	0.000	0.000	0.000
		Chain-hard	0.392	0.540	0.381	0.367	0.215
		Ecology	0.000	0.000	0.000	0.000	0.000
		Protein-reconstructed	0.354	0.348	0.264	0.144	0.752
		Healthcare	0.303	0.292	0.275	0.330	0.315
		Epidemiology	0.000	0.000	0.000	0.000	0.000
	20	ToyGraph	0.333	0.530	0.282	0.171	0.406
		Synthetic	0.523	0.574	0.573	0.578	0.558
		Synthetic-2	0.000	0.000	0.000	0.000	0.000
		Chain-hard	0.420	0.432	0.422	0.207	0.360
		Ecology	0.000	0.000	0.000	0.000	0.000
		Protein-reconstructed	0.393	0.300	0.140	0.399	0.540
		Healthcare	0.000	0.000	0.000	0.000	0.000
		Epidemiology	0.000	0.000	0.000	0.000	0.000
PA-GAP	100	ToyGraph	0.266	0.295	0.218	0.127	0.071
		Synthetic	0.149	0.154	0.136	0.125	0.162
		Synthetic-2	0.000	0.000	0.000	0.000	0.000
		Chain-hard	0.112	0.174	0.142	0.067	0.111
		Ecology	0.000	0.000	0.000	0.000	0.000
		Protein-reconstructed	0.033	0.034	0.011	0.047	0.045
		Healthcare	0.127	0.097	0.053	0.136	0.100
		Epidemiology	0.000	0.000	0.000	0.000	0.000
	50	ToyGraph	0.186	0.261	0.136	0.089	0.063
		Synthetic	0.162	0.165	0.123	0.186	0.150
		Synthetic-2	0.000	0.000	0.000	0.000	0.000
		Chain-hard	0.089	0.152	0.119	0.035	0.078
		Ecology	0.000	0.000	0.000	0.000	0.000
		Protein-reconstructed	0.033	0.034	0.011	0.046	0.045
		Healthcare	0.016	0.032	0.012	0.034	0.034
		Epidemiology	0.000	0.000	0.000	0.000	0.000
	20	ToyGraph	0.175	0.171	0.093	0.028	0.040
		Synthetic	0.137	0.137	0.116	0.141	0.116
		Synthetic-2	0.000	0.000	0.000	0.000	0.000
		Chain-hard	0.049	0.116	0.094	0.016	0.034
		Ecology	0.000	0.000	0.000	0.000	0.000
		Protein-reconstructed	0.034	0.034	0.011	0.043	0.047
		Healthcare	0.000	0.000	0.000	0.000	0.000
		Epidemiology	0.000	0.000	0.000	0.000	0.000

B.2 MCBO β sensitivity

The main benchmark keeps the MCBO wrapper setting $\beta = 10$, matching the runs reported in Table 6. To assess how much this default matters, we reran MCBO with $\beta \in \{0.5, 1, 2, 5\}$ and compare the resulting GAP and PA-GAP values with the main $\beta = 10$ results. The best-performing β is not consistent across datasets, budgets, or metrics: smaller values are often better on ToyGraph and Chain-hard, $\beta = 5$ is competitive on several PA-GAP rows, and $\beta = 10$ remains strongest on Protein-reconstructed. We therefore report the sensitivity sweep as an appendix diagnostic and avoid interpreting the default MCBO performance as an intrinsic property of the method.

B.3 Ranking Statistical Testing

The main text reports no pooled cross-method rank figures, because the compared pipelines do not all solve the same optimization problem (Table 5). The tables in this appendix are the benchmark’s only cross-method aggregation: they are retained as a *descriptive* reliability record, and we assess whether such rank

Table 19: Ranking reliability on the hard-intervention benchmark by budget. Ranks (1=best) are computed within each dataset–metric setting (8 datasets $\times$ 2 metrics, $N=16$ per budget). Entries show average rank $\pm$ s.d. and 95% bootstrap CI; lower is better. Bottom rows report the Friedman test, Nemenyi critical difference versus rank spread, and a conservative per-dataset Friedman test.

Method	$T=100$		$T=50$		$T=20$	
	Rank $\pm$ s.d.	95% CI	Rank $\pm$ s.d.	95% CI	Rank $\pm$ s.d.	95% CI
CoCaBO	3.38 $\pm$ 2.75	[2.12, 4.69]	3.44 $\pm$ 2.68	[2.22, 4.75]	3.88 $\pm$ 2.82	[2.56, 5.25]
cCBO	3.91 $\pm$ 2.23	[2.88, 5.03]	3.69 $\pm$ 2.32	[2.66, 4.84]	3.88 $\pm$ 2.24	[2.84, 4.97]
CBO	3.50 $\pm$ 1.59	[2.81, 4.31]	3.94 $\pm$ 1.18	[3.44, 4.50]	4.31 $\pm$ 1.35	[3.62, 4.94]
BO	4.12 $\pm$ 2.84	[2.78, 5.53]	3.94 $\pm$ 2.69	[2.66, 5.22]	3.81 $\pm$ 2.48	[2.62, 5.00]
CEO	4.78 $\pm$ 1.70	[3.94, 5.56]	4.88 $\pm$ 1.96	[3.94, 5.81]	4.25 $\pm$ 1.98	[3.31, 5.19]
DCBO	5.03 $\pm$ 2.22	[3.94, 6.06]	4.78 $\pm$ 2.27	[3.75, 5.84]	4.12 $\pm$ 2.18	[3.09, 5.22]
HCBO	5.38 $\pm$ 1.89	[4.44, 6.25]	5.44 $\pm$ 1.86	[4.50, 6.25]	5.84 $\pm$ 2.19	[4.75, 6.84]
MCBO	5.91 $\pm$ 1.89	[4.97, 6.72]	5.91 $\pm$ 2.19	[4.81, 6.88]	5.91 $\pm$ 2.01	[4.91, 6.81]
Friedman p	0.0272		0.0384		0.0462	
Nemenyi CD / spread	2.62 / 2.53		2.62 / 2.47		2.62 / 2.09	
Per-dataset p ($N=8$)	0.278		0.342		0.406	

summaries would support statistically separable method differences. For each fixed budget, methods are ranked within each dataset–metric setting, where rank 1 denotes the best mean score in that setting. We then summarize the average rank across settings and report bootstrap confidence intervals over settings. We also report a Friedman omnibus test and the Nemenyi critical difference for post-hoc rank separation. Because GAP and PA-GAP are computed from the same trajectories, we additionally report a conservative per-dataset Friedman test in which the two metric ranks are collapsed within each dataset. An important caveat applies to all of these tests: the pooled settings are not mutually independent. GAP and PA-GAP share the same underlying trajectories, and the three budgets $T \in \{100, 50, 20\}$ are nested truncations of the same runs rather than independent replications. The reported p -values and critical differences should therefore be read as descriptive reliability checks under these dependencies, not as exact inference over independent samples; the conservative per-dataset test partially mitigates the metric-sharing dependence but not the budget nesting.

B.3.1 Hard-intervention

The hard-intervention results show that average-rank orderings vary with budget. CoCaBO has the lowest average rank at $T = 100$ and $T = 50$, while BO has the lowest average rank at $T = 20$. However, these rank differences should be interpreted cautiously. Although the Friedman omnibus test detects rank heterogeneity at each budget, the rank spread is always smaller than the Nemenyi critical difference, so no pair of methods is statistically separable by the post-hoc test. The more conservative per-dataset Friedman test also does not reject at any budget. Thus, the ranking analysis supports the qualitative conclusion that performance is heterogeneous, but it does not support declaring a single hard-intervention method as uniformly superior.

B.3.2 Soft-intervention

The soft-intervention results show a consistent descriptive ordering: ACBO has the lowest average rank at all three budgets, followed by BO and then MCBO. Nevertheless, the statistical tests do not support a separable winner. The Friedman omnibus test does not reject at any budget, the rank spread remains below the Nemenyi critical difference, and the conservative per-dataset test also does not reject. We therefore treat the soft-intervention ranking as descriptive evidence that ACBO is often strong in this benchmark, rather than as a statistically conclusive superiority claim.

B.4 Average Reward Visualization

GAP, PA-GAP, and best-so-far curves summarize incumbent quality, but they do not show the quality of all evaluations made along the way. We therefore add average-reward trajectories as a complementary

Table 20: Ranking reliability on the soft-intervention benchmark by budget. Ranks (1=best) are computed within each dataset–metric setting (5 datasets $\times$ 2 metrics, $N=10$ per budget). Entries show average rank $\pm$ s.d. and 95% bootstrap CI; lower is better. Bottom rows report the Friedman test, Nemenyi critical difference versus rank spread, and a conservative per-dataset Friedman test.

Method	$T=100$		$T=50$		$T=20$	
	Rank $\pm$ s.d.	95% CI	Rank $\pm$ s.d.	95% CI	Rank $\pm$ s.d.	95% CI
ACBO	1.60 $\pm$ 0.84	[1.10, 2.10]	1.60 $\pm$ 0.70	[1.20, 2.00]	1.60 $\pm$ 0.66	[1.20, 2.00]
BO	1.80 $\pm$ 0.79	[1.40, 2.30]	1.90 $\pm$ 0.74	[1.50, 2.30]	2.00 $\pm$ 0.78	[1.55, 2.45]
MCBO	2.60 $\pm$ 0.52	[2.30, 2.90]	2.50 $\pm$ 0.85	[2.00, 3.00]	2.40 $\pm$ 0.77	[1.90, 2.80]
Friedman p	0.0608		0.1225		0.1690	
Nemenyi CD / spread	1.05 / 1.00		1.05 / 0.90		1.05 / 0.80	
Per-dataset p ($N=5$)	0.211		0.311		0.390	

diagnostic. For each run, rewards are normalized so that the reference optimum is 1 in the task’s natural optimization direction. This makes average-reward curves comparable across datasets while preserving their interpretation as the mean quality of the sequence of queried interventions. These plots are not used as the primary ranking criterion, because exploratory acquisitions can intentionally query uncertain points with poor immediate reward, but they help verify that the benchmark conclusions do not depend only on final incumbent values.

Figures 23 and 24 should be read together with the best-so-far trajectories and scalar metrics in the main text. When a method has strong best-so-far performance but weaker average reward, the optimizer is finding good incumbents while spending part of its budget on exploratory or poorly calibrated trials. Conversely, strong average reward with weaker final GAP indicates reliable early sampling without necessarily discovering the best final intervention. This reinforces the paper’s main evaluation message: CBO methods should be compared using multiple trajectory summaries rather than a single scalar metric.




Figure 23: Average-reward trajectories for the hard-intervention benchmark.




Figure 24: Average-reward trajectories for the soft-intervention/function-network benchmark.