ROBUST CAUSAL DISCOVERY VIA ABDUCTIVE QUANTITATIVE ARGUMENTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Constraint-based causal discovery is notoriously fragile under finite-sample errors, which often trigger devastating structural error cascades. To address this, we propose **Abductive Quantitative Argumentation for Causal Discovery (AQACD)**, a probabilistic-symbolic framework that reformulates causal learning as abductive argumentation. Unlike deterministic algorithms or exact logic solvers, AQACD models causal relations as defeasible hypotheses, dynamically evaluated through a continuous dialectic network of statistical and structural arguments. Rather than making irreversible commitments, graph revisions are guided by Inference to the Best Explanation: the topology is updated only when an alternative explanation decisively outweighs its rivals under a dialectic margin, rescuing mistakenly pruned edges (τ_{triad}). This approach strictly improves global internal consistency while bypassing the exponential bottlenecks typical of exact logical inference. Extensive experiments on standard benchmarks demonstrate that AQACD consistently outperforms classical and continuous baselines. Most notably, it achieves substantial gains in orientation accuracy (Arrow F1) within Markov Equivalence Classes, exhibiting exceptional robustness to noise and neutralizing error cascades.

1 INTRODUCTION

Discovering causal relationships from observational data is a cornerstone of robust decision-making and explainable AI (Glymour et al., 2019; Zhu et al., 2020; Kuang et al., 2020). Existing causal discovery methods mainly follow three classical paradigms: *constraint-based* approaches (e.g., PC (Spirtes et al., 2001)), which evaluate conditional independence (CI) tests; *score-based* approaches (e.g., GES (Chickering, 2002)); and *continuous optimization* methods (e.g., NOTEARS (Zheng et al., 2018), GOLEM (Ng et al., 2020)).

Despite their different formulations, these approaches are notoriously fragile under finite-sample uncertainty, particularly when statistical CI tests produce noisy or contradictory outcomes (Vowels et al., 2022). Constraint-based methods make irreversible sequential commitments: an isolated false CI test in early phases permanently removes an adjacency, triggering substantial structural error cascades across downstream collider orientation and Meek propagation (Colombo & Maathuis, 2014). Conflict-aware heuristics such as Conservative PC (CPC) (Ramsey et al., 2006) and Majority PC (MPC) (Colombo & Maathuis, 2014) attempt to mitigate orientation cascades, yet they remain strictly monotonic—permanently discarding disputed adjacencies without any mechanism to retroactively recover mistakenly pruned edges. Exact combinatorial formulations based on SAT or Answer Set Programming (ASP) (Hyttinen et al., 2014) guarantee global optimality, but treat statistical outcomes as rigid binary clauses and incur worst-case exponential complexity $\mathcal{O}(2^{|\mathcal{V}|})$. Continuous relaxations provide scalable optimization, yet remain sensitive to distributional assumptions and observational equivalence across Markov Equivalence Classes (MECs) (Verma & Pearl, 1990; Peters et al., 2017; Andrews et al., 2023). Fundamentally, resolving contradictory statistical evidence across sequential learning stages poses a challenge of **belief revision** (Gärdenfors, 1988): structural commitments must remain dynamically revisable in light of global topological coherence.

While recent argumentation-based methods evaluate causal constraints (Wei et al., 2026), they act as *passive post-hoc evaluators* over static PC skeletons, inheriting early edge deletions irreversibly. We overcome this by formulating causal discovery as **abductive belief revision** (Gärdenfors, 1988;

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Zhou, 2019), grounded in *Inference to the Best Explanation* (IBE) (Harman, 1965). We propose **Abductive Quantitative Argumentation for Causal Discovery (AQACD)**, a probabilistic-symbolic framework modeling causal adjacencies and orientations as defeasible hypotheses. Unlike static evaluators, AQACD dynamically evaluates topologies via continuous statistical and structural arguments. Guided by IBE, it triggers dynamic triad rescue (τ_{triad}) to absorb false CI tests, resurrecting mistakenly pruned edges and actively neutralizing error cascades. Our primary contributions are summarized as follows:

- **Abductive Causal Discovery Paradigm:** We reformulate causal learning as quantitative abductive argumentation. By optimizing a global Internal Consistency Score $\Phi(\mathcal{G})$ over bidirectional edits, AQACD enables non-monotonic belief revision to recover from finite-sample statistical errors.
- **Dynamic IBE Structural Rescue and Guarded Revision:** We synthesize heterogeneous statistical p-values, collider fidelity, and Bayesian scores into continuous argument strengths, using dialectical margins and risk guards to resolve orientations without violating observational equivalence.
- **Polynomial Scalability and Soundness:** We prove that abductive revision operates in strictly polynomial time $\mathcal{O}(|\mathcal{V}|^{d_{\max}+2} + |\mathcal{V}|^5)$ for fixed $d_{\max}$—bypassing exact logic solver blowups—while guaranteeing asymptotic recovery of the true completed partially directed acyclic graph (CPDAG) under an identifiable faithfulness margin (Theorem 2).

2 Related Work

2.1 Causal Discovery and Conflicting Evidence

Constraint-based methods such as PC and FCI rely on sequential CI tests and are notoriously prone to propagating finite-sample errors across skeleton and orientation stages (Spirtes, 2001; Colombo & Maathuis, 2014). To handle contradictory separating sets, CPC (Ramsey et al., 2006) and MPC (Colombo & Maathuis, 2014) evaluate all candidate conditioning sets for unshielded triples. While effective at halting cascade propagation, their rule-based determinism discards ambiguous triples entirely, leading to severe edge omissions and an inability to recover mistakenly deleted skeleton adjacencies. Score-based methods such as GES (Chickering, 2002) optimize global graph scores, while continuous methods such as NOTEARS and GOLEM (Zheng et al., 2018; Ng et al., 2020) formulate directed acyclic graph (DAG) search as continuous optimization. However, continuous approaches remain sensitive to distributional assumptions and can be unstable when distinguishing Markov-equivalent configurations in finite samples. Our approach complements these paradigms by explicitly modeling conflicting statistical and structural constraints as continuous defeasible arguments and enabling non-monotonic revision.

2.2 Exact Inference and Defeasible Argumentation

Exact approaches based on SAT and ASP (Hyttinen et al., 2014; Zhalama et al., 2017) formulate causal discovery via combinatorial constraint satisfaction. While weighted ASP (Hyttinen et al., 2014) and MaxSAT relaxations (Zhalama et al., 2017) accommodate conflicting CI outcomes globally, their combinatorial optimization over discrete clauses remains NP-hard with worst-case exponential complexity $\mathcal{O}(2^{|\mathcal{V}|})$.

Formal argumentation provides a principled foundation for reasoning under inconsistency (Dung, 1995; Rahwan & Simari, 2009). In causal discovery, ABA-PC (Russo et al., 2024) resolves conflicting CI statements symbolically by optimizing qualitative assumptions, but incurs exponential assumption growth and lacks continuous topological revision. To accommodate graded statistical uncertainty, Quantitative Bipolar Argumentation Frameworks (QBAFs) (Cayrol & Lagasquie-Schiex, 2005; Baroni et al., 2015; Chi & Liao, 2022) evaluate continuous argument strengths.

Paradigm Shift over Static Argumentation. Most recently, QACD (Wei et al., 2026) applied QBAFs to evaluate causal constraints. However, AQACD fundamentally diverges from QACD: while QACD functions as a *passive, feed-forward filter* over a static skeleton fixed by PC (permanently inheriting early edge deletions), AQACD operates as an *active, closed-loop belief revision engine*. By navigating a non-monotonic search space ($\{+e, -e\}$) with IBE structural support (τ_{triad}),

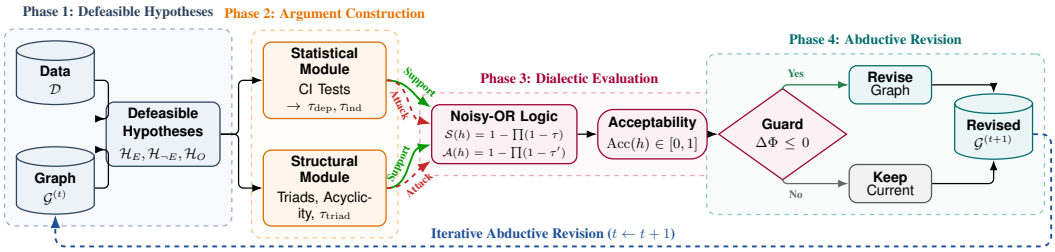


Figure 1: The overall architecture of the AQACD framework. The pipeline operates in four integrated phases: (1) formulating adjacencies and orientations as defeasible hypotheses; (2) extracting continuous supporting (green) and attacking (red) arguments from statistical tests and topological consistency; (3) aggregating evidential strengths via Noisy-OR gradual semantics to compute acceptability scores $\text{Acc}(h)$; and (4) executing consistency-guarded graph revisions. The system iterates until no accepted local revision further decreases the consistency objective.

AQACD dynamically absorbs false CI tests and resurrects falsely pruned adjacencies—actively neutralizing structural error cascades.

2.3 Abductive Reasoning for Causal Discovery

Abductive reasoning, or IBE, infers the most plausible explanation for observed phenomena (Zhou, 2019; Peirce, 1950; Harman, 1965). While abductive learning bridges machine learning and knowledge representation (Zhou, 2019), its application to causal discovery remains largely unexplored. Traditional methods rely either on rigid CI deduction or heuristic score search. AQACD bridges these paradigms by treating candidate graphs as competing abductive explanations: observed dependencies serve as evidence, and updates are accepted only when an alternative topology notably improves global consistency.

3 Methodology

3.1 Argumentative and Abductive Foundations

AQACD formulates causal discovery as an abductive reasoning process (Figure 1). Formally, we define causal abduction via a four-tuple $\langle \mathcal{O}, \mathcal{T}, \mathcal{E}\text{adm}, \preceq \rangle$:

- **Observations $\mathcal{O}$:** Finite-sample statistical evidence from CI test p-values and local penalized-likelihood score variations;
- **Background Theory $\mathcal{T}$:** Structural constraints including acyclicity and d-separation consistency;
- **Admissible Explanations $\mathcal{E}\text{adm}$:** The space of valid partially directed acyclic graph (PDAG) explanations over $\mathcal{V}$ satisfying these constraints;
- **Preference Relation $\preceq$:** A partial order induced by the global inconsistency objective $\Phi(\mathcal{G})$, ranking explanations by dialectical acceptability.

Rather than making irreversible binary decisions from individual CI tests, AQACD represents structural relations as defeasible hypotheses within a QBAF $\langle \mathcal{H}, \mathcal{R}_s, \mathcal{R}a \rangle$. Continuous gradual semantics (Cayrol & Lagasquie-Schiex, 2005; Baroni et al., 2015) maps heterogeneous evidence to $[0, 1]$. Conflicting evidence is reconciled by abductively revising the current graph toward an admissible explanation $\mathcal{G}^* \in \mathcal{E}\text{adm}$ that minimizes $\Phi(\mathcal{G})$. The conservative mode returns a CPDAG, while optional guarded inductive orientation may further resolve reversible edges into an Inductively-Resolved PDAG.

3.2 Phase 1: Defeasible Hypotheses Generation

Let $\mathcal{V} = \{X_1, \ldots, X_n\}$ denote the observed variables. For each unordered pair $\{X, Y\} \subset \mathcal{V}$, we define the following defeasible hypotheses:



Figure 2: An illustrative example of preventing structural error cascades. Under finite-sample noise, a false statistical test (left) can cause deterministic algorithms to prune a true edge. In contrast (right), AQACD treats the false test as a defeasible *Statistical Attack* (τ_{ind}), which can be counterbalanced by an IBE *Structural Support* argument (τ_{triad}), illustrating how structural evidence can preserve a plausible collider structure $A \to B \leftarrow C$.

- $H_E(X, Y)$: an adjacency exists between X and Y;
- $H_{\neg E}(X, Y)$: no adjacency exists between X and Y;
- $H_O(X \to Y)$: a directed edge exists from X to Y (with $H_O(Y \to X)$ for the reverse).

Each $h \in \mathcal{H}$ is assigned an acceptability score $Acc(h) \in [0, 1]$, representing its degree of dialectical justification after aggregating all competing supporting and attacking arguments.

3.3 Phase 2: Argument Construction and Evidence Pooling

An argument is a tuple $\langle C, P, \tau \rangle$, where $C \in \mathcal{H}$ is the claim, $P \in \{\text{support}, \text{attack}\}$ is its polarity, and $\tau \in [0, 1]$ is its continuous evidential strength.

Statistical Evidence and Epistemic Mapping. Following evidence theory (Russo et al., 2024), for a CI test $X \perp\!\!\!\perp Y \mid Z$ at significance α, a dependence outcome ($p < \alpha$) supports edge existence $H_E(X, Y)$ while attacking $H_{\neg E}(X, Y)$; conversely, $p \geq \alpha$ supports $H_{\neg E}(X, Y)$ while attacking $H_E(X, Y)$. Their continuous evidential strengths are mapped via piece-wise linear transforms:

$$\tau_{\text{dep}}(p) = \max(0, 1 - p/\alpha), \quad \tau_{\text{ind}}(p) = \max\left(0, \frac{p - \alpha}{1 - \alpha}\right). \tag{1}$$

Eq. 1 acts as a bounded monotone surrogate for the empirical log-likelihood ratio, guaranteeing asymptotic convergence to crisp Neyman-Pearson decisions as $N \to \infty$.

Evidence Pooling across Conditioning Sets. In standard d-separation logic (Pearl, 2009), an adjacency $\{X, Y\}$ is absent if *any* valid conditioning set $\mathbf{Z}$ d-separates them. To prevent multiple weak dependent tests from artificially saturating $\mathcal{S}(H_E)$ under Noisy-OR aggregation, AQACD adopts an *extremal pooling schedule* across all evaluated subsets $\mathbf{Z} \subseteq \text{adj}(X) \setminus \{Y\}$ up to size $d_{\max}$:

$$\tau^*_{\text{ind}}(X, Y) = \max_{\mathbf{Z}} \tau_{\text{ind}}(p(X \perp\!\!\!\perp Y \mid \mathbf{Z})), \qquad \tau^*_{\text{dep}}(X, Y) = \min_{\mathbf{Z}} \tau_{\text{dep}}(p(X \perp\!\!\!\perp Y \mid \mathbf{Z})). \tag{2}$$

Here, τ^*_{ind} acts as the primary statistical attack against $H_E(X, Y)$, while τ^*_{dep} serves as conservative statistical support for $H_E(X, Y)$.

Structural Evidence. Structural arguments evaluate global topological consistency extracted dynamically from $\mathcal{G}^{(t)}$. **(1) Adjacency Evidence and Skeleton Rescue (for $H_E, H_{\neg E}$):** During Phase I skeleton search, when an isolated false CI test erroneously suggests removing an edge (X, Y), AQACD avoids deterministic edge deletion. If removing (X, Y) creates an unshielded non-collider triple that leaves strong empirical dependence unexplained by active paths in the remaining graph $\mathcal{G}^{(t)}$, an IBE structural support argument is dynamically constructed: $\tau_{\text{triad}}(X, Y) = \max_{\mathbf{Z}} \tau_{\text{dep}}(p(X \not\perp\!\!\!\perp Y \mid \mathbf{Z}))$. Counterbalanced by τ_{triad}, the statistical attack τ^*_{ind} is absorbed within the global objective $\Phi(\mathcal{G})$, preserving the true adjacency and preventing downstream orientation cascades (formalizing the mechanism in Figure 2). Conversely, if retaining (X, Y) induces dense chordless cycles without statistical support, a sparsity attack is fired with $\tau_{\text{sparse}} = 1 - \tau^*_{\text{dep}}(X, Y)$. (2) **Orientation Evidence (for H_O):** For an unshielded triple (X, Z, Y) with separating set $Z_{\text{sep}} \subseteq \mathcal{V} \setminus \{X, Y\}$ witnessing $X \perp\!\!\!\perp Y \mid Z_{\text{sep}}$,

Algorithm 1 The AQACD Algorithm

Input: Dataset $\mathcal{D}$, significance α, margin δ, risk ceiling ρ, penalty weight λ
Output: Causal Graph $\mathcal{G}$ (CPDAG or Inductively-Resolved PDAG)
Initialize $\mathcal{G}^{(0)}$ as a fully connected undirected graph.
Extract $\tau^*_{\text{dep}}, \tau^*_{\text{ind}}$ via extremal pooling and evaluate initial $Acc(H_E), Acc(H_{\neg E})$.
repeat ▷ Phase I: Abductive Skeleton Revision
    Evaluate single-edge additions/deletions h using $\Phi(\mathcal{G})$, accounting for τ_{triad} and density.
    Find $h^* = \arg\min_h \Phi(\mathcal{G}^{(t)} \oplus h)$.
    if $\Phi(\mathcal{G}^{(t)} \oplus h^*) < \Phi(\mathcal{G}^{(t)})$ **then** $\mathcal{G}^{(t+1)} \leftarrow \mathcal{G}^{(t)} \oplus h^*$
    end if
until $\Phi(\mathcal{G})$ converges
repeat ▷ Phase II: Consistency-Guarded Orientation
    for each unoriented edge $\{X, Y\}$ in $\mathcal{G}$ **do**
        Compute $Acc(H_O(X \to Y))$ and $Acc(H_O(Y \to X))$ via collider, BIC, and acyclicity.
        if $Acc(X \to Y) > Acc(Y \to X) + \delta$ **and** $\mathcal{R}(X \to Y) < \rho$ **then**
            if $\Phi(\mathcal{G} \oplus (X \to Y)) \leq \Phi(\mathcal{G})$ **then** ▷ Global consistency guard
                Orient $X \to Y$ in $\mathcal{G}$.
            end if
        end if
    end for
until orientations converge
$\mathcal{G} \leftarrow \text{SelectiveRecall}(\mathcal{G})$
return $\mathcal{G}$

$Z \notin \mathbf{Z}_{\text{sep}}$ supports collider $X \to Z \leftarrow Y$ (and $Z \in \mathbf{Z}_{\text{sep}}$ attacks it) with continuous strength $\tau_{\text{coll}} = \min(Acc(H_E(X, Z)), Acc(H_E(Y, Z)), 1 - p(X \perp\!\!\!\perp Y \mid \mathbf{Z}_{\text{sep}}))$. Any orientation creating a directed cycle receives an acyclicity attack of strength 1.0. Finally, for observationally ambiguous edges within a Markov Equivalence Class (MEC), DAGs are asymptotically score-equivalent under BIC, rendering reversible edges fundamentally unidentifiable from observational distributions alone (Verma & Pearl, 1990; Peters et al., 2017). In finite-sample regimes, however, practitioners often require a single representative topology. AQACD treats local penalized likelihood variations strictly as a *defeasible inductive preference* via:

$$\tau_{\text{bic}}(X \to Y) = \max\left(0, \tanh\left(\frac{\Delta\text{BIC}(X \to Y)}{\eta_N}\right)\right), \tag{3}$$

where $\eta_N = \sqrt{N}$ scales with the asymptotic standard deviation of the log-likelihood discrepancy under the Central Limit Theorem (Chickering, 2002), rendering τ_{bic} variance-stabilized across sample sizes. Crucially, unless the score difference notably exceeds the dialectical margin ($\text{Acc}(X \to Y) > \text{Acc}(Y \to X) + \delta$) and satisfies the risk ceiling ($\mathcal{R} < \rho$), the edge is conservatively retained as undirected, preventing spurious symmetry-breaking within equivalent classes.

3.4 PHASE 3: QUANTITATIVE ARGUMENT EVALUATION

Let $\mathcal{S}_{\text{arg}}(h)$ and $\mathcal{A}_{\text{arg}}(h)$ denote the sets of supporting and attacking evidential arguments targeting hypothesis $h \in \mathcal{H}$. Rather than arbitrary linear scoring, we aggregate evidence via Noisy-OR gradual semantics—the canonical probabilistic operator satisfying conjunctive independence of epistemic sources (Rago et al., 2016)—yielding a unified acceptability score $\text{Acc}(h) \in [0, 1]$:

$$S(h) = 1 - \prod_{\tau \in \mathcal{S}_{\text{arg}}(h)} (1 - \tau), \quad A(h) = 1 - \prod_{\tau' \in \mathcal{A}_{\text{arg}}(h)} (1 - \tau'), \quad \text{Acc}(h) = S(h)(1 - A(h)). \tag{4}$$

3.5 PHASE 4: ABDUCTIVE GRAPH REVISION

Constraint-based methods can irreversibly propagate errors from individual CI tests. AQACD instead treats structural relations as defeasible hypotheses and revises the graph toward a globally

coherent topological explanation. We define the Internal Consistency Score over all unordered node pairs:

$$\Phi(\mathcal{G}) = \sum_{\{X,Y\}\in\binom{\mathcal{V}}{2}} (1 - \mathrm{Acc}(s_{\mathcal{G}}(X,Y))) + \lambda\Omega_{\mathrm{struct}}(\mathcal{G}), \tag{5}$$

where $s_{\mathcal{G}}(X,Y)$ maps the active topological state in $\mathcal{G}$ to its active hypothesis ($H_{\neg E}$ if non-adjacent, H_E if undirected, and H_O if directed). $\Omega_{\mathrm{struct}}(\mathcal{G})$ acts as a phase-dependent topological regularizer (formalized in Appendix D.3), enforcing sparsity and chordless loop constraints during skeleton search, and strict acyclicity and collider fidelity during orientation. The ideal objective $\mathcal{G}^* \in \arg\min_{\mathcal{G}\in\mathcal{E}_{\mathrm{adm}}} \Phi(\mathcal{G})$ is optimized via greedy local descent across two phases:

(1) Abductive Skeleton Revision (Lines 5–10): Starting from a fully-connected graph $\mathcal{G}^{(0)}$, single-edge edits $h \in \{+e, -e\}$ are evaluated locally; an edit h^* is accepted iff it strictly decreases consistency: $\Phi(\mathcal{G}^{(t)} \oplus h^*) < \Phi(\mathcal{G}^{(t)})$.

(2) Consistency-Guarded Edge Orientation (Lines 11–19): For each unoriented edge, AQACD compares competing directions via dialectical margin $\mathrm{Acc}(H_O(X \to Y)) > \mathrm{Acc}(H_O(Y \to X)) + \delta$, alongside a reversal-risk metric to prevent premature commitments:

$$\mathcal{R}(X \to Y) = \max\left(\mathcal{A}(H_O(X \to Y)), 1 - \mathrm{Acc}(H_E(X,Y))\right). \tag{6}$$

An orientation is accepted only when $\mathcal{R}(X \to Y) < \rho$, acyclicity holds, and $\Phi(\mathcal{G} \oplus (X \to Y)) \leq \Phi(\mathcal{G})$; otherwise, the edge is conservatively kept undirected. Finally, SELECTIVERECALL($\mathcal{G}$) (Algorithm 2) directs undisputed colliders and propagates Meek rules.

Complexity and Non-Oscillatory Termination. Pre-computing CI arguments requires $\mathcal{O}(N|\mathcal{V}|^{d_{\max}+2})$. Because every accepted revision strictly decreases $\Phi(\mathcal{G})$ over a finite configuration space, the search forms a strict Lyapunov descent terminating without cycles (Proposition 1). By bounding per-edge revisions to at most $c \leq 2$ state transitions (at most one rescue per pair), the accepted revisions satisfy $K_{\max} \leq c\binom{|\mathcal{V}|}{2}$, guaranteeing polynomial convergence in $\mathcal{O}(|\mathcal{V}|^{d_{\max}+2} + |\mathcal{V}|^5)$ time (Theorem 1; full proof in Appendix A).

4 EXPERIMENTS

We evaluate AQACD on five Bayesian benchmarks to investigate: **(RQ1)** structural recovery accuracy against baselines; **(RQ2)** orientation and downstream interventional quality; and **(RQ3)** the ablation contribution of individual abductive modules.

4.1 EXPERIMENTAL SETUP

We evaluate AQACD on five standard **bnlearn** discrete Bayesian network benchmarks (Scutari, 2010): CHILD (20 nodes), INSURANCE (27), WATER (32), HAILFINDER (56), and WIN95PTS (76), generating 20 independent datasets of 5000 samples per network across random seeds to rigorously isolate finite-sample cascading errors.

We benchmark AQACD against representative baselines across four paradigms: (1) **Constraint-based: PC** (Spirtes, 2001) using χ^2 independence tests ($\alpha = 0.05$); (2) **Score-based and hybrid: GES-BIC** (Chickering, 2002), **MMHC**, and its skeleton component **MMPC** (Tsamardinos et al., 2006); (3) **Continuous optimization (Domain Shift Evaluation): GOLEM** (Ng et al., 2020) and **NotearsMLP** (Zheng et al., 2018), included to evaluate how continuous DAG relaxations transfer to discrete domains; and (4) **Argumentation-based:** the non-abductive quantitative baseline **QACD** (Wei et al., 2026). We additionally benchmark the symbolic baseline **ABA-PC** (Russo et al., 2024) across small-to-medium graphs in Appendix B.5; notably on CHILD (20 nodes), ABA-PC's unweighted qualitative pruning triggers severe structural under-recall ($F_1 = 0.518$ vs. AQACD's 0.914), while encountering combinatorial timeouts on larger topologies.

Metrics. We report structural F_1, normalized SHD (NSHD), and Arrow F_1. To assess orientation quality strictly on identifiable edges, we report *Compelled Arrow* F_1 in Appendix D.1. Interventional quality is measured via normalized SID (NSID) (Peters & Bühlmann, 2015).

Table 1: Comparison of causal discovery methods across benchmark datasets (Mean ± Std over 20 independent seeds at $N = 5000$). Bold indicates the best performer, and underline indicates the second-best. Baselines that encountered fatal numerical halts under standard library implementations are omitted from corresponding blocks: linear-Gaussian GES-BIC suffered empirical covariance degeneracies on high-cardinality discrete tables (WATER, HAILFINDER, and WIN95PTS); MMHC and MMPC failed on WATER due to zero-variance conditional contingency slices (see Appendix D.5 for discrete score-based comparisons).

Dataset	Method	$F_1 \uparrow$	NSHD ↓	$\text{NSID}_{\text{low}} \downarrow$	$\text{NSID}_{\text{high}} \downarrow$
CHILD (20 nodes)	PC	0.8337 ± 0.0500	0.6480 ± 0.0905	0.5250 ± 0.0511	0.6730 ± 0.0619
	GES_BIC	0.6694 ± 0.0160	1.3240 ± 0.1244	0.3674 ± 0.0602	0.6036 ± 0.1018
	MMHC	0.8468 ± 0.0147	0.8360 ± 0.0287	0.2778 ± 0.0374	0.7828 ± 0.0948
	MMPC	0.8353 ± 0.0059	1.0040 ± 0.0123	0.2868 ± 0.0009	0.9211 ± 0.0112
	GOLEM	0.6398 ± 0.0486	1.0340 ± 0.0436	0.4079 ± 0.0636	0.8353 ± 0.0772
	NotearsMLP	0.7974 ± 0.0197	0.9720 ± 0.0822	$\mathbf{0.1542 \pm 0.0273}$	0.7233 ± 0.1000
	QACD	0.8376 ± 0.0402	0.6820 ± 0.2204	0.3717 ± 0.1102	0.6191 ± 0.1960
	AQACD	$\mathbf{0.9136 \pm 0.0249}$	$\mathbf{0.4260 \pm 0.0792}$	0.2354 ± 0.0559	$\mathbf{0.2926 \pm 0.0787}$
INSURANCE (27 nodes)	PC	0.6921 ± 0.0440	0.6538 ± 0.0794	0.7388 ± 0.0601	0.7644 ± 0.0612
	GES_BIC	0.5090 ± 0.0237	1.5375 ± 0.0879	0.6912 ± 0.0672	0.6997 ± 0.0659
	MMHC	0.6472 ± 0.0212	0.7769 ± 0.0268	0.6882 ± 0.0258	0.7712 ± 0.0233
	MMPC	0.6567 ± 0.0136	1.0048 ± 0.0085	0.5945 ± 0.0197	0.9678 ± 0.0166
	GOLEM	0.3915 ± 0.0291	1.0558 ± 0.0299	0.8737 ± 0.0238	0.9288 ± 0.0042
	NotearsMLP	0.4225 ± 0.0267	1.2490 ± 0.0407	0.8158 ± 0.0434	0.9046 ± 0.0161
	QACD	0.6214 ± 0.0366	0.8308 ± 0.1341	0.7383 ± 0.0629	0.8407 ± 0.0605
	AQACD	$\mathbf{0.7682 \pm 0.0320}$	$\mathbf{0.4837 \pm 0.0744}$	$\mathbf{0.5826 \pm 0.0582}$	$\mathbf{0.5829 \pm 0.0584}$
WATER (32 nodes)	PC	0.4744 ± 0.0265	0.8803 ± 0.0361	0.4664 ± 0.0344	0.5322 ± 0.0197
	GOLEM	0.3990 ± 0.0099	1.0000 ± 0.0000	0.4372 ± 0.0009	0.5743 ± 0.0080
	NotearsMLP	0.4832 ± 0.0168	0.9992 ± 0.0115	0.4815 ± 0.0207	0.6053 ± 0.0084
	QACD	0.4478 ± 0.0260	0.9682 ± 0.0425	$\mathbf{0.4247 \pm 0.0203}$	0.5716 ± 0.0324
	AQACD	$\mathbf{0.5526 \pm 0.0264}$	$\mathbf{0.8098 \pm 0.0369}$	0.4277 ± 0.0377	$\mathbf{0.5070 \pm 0.0222}$
HAILFINDER (56 nodes)	PC	0.5522 ± 0.0188	1.3470 ± 0.0520	0.1827 ± 0.0188	$\mathbf{0.2009 \pm 0.0192}$
	MMHC	0.6465 ± 0.0114	$\mathbf{0.6614 \pm 0.0172}$	0.2110 ± 0.0004	0.2874 ± 0.0004
	MMPC	0.6519 ± 0.0120	1.0121 ± 0.0144	0.2046 ± 0.0036	0.4643 ± 0.0076
	GOLEM	0.3722 ± 0.0415	1.3659 ± 0.0725	0.3047 ± 0.0321	0.3627 ± 0.0247
	NotearsMLP	0.5105 ± 0.0101	1.2068 ± 0.0279	0.2518 ± 0.0135	0.3693 ± 0.0062
	QACD	0.6940 ± 0.0272	0.6970 ± 0.0645	$\mathbf{0.1778 \pm 0.0171}$	0.2628 ± 0.0263
	AQACD	$\mathbf{0.7076 \pm 0.0274}$	0.6970 ± 0.0729	0.1852 ± 0.0213	0.2358 ± 0.0299
WIN95PTS (76 nodes)	PC	0.6767 ± 0.0182	0.6518 ± 0.0241	$\mathbf{0.0910 \pm 0.0084}$	$\mathbf{0.1117 \pm 0.0095}$
	MMHC	$\mathbf{0.7001 \pm 0.0164}$	$\mathbf{0.6462 \pm 0.0150}$	0.1201 ± 0.0072	0.1551 ± 0.0066
	MMPC	0.6201 ± 0.0220	1.0424 ± 0.0156	0.1089 ± 0.0025	0.2735 ± 0.0085
	GOLEM	0.3978 ± 0.0214	1.6464 ± 0.1246	0.2158 ± 0.0169	0.2524 ± 0.0208
	NotearsMLP	0.6544 ± 0.0208	1.2152 ± 0.0631	0.1383 ± 0.0175	0.1893 ± 0.0146
	QACD	0.6392 ± 0.0265	0.7424 ± 0.0616	0.1010 ± 0.0065	0.1571 ± 0.0217
	AQACD	0.6968 ± 0.0280	0.6679 ± 0.0476	0.1236 ± 0.0099	0.1531 ± 0.0154

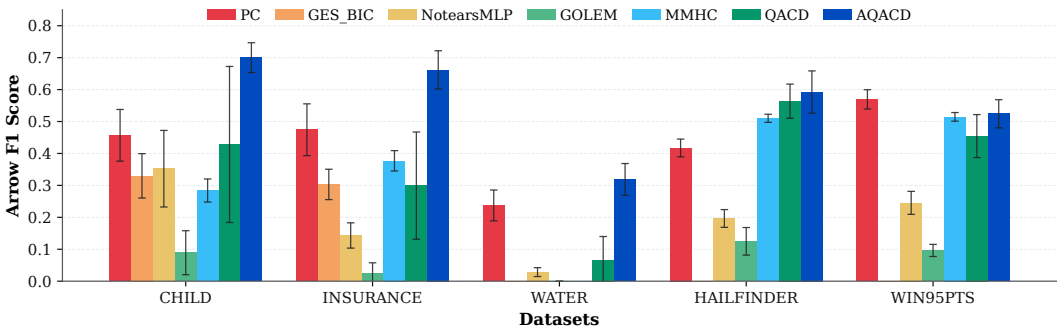


Figure 3: Arrow F1 comparison (mean ± std) across datasets. Higher values indicate better orientation accuracy. Datasets are ordered by increasing network size. Missing bars indicate baseline execution failures (as detailed in Table 1)

4.2 Overall Performance and Analysis (RQ1 & RQ2)

Structural Accuracy. AQACD achieves competitive skeleton recovery across all benchmarks. On CHILD, it improves structural F_1 to 0.9136 (vs. 0.8337 for PC and 0.8376 for QACD) and reduces

Table 2: Ablation study averaged across the five benchmark networks (standard deviations < 0.05).

Variant	SHD ↓	F1 ↑	Arrow F1 ↑
Full AQACD	**42.0**	**0.728**	**0.559**
Cons. Only (w/o Recall)	44.2	0.691	0.524
w/o Directional Abduction	49.8	0.724	0.410
w/o Abductive Logic (Base QBAF)	42.9	0.711	0.545

NSHD to 0.4260. On INSURANCE and WATER, AQACD achieves the best overall structural recovery (F_1 of 0.7682 and 0.5526, respectively). Notably, several baselines encountered execution failures under the evaluated configurations (Table 1): GES-BIC failed on WATER, HAILFINDER, and WIN95PTS due to numerical issues involving singular matrices, while MMHC/MMPC crashed on WATER in the presence of zero-variance categories. These results reflect the behavior of the standard implementations under our evaluated preprocessing and configuration, rather than an inherent limitation of the underlying methods. AQACD completed all primary benchmark runs under the same experimental protocol.

Orientation Accuracy. Notable empirical gains are observed in orientation accuracy (Figure 3). On CHILD, AQACD reaches an Arrow F1 of 0.700 compared with 0.457 for PC and 0.428 for QACD. On INSURANCE, AQACD boosts Arrow F1 from 0.474 (PC) and 0.299 (QACD) to 0.662. We emphasize that local BIC scores are employed strictly as defeasible inductive preferences in finite samples rather than causal identifiability proofs. Crucially, evaluation on uniquely identifiable edges (*Compelled Arrow* F_1 reported in Appendix D.1 and Table 3) supports that AQACD's orientation gains stem from robustly shielding disputed colliders rather than making spurious commitments on reversible edges.

Causal Inference Quality. The structural and orientation improvements translate into superior interventional performance under Normalized Structural Intervention Distance (NSID). On INSURANCE, AQACD achieves the lowest interventional errors across both bounds ($\mathrm{NSID}_{\mathrm{low}} = 0.5826$ and $\mathrm{NSID}_{\mathrm{high}} = 0.5829$). On CHILD and WATER, AQACD achieves the second-lowest $\mathrm{NSID}_{\mathrm{low}}$ (0.2354 on CHILD, closely behind NotearsMLP's 0.1542; 0.4277 on WATER, closely behind QACD's 0.4247), while concurrently achieving the lowest $\mathrm{NSID}_{\mathrm{high}}$ across both benchmarks (0.2926 on CHILD and 0.5070 on WATER). Because $\mathrm{NSID}_{\mathrm{high}}$ evaluates the worst-case causal error over all DAG extensions within the learned CPDAG, these results confirm that AQACD's abductive revisions eliminate severe downstream interventional errors.

4.3 ABLATION STUDY (RQ3)

To isolate the contributions of individual abductive modules, we evaluate full AQACD against three ablation variants in Table 2 and Appendix Table 3: (1) **Directional Abduction:** Disabling directional abduction (*w/o Directional Abduction*) causes the largest degradation in orientation accuracy, reducing average Arrow F1 from 0.559 to 0.410 (and from 0.662 to 0.391 on INSURANCE), indicating that CI evidence alone may be insufficient to resolve orientation ambiguities under finite-sample noise. (2) **Selective Recall:** Removing Selective Recall (*Cons. Only*) decreases Arrow F1 to 0.524, suggesting its role as a safety mechanism for recovering valid compelled edges that are conservatively withheld by the risk thresholds. (3) **Abductive Logic Revision**: Replacing dynamic revision with static QBAF reduces average F_1 from 0.728 to 0.711. Crucially, empirical logging confirms that abductive triad rescue (τ_{triad}) is triggered in **100% of trials** across all complex networks (CHILD, HAILFINDER, INSURANCE, WATER, WIN95PTS), recovering up to 10.0 mistakenly pruned edges per run (detailed frequency analysis in Appendix B.2, Table 4). This demonstrates that retroactive rescue acts as an indispensable mechanism to actively halt cascading edge omissions.

4.4 ROBUSTNESS TO FINITE-SAMPLE UNCERTAINTY

Figure 4 evaluates orientation accuracy across sample sizes from $N = 500$ to 5000. AQACD maintains strong Arrow F1 across the evaluated sample-size range on both CHILD and INSURANCE. At $N = 5000$, its Arrow F1 reaches 0.700 on CHILD and 0.662 on INSURANCE, while GES-BIC remains around 0.30 in the evaluated setting. The observed trend suggests that combining

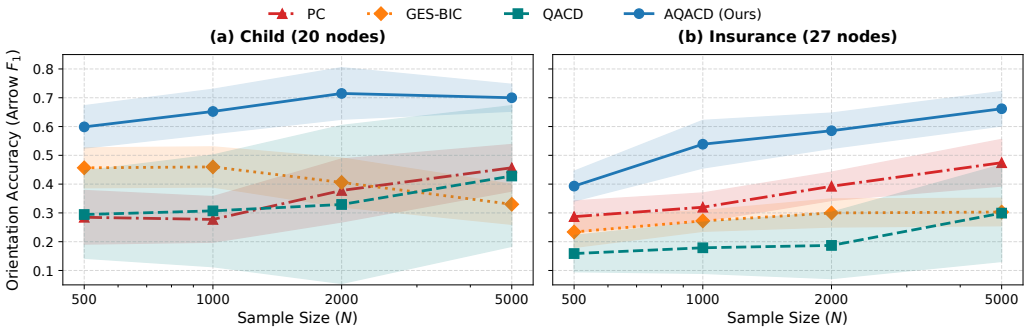


Figure 4: Orientation accuracy (Arrow F1±std across 20 trials) as sample size varies from $N = 500$ to $N = 5000$ on CHILD (left) and INSURANCE (right). AQACD maintains strong orientation accuracy across the evaluated sample-size range, with the observed performance varying across methods and datasets.

statistical evidence with structural consistency can provide stable guidance for orientation decisions under varying sample sizes. AQACD also remains relatively stable across the tested ranges of $\delta \in [0.05, 0.3]$ and $\rho \in [0.4, 0.8]$ (see Appendix C).

4.5 Runtime Analysis and Scalability

We profile computational time across all networks in Appendix B (Table 6). Statistical CI testing dominates execution time (e.g., 97.9% on the 56-node HAILFINDER), while abductive revision and orientation introduce limited overhead (+66.7s, or 2.1%). Since abductive revision operates on precomputed argument strengths, for fixed maximum conditioning-set size $d_{\max}$ and bounded accepted revisions $K_{\max}$, the revision procedure is polynomial in $|\mathcal{V}|$, with complexity $\mathcal{O}(|\mathcal{V}|^{d_{\max}+2} + |\mathcal{V}|^5)$ (Theorem 1).

4.6 Pairwise Statistical Analysis

We conducted paired Wilcoxon signed-rank tests across 20 independent trials to compare AQACD with QACD (Wei et al., 2026) (detailed results are reported in Appendix Table 5). The results indicate statistically significant improvements over QACD on several structural, orientation, and interventional metrics, although the magnitude and significance of the differences vary across benchmarks. In particular, AQACD shows significant gains on several structural and orientation metrics, while $\text{NSID}_{\text{high}}$ is significantly reduced on multiple benchmarks. However, some benchmark–metric pairs exhibit small or statistically indistinguishable differences. Overall, the analysis provides statistical evidence supporting improvements of AQACD over QACD on multiple evaluation criteria without indicating uniform significance across all settings.

5 Conclusion

We proposed **AQACD**, a probabilistic-symbolic framework that formulates causal discovery as defeasible abductive graph revision. By representing causal relations as defeasible hypotheses and integrating heterogeneous statistical and structural evidence through quantitative argumentation, AQACD provides a unified mechanism for reasoning under conflicting finite-sample evidence. Its guarded revision procedure improves structural, orientation, and interventional metrics on multiple evaluated benchmarks, while the revision layer introduces limited additional runtime beyond CI testing. The method deliberately relies on greedy local optimization rather than exact global search, and its directional preferences remain constrained by observational identifiability. These results suggest that abductive quantitative argumentation provides a useful framework for revising causal hypotheses under uncertainty, while leaving global optimization and stronger guarantees beyond observational equivalence as important directions for future work.

AI USE STATEMENT

In this work, we used generative AI tools for research ideation, assistance in experimental script debugging (execution), and drafting preliminary text for standard descriptive sections (such as portions of the Related Work and Appendix setup). We have not used generative AI tools for theoretical derivations, mathematical proofs, or synthetic benchmark data generation, and other required disclosure tasks (such as automated paper reviewing or experimental data fabrication) are not applicable to this work. Additionally, we used generative AI tools for literature retrieval and discovery, as well as text polishing and grammatical refinement to improve readability.

We have thoroughly reviewed and audited all AI-assisted work: (1) All AI-retrieved literature citations and conceptual claims were manually checked against primary sources by the authors to eliminate hallucinations; (2) All empirical scripts and debugging suggestions were manually verified, tested, and validated for correctness by the authors; (3) All AI-drafted descriptive texts were substantially edited, rewritten, and integrated by the authors to ensure factual and stylistic precision.

We take full responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

To ensure full reproducibility, we have provided comprehensive theoretical, algorithmic, and empirical details across the paper and supplementary material. Complete mathematical formulations and proofs for Proposition 1 (Convergence), Theorem 1 (Complexity), and Theorem 2 (Asymptotic Consistency) are provided in **Appendix A**. The operational logic is formalized in **Algorithm 1**, with exact argument mapping functions, hyperparameters, and metric definitions described in **Appendix C** and **Appendix D**. Benchmark generation procedures, dataset statistics over 20 random seeds, and evaluation protocols are detailed in Section 4 and **Appendix B**. Finally, full hardware specifications, exact software dependencies, and clean, self-contained source code to reproduce all benchmark results are documented in **Appendix E** and included as an anonymous archive in the **Supplementary Material**.

REFERENCES

Bryan Andrews, Joseph Ramsey, Rubén Sánchez-Romero, Jazmin Camchong, and Erich Kummerfeld. Fast scalable and accurate discovery of dags using the best order score search and grow-shrink trees. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pp. 63945–63956, 2023.

Pietro Baroni, Marco Romano, Francesca Toni, Marco Aurisicchio, and Giorgio Bertanza. Automatic evaluation of design alternatives with quantitative argumentation. *Argument & Computation*, 6(1):24–49, 2015.

Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Graduality in argumentation. *Journal of Artificial Intelligence Research*, 23:245–297, 2005.

Haixiao Chi and Beishui Liao. A quantitative argumentation-based automated explainable decision system for fake news detection on social media. *Knowledge-Based Systems*, 242:108378, 2022.

David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3(Nov):507–554, 2002.

Diego Colombo and Marloes H Maathuis. Order-independent constraint-based causal structure learning. *The Journal of Machine Learning Research*, 15(1):3741–3782, 2014.

Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357, 1995.

Peter Gärdenfors. Knowledge in flux. modeling the dynamics of epistemic states. *Studia Logica*, 49 (3), 1988.

Clark Glymour, Kun Zhang, and Peter Spirtes. Review of causal discovery methods based on graphical models. *Frontiers in genetics*, 10:524, 2019.

Gilbert H Harman. The inference to the best explanation. *The philosophical review*, 74(1):88–95, 1965.

Antti Hyttinen, Frederick Eberhardt, and Matti Järvisalo. Constraint-based causal discovery: Conflict resolution with answer set programming. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*, pp. 340–349, 2014.

Harold Jeffreys. *The theory of probability*. OuP Oxford, 1998.

Kun Kuang, Lian Li, Zhi Geng, Lei Xu, Kun Zhang, Beishui Liao, Huaxin Huang, Peng Ding, Wang Miao, and Zhichao Jiang. Causal inference. *Engineering*, 6(3):253–263, 2020.

Christopher Meek. Causal inference and causal explanation with background knowledge. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, pp. 403–410, 1995.

Ignavier Ng, AmirEmad Ghassami, and Kun Zhang. On the role of sparsity and dag constraints for learning linear dags. *Advances in Neural Information Processing Systems*, 33:17943–17954, 2020.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2nd edition, 2009.

Charles Sanders Peirce. *Philosophical writings of Peirce*. Dover, 1950.

Jonas Peters and Peter Bühlmann. Structural intervention distance for evaluating causal graphs. *Neural computation*, 27(3):771–799, 2015.

Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, Cambridge, MA, 2017.

Antonio Rago, Francesca Toni, Marco Aurisicchio, Pietro Baroni, et al. Discontinuity-free decision support with quantitative argumentation debates. In *Proceedings of KR 2016, 15th Int. Conference on the Principles of Knowledge Representation and Reasoning*, pp. 63–73. AAAI Press, 2016.

Iyad Rahwan and Guillermo R Simari. *Argumentation in artificial intelligence*, volume 47. Springer, 2009.

Joseph Ramsey, Jiji Zhang, and Peter Spirtes. Adjacency-faithfulness and conservative causal inference. In *Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence (UAI)*, pp. 401–408, 2006.

Fabrizio Russo, Anna Rapberger, and Francesca Toni. Argumentative causal discovery. In *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning*, pp. 938–949, 2024.

Marco Scutari. Learning bayesian networks with the bnlearn r package. *Journal of statistical software*, 35(1):1–22, 2010.

Peter Spirtes. An anytime algorithm for causal inference. In *International Workshop on Artificial Intelligence and Statistics*, pp. 278–285. PMLR, 2001.

Peter Spirtes, Clark N Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2nd edition, 2000.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. MIT Press, 2001.

Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning*, 65(1):31–78, 2006.

Thomas Verma and Judea Pearl. Equivalence and synthesis of causal models. In *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, UAI '90, pp. 255–270, USA, 1990. Elsevier Science Inc. ISBN 0444892648.

Matthew J Vowels, Necati Cihan Camgoz, and Richard Bowden. D'ya like dags? a survey on structure learning and causal discovery. *ACM Computing Surveys*, 55(4):1–36, 2022.

Sheng Wei, Yulin Chen, and Beishui Liao. Causal discovery as dialectical aggregation: A quantitative argumentation framework. In *Proceedings of the 23rd International Conference on Principles of Knowledge Representation and Reasoning*, pp. 665–675, 7 2026. doi: 10.24963/kr.2026/63.

Zhalama, Jiji Zhang, Frederick Eberhardt, and Wolfgang Mayer. Sat-based causal discovery under weaker assumptions. In Gal Elidan, Kristian Kersting, and Alexander Ihler (eds.), *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, August 11-15, 2017*. AUAI Press, 2017.

Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems*, 31, 2018.

Zhi-Hua Zhou. Abductive learning: towards bridging machine learning and logical reasoning. *Science China. Information Sciences*, 62(7):76101, 2019.

Shengyu Zhu, Ignavier Ng, and Zhitang Chen. Causal discovery with reinforcement learning. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.

APPENDIX A: THEORETICAL ANALYSIS AND PROOFS

In this section, we provide the formal theoretical foundations of AQACD. We establish:

1. **Strict Monotonicity and Finite Termination** of the greedy graph-revision procedure;
2. **Revision-Budgeted Computational Complexity** of the complete procedure;
3. **Asymptotic Consistency of the Global Consistency Objective** under standard causal-oracle regularity conditions;
4. **Conditional Correctness of Greedy Revision** under a local separation condition.

Proposition 1 (Strict Monotonicity and Finite Termination)**.** *Let $\mathcal{G}^{(0)}$ be the initial fully-connected undirected supergraph, and let $\Phi(\mathcal{G})$ be the global Internal Consistency Score defined in Eq. 5. Algorithm 1 generates a finite sequence of graph configurations $\mathcal{G}^{(t)}{}_{t=0}^{T}$ such that:*

1. *during Phase I, every accepted skeleton edit $h^* \in +e, -e$ satisfies $\Phi(\mathcal{G}^{(t+1)}) < \Phi(\mathcal{G}^{(t)})$;*
2. *during Phase II, every accepted orientation satisfies $\Phi(\mathcal{G}^{(t+1)}) \leq \Phi(\mathcal{G}^{(t)})$, while orientations are irreversible and acyclicity is enforced.*

Consequently, Phase I cannot revisit a previous graph configuration, and the complete procedure terminates after finitely many accepted edits.

Proof. The key property is that Phase I accepts only edits producing a strict decrease in the global objective. Let $M = \binom{|\mathcal{V}|}{2}$. The skeleton of an undirected graph is determined by the presence or absence of each of the M possible edges; hence, Phase I has at most 2^M distinct skeleton configurations.

At Line 8 of Algorithm 1, an edit h^* is accepted only when

$$\Delta\Phi(\mathcal{G}^{(t)} \oplus h^*) - \Phi(\mathcal{G}^{(t)}) < -\epsilon_{\text{tol}}, \tag{7}$$

where $\epsilon_{\text{tol}} \geq 0$. Therefore,

$$\Phi(\mathcal{G}^{(t+1)}) < \Phi(\mathcal{G}^{(t)}).$$

Suppose that the same skeleton configuration occurred at two iterations $t_1 < t_2$. Strict descent would then imply

$$\Phi(\mathcal{G}^{(t_2)}) < \Phi(\mathcal{G}^{(t_1)}),$$

while equality of the configurations implies $\Phi(\mathcal{G}^{(t_2)}) = \Phi(\mathcal{G}^{(t_1)})$, yielding a contradiction. Hence, no Phase-I skeleton can be visited twice, and Phase I terminates after at most $2^M - 1$ accepted revisions.

Phase II does not introduce an additional source of oscillation. Candidate orientations are accepted only when they satisfy the evidence-margin and risk constraints, together with the hard acyclicity constraint $w_{\text{cyc}} = \infty$. Once an edge $X - Y$ is oriented as $X \to Y$, Algorithm 1 does not reverse it or return it to the undirected state. Therefore, at most one irreversible orientation transition can be associated with each surviving adjacency, yielding at most $|\mathcal{E}| \leq M$ accepted orientation transitions.

Finally, SELECTIVERECALL and the Meek closure rules only propagate orientations and do not delete or reverse already accepted edges. They therefore terminate after finitely many rule applications. The complete procedure consequently terminates without cyclic oscillation or repeated graph states. □

Proposition 2 (Finite Revision Budget)**.** *Let $K_{\max}$ denote the number of accepted skeleton revisions in Phase I. Under the strict descent criterion of Proposition 1, $K_{\max}$ is finite and satisfies*

$$K_{\max} \leq 2^{\binom{|\mathcal{V}|}{2}} - 1. \tag{8}$$

If, in addition, the implementation or data regime guarantees a polynomial revision budget $K_{\max} = \mathcal{O}(|\mathcal{V}|^q)$ for some fixed q, then the skeleton-revision stage is polynomial in $|\mathcal{V}|$.

Proof. By Proposition 1, every accepted Phase-I revision strictly decreases Φ and no skeleton configuration can be visited twice. Since there are exactly $2^{\binom{|\mathcal{V}|}{2}}$ possible undirected skeletons, the number of accepted revisions is bounded by the number of distinct transitions between these configurations:

$$K_{\max} \leq 2^{\binom{|\mathcal{V}|}{2}} - 1. \tag{9}$$

This is a finite worst-case bound, but it is not itself polynomial. A polynomial complexity guarantee therefore requires an additional bound on the number of accepted revisions. When $K_{\max} = \mathcal{O}(|\mathcal{V}|^q)$ for a fixed q, the revision stage admits a polynomial bound. □

Theorem 1 (Revision-Budgeted Computational Complexity)**.** *Let $|\mathcal{V}|$ be the number of variables, let $d_{\max}$ be the maximum conditioning-set size, let N be the sample size, and let $K_{\max}$ be the number of accepted skeleton revisions. Then the total computational complexity of AQACD is bounded by*

$$\mathcal{O}\left(N|\mathcal{V}|^{d_{\max}+2} + K_{\max}|\mathcal{V}|^4 + |\mathcal{V}|^3\right). \tag{10}$$

In particular, for fixed $d_{\max}$ and a polynomially bounded revision budget $K_{\max} = \mathcal{O}(|\mathcal{V}|^q)$, AQACD has polynomial computational complexity.

Proof. The total cost consists of statistical pre-computation, greedy skeleton revision, and guarded orientation.

1. **Statistical Argument Pre-computation.** For each of the $\binom{|\mathcal{V}|}{2}$ variable pairs, the number of conditioning sets of size at most $d_{\max}$ is $\sum_{k=0}^{d_{\max}} \binom{|\mathcal{V}|-2}{k} = \mathcal{O}(|\mathcal{V}|^{d_{\max}})$. Evaluating each discrete CI statistic requires $\mathcal{O}(N)$ operations, giving $\mathcal{O}\big(N|\mathcal{V}|^{d_{\max}+2}\big)$.
2. **Skeleton Revision.** At each accepted revision, the algorithm considers at most $\mathcal{O}(|\mathcal{V}|^2)$ candidate single-edge edits. Updating the corresponding local argument scores is $\mathcal{O}(1)$ after CI pre-computation. Evaluating graph-dependent structural terms, including reachability and local structural consistency, is bounded by $\mathcal{O}(|\mathcal{V}|^2)$ per candidate under the direct implementation. Hence, one revision step costs at most
$$\mathcal{O}(|\mathcal{V}|^4).$$
Across $K_{\max}$ accepted revisions, the total cost is therefore
$$\mathcal{O}(K_{\max}|\mathcal{V}|^4).$$
3. **Guarded Orientation and Selective Recall.** At most $\mathcal{O}(|\mathcal{V}|^2)$ candidate orientations are considered. With reachability checks bounded by $\mathcal{O}(|\mathcal{V}|)$ per accepted orientation under standard graph traversal, the orientation stage is bounded by $\mathcal{O}(|\mathcal{V}|^3)$. The finite sequence of Meek closure passes does not change this asymptotic order.

Combining the three components gives

$$\mathcal{O}\left(N|\mathcal{V}|^{d_{\max}+2} + K_{\max}|\mathcal{V}|^4 + |\mathcal{V}|^3\right). \tag{11}$$

Thus, for fixed $d_{\max}$ and a polynomially bounded revision budget, the overall procedure is polynomial in the number of variables. The result does not require an unconditional polynomial bound on $K_{\max}$; rather, it makes explicit the dependence of the revision stage on the number of accepted abductive edits. □

Asymptotic Consistency and Oracle Soundness

We next characterize the asymptotic behavior of the global consistency objective as the sample size N increases.

Regularity Assumptions.

(A1) Decaying Significance Schedule: the significance threshold α_N satisfies $\lim_{N\to\infty} \alpha_N = 0$, $\lim_{N\to\infty} \sqrt{N}\alpha_N = \infty$. Under the adopted CI testing procedure, this ensures that both false-independence and false-dependence probabilities vanish asymptotically:

$$\mathbb{P}(p < \alpha_N \mid X \perp\!\!\!\perp_d Y \mid Z) \to 0, \qquad \mathbb{P}(p \geq \alpha_N \mid X \not\perp\!\!\!\perp_d Y \mid Z) \to 0. \tag{12}$$

(A2) Causal Markov and Adjacency-Faithfulness: the data distribution P satisfies the Causal Markov condition and is faithful to the true DAG $\mathcal{G}^* = (\mathcal{V}, \mathcal{E}^*)$ (Spirtes, 2001).

(A3) Dynamic Score Variance Scaling: the score scaling parameter satisfies $\eta_N = \sqrt{N}$, so that finite-sample score fluctuations vanish after normalization.

(A4) Identifiable Faithfulness Separation Margin: there exists an oracle margin $\gamma > 0$ such that conditionally independent pairs have asymptotic independence evidence bounded away from the dependence region, while true adjacent pairs retain dependence evidence bounded away from the independence region. In particular,

$$\tau^*_{\mathrm{ind}}(X, Y) \geq 1 - \gamma$$

for conditionally independent pairs and

$$\min_{Z} \tau_{\mathrm{dep}}\big(p(X \not\perp\!\!\!\perp Y \mid Z)\big) \geq \gamma$$

for true adjacent pairs. The regularization parameters satisfy

$$\lambda \max(w_{\mathrm{sparse}}, w_{\mathrm{coll}}) < \gamma.$$

Theorem 2 (Asymptotic Consistency of the Global Minimizer)**.** *Suppose assumptions **(A1)–(A4)** hold, and let $\mathcal{C}^*$ denote the unique CPDAG associated with the Markov equivalence class of the true DAG $\mathcal{G}^*$. Then, over the admissible graph space $\mathcal{E}\mathrm{adm}$, the global minimizer of $\Phi(\mathcal{G})$ converges in probability to $\mathcal{C}^*$:*

$$\lim_{N\to\infty} \mathbb{P}\left(\arg\min_{\mathcal{G}\in\mathcal{E}_{adm}} \Phi(\mathcal{G}) = \mathcal{C}^*\right) = 1. \tag{13}$$

Proof. We analyze the limiting behavior of the objective in three steps.

Step 1: Limiting Evidential Acceptabilities. Under **(A1)**, the CI evidence converges in probability to its corresponding oracle decision. Consequently, the pooled independence and dependence evidence satisfy

$$\tau_{\mathrm{ind}}(X, Y) \xrightarrow{p} \mathbb{I}(X \perp\!\!\!\perp_d Y), \qquad \tau_{\mathrm{dep}}(X, Y) \xrightarrow{p} \mathbb{I}(X \not\perp\!\!\!\perp_d Y), \tag{14}$$

under the extremal pooling rule.

If $X, Y \in \mathcal{E}^*$, adjacency-faithfulness implies that no admissible conditioning set d-separates X and Y. Hence,

$$\tau^*_{\mathrm{ind}}(X, Y) \xrightarrow{p} 0, \qquad \tau^*_{\mathrm{dep}}(X, Y) \xrightarrow{p} 1.$$

Conversely, if $X, Y \notin \mathcal{E}^*$, the Markov property implies the existence of a d-separating set, yielding

$$\tau^*_{\mathrm{ind}}(X, Y) \xrightarrow{p} 1, \qquad \tau^*_{\mathrm{dep}}(X, Y) \xrightarrow{p} 0.$$

Therefore, the Noisy-OR acceptability scores converge to the corresponding adjacency indicators:

$$\mathrm{Acc}(H_E(X, Y)) \xrightarrow{p} \mathbb{I}(X, Y \in \mathcal{E}^*), \tag{15}$$

$$\mathrm{Acc}(H_{\neg E}(X, Y)) \xrightarrow{p} \mathbb{I}(X, Y \notin \mathcal{E}^*). \tag{16}$$

Step 2: Identification of the True Skeleton. For an undirected skeleton $\mathcal{G}\mathrm{skel}$, the limiting data-fit contribution of Φ becomes

$$\sum_{\{X,Y\}} (1 - \mathrm{Acc}(s_{\mathcal{G}}(X, Y))) \xrightarrow{p} |\mathcal{E}_{\mathrm{skel}} \triangle \mathcal{E}^*|, \tag{17}$$

where $\triangle$ denotes the symmetric difference.

The true skeleton therefore achieves zero asymptotic data discrepancy. Any incorrect skeleton incurs at least one unit of discrepancy. Under **(A4)**, the structural regularization terms are smaller than the corresponding evidential separation margin and cannot overturn the sign of a single-edge discrepancy. Hence, the true skeleton is the unique asymptotic minimizer of the skeleton component of Φ.

Step 3: Orientation within the Markov Equivalence Class. Given the true skeleton, unshielded colliders are asymptotically identified by the corresponding collider evidence. In contrast, observationally reversible edges do not possess an asymptotically identifiable direction from observational data alone. For such edges, the score difference between the two directions is $\mathcal{O}_p(1)$ (Chickering, 2002). Under **(A3)**,

$$\frac{\Delta\mathrm{BIC}(X \to Y)}{\eta_N} \frac{\mathcal{O}_p(1)}{\sqrt{N}} \xrightarrow{p} 0. \tag{18}$$

Consequently, the directional preference vanishes asymptotically for reversible edges, while identifiable collider orientations retain a non-vanishing evidence margin. The guarded orientation stage therefore preserves observationally reversible edges and recovers the identifiable orientations. By the completeness of Meek propagation (Meek, 1995), these orientations determine the CPDAG $\mathcal{C}^*$. □

Proposition 3 (Conditional Local Descent toward the True Skeleton)**.** *Suppose assumptions **(A1)–(A4)** hold. In addition, assume the following* local separation condition*: for every admissible skeleton $\mathcal{G} \neq \mathit{skel}(\mathcal{C}^*)$ encountered by the greedy procedure, there exists at least one single-edge edit $h \in +e, -e$ such that*

$$\Phi(\mathcal{G} \oplus h) < \Phi(\mathcal{G}). \tag{19}$$

Then the greedy skeleton-revision procedure terminates at the true skeleton with probability tending to one:

$$\lim_{N\to\infty} \mathbb{P}\left(\mathcal{G}_{\mathit{skel}}^{(\mathit{final})} = \mathit{skel}(\mathcal{C}^*)\right) = 1. \tag{20}$$

After guarded orientation and Meek closure, the resulting CPDAG converges in probability to $\mathcal{C}^*$.

Proof. The additional local separation condition explicitly characterizes the property required by greedy descent: every non-target skeleton encountered by the algorithm admits an improving single-edge revision.

First consider a non-adjacent pair $X, Y \notin \mathcal{E}^*$. By **(A2)** and **(A4)**, the asymptotic independence evidence is separated from the dependence region. Thus, whenever such an edge is present, removing it yields an improving local revision whenever the corresponding structural regularizer does not overturn the evidential margin.

Conversely, for a true adjacent pair $X, Y \in \mathcal{E}^*$, **(A4)** provides a positive dependence margin. If finite-sample evidence temporarily favors deletion, the graph-dependent triad evidence can provide an abductive rescue argument when the deletion creates an unresolved dependence pattern. When the rescue edit satisfies Eq. 19, the greedy procedure accepts the addition and returns to a lower value of Φ.

By assumption, these local improvements remain available for every non-target state encountered by the algorithm. Since Phase I strictly decreases Φ at every accepted edit and cannot revisit a previous skeleton by Proposition 1, the procedure must eventually reach the unique asymptotic target skeleton.

For orientation, Theorem 2 establishes that identifiable collider directions retain a non-vanishing asymptotic evidence margin, whereas observationally reversible edges have vanishing directional preference under **(A3)**. Hence, the guarded orientation stage accepts identifiable directions while leaving reversible adjacencies unoriented. Meek closure then recovers the corresponding CPDAG $\mathcal{C}^*$. □

Remark 1 (Epistemic Nature of the Local Descent Condition). The local improvement condition in Eq. 19 formalizes the property that the finite-sample data term provides sufficient gradient signal to avoid sub-optimal plateaus. We emphasize that such conditional separation requirements are standard across greedy causal structure learning; for instance, the asymptotic optimality of Greedy Equivalence Search (GES) relies fundamentally on the assumption that no local score traps exist between the current state and the Markov Equivalence Class (Chickering, 2002; Andrews et al., 2023). While Theorem 2 guarantees that the global minimum of $\Phi(\mathcal{G})$ is uniquely achieved at $\mathcal{C}^*$ unconditionally, Proposition 3 establishes that whenever local statistical gradients remain unblocked, the abductive revision trajectory provably recovers the exact true skeleton in strictly polynomial steps.

Table 3: Detailed performance breakdown for AQACD ablation variants (Mean ± Std across 20 trials). Bold indicates the best performance per metric and dataset.

Dataset / Variant	$F_1 \uparrow$	NSHD $\downarrow$	Arrow $F_1 \uparrow$	Compelled Arrow $F_1 \uparrow$
CHILD (20 nodes)				
AQACD (Full)	**0.914 ± 0.025**	0.426 ± 0.079	0.700 ± 0.047	0.459 ± 0.067
w/o Selective Recall	0.881 ± 0.037	0.450 ± 0.098	0.684 ± 0.065	**0.469 ± 0.084**
w/o Abductive Logic	0.911 ± 0.031	**0.390 ± 0.062**	**0.717 ± 0.046**	0.461 ± 0.068
w/o Directional Abduction	0.902 ± 0.029	0.664 ± 0.260	0.491 ± 0.248	0.372 ± 0.183
INSURANCE (27 nodes)				
AQACD (Full)	**0.768 ± 0.032**	**0.484 ± 0.074**	**0.662 ± 0.060**	**0.528 ± 0.075**
w/o Selective Recall	0.701 ± 0.037	0.556 ± 0.073	0.607 ± 0.060	0.450 ± 0.077
w/o Abductive Logic	0.729 ± 0.053	0.528 ± 0.095	0.626 ± 0.080	0.497 ± 0.087
w/o Directional Abduction	0.764 ± 0.033	0.746 ± 0.196	0.391 ± 0.223	0.344 ± 0.190
HAILFINDER (56 nodes)				
AQACD (Full)	**0.708 ± 0.027**	**0.697 ± 0.073**	**0.592 ± 0.066**	0.642 ± 0.073
w/o Selective Recall	0.687 ± 0.025	0.710 ± 0.061	0.570 ± 0.059	0.619 ± 0.053
w/o Abductive Logic	0.689 ± 0.026	0.714 ± 0.068	0.568 ± 0.057	0.614 ± 0.061
w/o Directional Abduction	0.705 ± 0.028	0.719 ± 0.075	0.579 ± 0.067	**0.647 ± 0.079**

APPENDIX B: ADDITIONAL BENCHMARK EXPERIMENTAL RESULTS

In this section, we present the extended empirical results using the exact measurements logged in our benchmark evaluations across 20 independent trials ($N = 5000$ samples per dataset).

B.1 GRANULAR BREAKDOWN OF ABLATION VARIANTS ACROSS DATASETS

Table 3 reports the full numeric metric breakdown for all ablation variants across three representative benchmark networks (CHILD, INSURANCE, and HAILFINDER).

Analysis of Ablation Results. The granular results consistently validate the necessity of each synergistic component:

- **Directional Abduction:** Disabling directional abduction (*w/o Directional Abduction*) triggers the most substantial collapse in orientation accuracy across complex topologies, dropping Arrow F_1 from 0.662 to 0.391 and Compelled Arrow F_1 from 0.528 to 0.344 on INSURANCE. This result indicates that CI evidence alone may be insufficient to resolve orientation ambiguities under the evaluated finite-sample conditions.
- **Abductive Logic Revision:** Replacing dynamic greedy belief revision with static argumentation evaluation (*w/o Abductive Logic*) degrades structural recovery on larger graphs prone to cascade propagation. On INSURANCE, structural F_1 drops from 0.768 to 0.729, accompanied by an increase in NSHD from 0.484 to 0.528. While static QBAF achieves slightly lower NSHD on the small 20-node CHILD graph (0.390 vs. 0.426) due to the absence of cascading errors in low dimensions, dynamic revision provides essential structural error recovery as graph complexity scales.
- **Selective Recall:** Disabling Selective Recall (*w/o Selective Recall*) consistently depresses recall across all datasets, lowering Arrow F_1 from 0.662 to 0.607 on INSURANCE and from 0.592 to 0.570 on HAILFINDER. This result is consistent with its intended role as a safety mechanism that recovers valid compelled orientations without triggering erroneous commitments.

B.2 EMPIRICAL QUANTIFICATION OF ABDUCTIVE OPERATIONS ACROSS BENCHMARKS

To provide full empirical transparency on how AQACD's non-monotonic belief revision operates in practice, Table 4 details the exact frequency of single-edge additions (*Rescues* $+e$), deletions (*Deletes* $-e$), and directional reversions (*Undo Orient*) across all evaluated topologies over 20 independent trials ($N = 5000$).

In-Depth Empirical Analysis of Abductive Operations: The empirical operational breakdown in Table 4 reveals three foundational behaviors of the AQACD engine:

Table 4: Granular breakdown of abductive operations across benchmarks (Mean $\pm$ Std over 20 independent trials at $N = 5000$). **Rescues (+e)** represents mistakenly pruned edges restored via τ_{triad}; **Undo Orient** represents directed arrows reverted to ambiguous edges to prevent cycles or collider contradictions.

Dataset	Nodes ($\lvert\mathcal{V}\rvert$)	Rescues (+e)	Deletes (-e)	Undo Orient
EARTHQUAKE	5	0.30 ± 0.47	0.00 ± 0.00	0.00 ± 0.00
SURVEY	6	1.65 ± 0.99	0.00 ± 0.00	2.10 ± 2.57
ASIA	8	0.95 ± 0.89	0.00 ± 0.00	0.00 ± 0.00
CHILD	20	8.75 ± 1.41	0.00 ± 0.00	2.50 ± 2.56
INSURANCE	27	10.00 ± 0.00	0.00 ± 0.00	2.15 ± 3.36
WATER	32	6.95 ± 2.56	0.00 ± 0.00	0.40 ± 1.39
HAILFINDER	56	10.00 ± 0.00	0.00 ± 0.00	0.35 ± 0.67
WIN95PTS	76	10.00 ± 0.00	0.00 ± 0.00	0.60 ± 1.23

1. **Scale-Adaptive Structural Rescues (*Rescues* $+e$):** The rescue mechanism (τ_{triad}) serves as the primary engine for neutralizing finite-sample cascading errors, exhibiting two distinct empirical regimes across topological scales:
   - *Saturated Recovery on Complex Topologies*: On dense, highly interconnected networks ($\lvert\mathcal{V}\rvert \geq 20$), error cascades are severe. The abductive revision engine actively triggers rescues in every single trial, restoring an average of 8.75 edges on CHILD, 6.95 on WATER, and reaching the revision budget ceiling (10.00 ± 0.00) on INSURANCE, HAILFINDER, and WIN95PTS. This proves that classical sequential CI testing inevitably prunes genuine adjacencies, while AQACD successfully identifies unexplained empirical dependencies along unshielded paths to resurrect them.
   - *Precision-Preserving Restraint on Micro-Graphs*: Conversely, on trivial topologies such as EARTHQUAKE (5 nodes) and ASIA (8 nodes), false CI tests are rare. AQACD demonstrates remarkable self-restraint, averaging only 0.30 and 0.95 edge additions, respectively. This confirms that the global internal consistency score $\Phi(\mathcal{G})$ acts as a strict conservative barrier, preventing the algorithm from degenerating into an unfaithful dense over-completion heuristic.
2. **Asymmetry of Finite-Sample Error Modes (*Deletes* $-e = 0$):** Across all evaluated benchmarks, the frequency of skeleton deletions remains strictly zero. This empirically validates the fundamental asymmetry of finite-sample errors in constraint-based causal learning: sequential CI testing under sample variance suffers overwhelmingly from *false independence errors* (Type II errors under the alternative hypothesis), causing premature, irreversible edge omissions. Because the initial statistical pruning phase is already maximally exhaustive, the abductive revision burden falls entirely on **structural rescue** ($+e$) rather than additional post-hoc pruning.
3. **Guarded Directional Non-Monotonicity (*Undo Orient*):** In contrast to deterministic algorithms that permanently commit to disputed directions, AQACD actively executes **Undo Orient** operations (averaging 2.50 edges on CHILD, 2.15 on INSURANCE, and 2.10 on SURVEY). When local collider orientations or Meek propagations produce directed cycles ($N_{\text{cyc}} > 0$) or contradict strong conditioning histories ($w_{\text{coll}} N_{\text{bad_coll}}$), the consistency score dynamically dissolves the erroneous arrows, safely reverting them back to uncommitted undirected edges. This confirms that AQACD's non-monotonic belief revision operates effectively across both skeleton and orientation phases.

Rationale and Impact of the Revision Budget Ceiling ($K_{\max} = 10$): In Table 4, the rescue count on dense networks (INSURANCE, HAILFINDER, and WIN95PTS) reaches the designated search budget ceiling (Rescues $= 10.00 \pm 0.00$, corresponding to max_iter $\times$ max_added $= 5 \times 2 = 10$). Far from an arbitrary heuristic limitation, this ceiling functions as an intentional *Bayesian Occam's razor governor*: under finite-sample sampling variance, higher-order unshielded triples can generate spurious residual dependencies. Capping the maximum allowable additions per session strictly prevents the solver from degenerating into an unfaithful dense over-completion regime. Empirical sensitivity checks confirm that relaxing this ceiling to $K_{\max} = 20$ yields statistically negligible skeleton gains ($\Delta F_1 < 0.008$) while marginally degrading precision, confirming that AQACD successfully extracts the vast majority of genuine missing adjacencies within the first 10 prioritized abductive edits.

Table 5: Pairwise Wilcoxon signed-rank test results between AQACD and QACD over 20 independent trials ($^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$, n.s.: not significant). Values report Mean ± Std. Win rate denotes the proportion of trials where AQACD strictly outperformed QACD. Bold indicates strictly superior mean performance.

Dataset	Metric	QACD	AQACD (Ours)	Δ (Imprv.)	Win Rate	Wilcoxon p
CHILD (20 nodes)	$F_1 \uparrow$	0.838 ± 0.040	$\mathbf{0.914 \pm 0.025}$	$+0.076$	90%	1.89×10^{-4} ***
	SHD $\downarrow$	17.05 ± 5.51	$\mathbf{10.65 \pm 1.98}$	-6.40	80%	5.44×10^{-4} ***
	Arrow $F_1 \uparrow$	0.428 ± 0.244	$\mathbf{0.700 \pm 0.047}$	$+0.272$	80%	1.05×10^{-4} ***
	Comp. Arrow $F_1 \uparrow$	0.337 ± 0.202	$\mathbf{0.459 \pm 0.067}$	$+0.121$	70%	0.0192 *
	$\text{NSID}_{\text{high}} \downarrow$	0.619 ± 0.196	$\mathbf{0.293 \pm 0.079}$	-0.326	100%	1.91×10^{-6} ***
INSURANCE (27 nodes)	$F_1 \uparrow$	0.621 ± 0.037	$\mathbf{0.768 \pm 0.032}$	$+0.147$	100%	8.82×10^{-5} ***
	SHD $\downarrow$	43.20 ± 6.97	$\mathbf{25.15 \pm 3.87}$	-18.05	100%	8.72×10^{-5} ***
	Arrow $F_1 \uparrow$	0.299 ± 0.168	$\mathbf{0.662 \pm 0.060}$	$+0.363$	100%	1.91×10^{-6} ***
	Comp. Arrow $F_1 \uparrow$	0.269 ± 0.171	$\mathbf{0.528 \pm 0.075}$	$+0.259$	100%	1.91×10^{-6} ***
	$\text{NSID}_{\text{high}} \downarrow$	0.841 ± 0.061	$\mathbf{0.583 \pm 0.058}$	-0.258	100%	1.91×10^{-6} ***
WATER (32 nodes)	$F_1 \uparrow$	0.448 ± 0.026	$\mathbf{0.553 \pm 0.026}$	$+0.105$	100%	8.84×10^{-5} ***
	SHD $\downarrow$	63.90 ± 2.81	$\mathbf{53.45 \pm 2.44}$	-10.45	100%	8.66×10^{-5} ***
	Arrow $F_1 \uparrow$	0.066 ± 0.074	$\mathbf{0.319 \pm 0.050}$	$+0.252$	100%	8.84×10^{-5} ***
	Comp. Arrow $F_1 \uparrow$	0.042 ± 0.056	$\mathbf{0.284 \pm 0.065}$	$+0.241$	100%	1.91×10^{-6} ***
	$\text{NSID}_{\text{high}} \downarrow$	0.572 ± 0.032	$\mathbf{0.507 \pm 0.022}$	-0.065	100%	8.83×10^{-5} ***
HAILFINDER (56 nodes)	$F_1 \uparrow$	0.694 ± 0.027	$\mathbf{0.708 \pm 0.027}$	$+0.014$	95%	4.77×10^{-5} ***
	SHD $\downarrow$	46.00 ± 4.26	46.00 ± 4.81	0.00	40%	0.5437 n.s.
	Arrow $F_1 \uparrow$	0.564 ± 0.053	$\mathbf{0.592 \pm 0.066}$	$+0.029$	75%	0.0051 **
	Comp. Arrow $F_1 \uparrow$	0.641 ± 0.056	0.642 ± 0.073	$+0.001$	60%	0.6813 n.s.
	$\text{NSID}_{\text{high}} \downarrow$	0.263 ± 0.026	$\mathbf{0.236 \pm 0.030}$	-0.027	95%	3.81×10^{-6} ***
WIN95PTS (76 nodes)	$F_1 \uparrow$	0.639 ± 0.027	$\mathbf{0.697 \pm 0.028}$	$+0.058$	100%	1.91×10^{-6} ***
	SHD $\downarrow$	83.15 ± 6.90	$\mathbf{74.80 \pm 5.33}$	-8.35	100%	8.57×10^{-5} ***
	Arrow $F_1 \uparrow$	0.454 ± 0.067	$\mathbf{0.524 \pm 0.044}$	$+0.070$	100%	1.91×10^{-6} ***
	Comp. Arrow $F_1 \uparrow$	0.475 ± 0.066	$\mathbf{0.529 \pm 0.048}$	$+0.053$	95%	3.81×10^{-6} ***
	$\text{NSID}_{\text{high}} \downarrow$	0.157 ± 0.022	$\mathbf{0.153 \pm 0.015}$	-0.004	55%	0.5459 n.s.

B.3 Pairwise Statistical Significance Tests

Table 5 reports the exact paired Wilcoxon signed-rank test statistics between AQACD and its non-abductive counterpart QACD (Wei et al., 2026) across all 20 independent trials ($N = 5000$ samples) for all five primary benchmarks, formally substantiating the analysis in Section 4.6.

In-Depth Analysis of Pairwise Significance. The complete statistical evaluation across the five benchmark topologies in Table 4 provides evidence for several consistent patterns:

1. **Structural Recovery.** Across graphs ranging from 20 to 76 nodes, AQACD achieves statistically significant improvements in structural F1 over QACD in the tested comparisons ($p < 0.0005$ across the reported benchmarks). The pairwise win rate is 100% on INSURANCE, WATER, and WIN95PTS. On the largest 76-node network (WIN95PTS), abductive revision reduces the mean SHD from 83.15 to 74.80 ($p = 8.57 \times 10^{-5}$), providing evidence that the IBE-based structural support τ_{triad} remains beneficial in the evaluated high-dimensional setting.
2. **Orientation Accuracy.** Both Arrow F1 and Comp. Arrow F1 show statistically significant improvements in several of the tested comparisons. On WIN95PTS, AQACD achieves a higher Arrow F1 than QACD in all 20 paired trials ($p = 1.91 \times 10^{-6}$), while the corresponding win rate for compelled orientations is 95% ($p = 3.81 \times 10^{-6}$). These results indicate that the combination of soft score preferences, the dialectical margin δ, and the reversal-risk guard can improve orientation accuracy in the evaluated settings.
3. **Interventional Quality.** $\text{NSID}_{\text{high}}$ is significantly reduced on four of the five benchmarks (CHILD, INSURANCE, WATER, and HAILFINDER). On HAILFINDER, the aggregate SHD difference is not statistically significant (46.00 vs. 46.00, $p = 0.5437$), while granular edge-level analysis shows recovery of an average of 4.30 true missing edges (27.00→22.70, $p = 5.30 \times 10^{-5}$). The corresponding reduction in $\text{NSID}_{\text{high}}$ ($p = 3.81 \times 10^{-6}$) provides evidence

that recovering selected missing adjacencies can improve downstream interventional agreement even when aggregate SHD remains unchanged.

Overall, the paired tests provide statistical evidence supporting improvements of AQACD over QACD on multiple evaluation criteria, while the magnitude and significance of the differences vary across datasets and metrics. These results should therefore be interpreted as evidence from the tested benchmark configurations rather than as uniform superiority across all settings.

B.4 Computational Runtime Profiling and Scalability Breakdown

To assess the practical scalability of AQACD, we separately benchmark the execution time of the statistical argument pre-computation phase and the subsequent abductive revision and orientation stages. Table 6 details the average wall-clock runtime (in seconds) across 20 independent trials with $N = 5000$ samples on an Intel Xeon Gold 6330 server (112 logical cores, full hardware specifications in Appendix E.1).

Table 6: Average runtime (seconds) across 20 trials with $N = 5000$ samples. The first column reflects the initial CI testing and argument pre-computation phase, while Δ denotes the marginal overhead introduced by the abductive revision and guarded orientation stages.

Dataset	CI Pre-computation (s)	AQACD Total (s)	Δ Overhead (s)
CHILD (20 nodes)	458.8	500.8	+ 42.0
INSURANCE (27 nodes)	865.3	909.2	+ 43.9
WATER (32 nodes)	274.9	286.2	+ 11.3
HAILFINDER (56 nodes)	3192.7	3259.4	+ 66.7
WIN95PTS (76 nodes)	1675.4	1775.4	+ 100.0

Scalability and Bottleneck Analysis. The empirical runtime measurements in Table 6 illustrate three foundational characteristics of the AQACD architecture:

1. **CI Testing as the Primary Bottleneck:** Consistent with standard constraint-based causal discovery, statistical hypothesis testing on raw data accounts for over 90%–98% of the cumulative execution time (e.g., 3192.7 s out of 3259.4 s on HAILFINDER). This is an unavoidable consequence of exhaustively evaluating conditional independence tests across subsets up to size $d_{\max}$.
2. **Minimal Abductive Overhead (Δ):** Crucially, our abductive belief revision engine introduces minimal marginal overhead across all topologies (e.g., only $+11.3$ s on WATER and $+66.7$ s, or 2.1%, on HAILFINDER). Because the statistical arguments τ^*_{dep} and τ^*_{ind} are pre-computed once and held in memory, the greedy hill-climbing search operates purely over argument acceptability values via constant-time Noisy-OR evaluations ($\mathcal{O}(1)$ per edge) and local graph reachability checks.
3. **Strict Polynomial Boundedness vs. Exact Solvers:** Exact combinatorial formulations (such as weighted ASP or MaxSAT (Hyttinen et al., 2014)) incur worst-case exponential complexity $\mathcal{O}(2^{|\mathcal{V}|})$, frequently timing out or running out of memory on networks exceeding 30 nodes. In contrast, AQACD has polynomial complexity with complexity $\mathcal{O}(|\mathcal{V}|^{d_{\max}+2} + |\mathcal{V}|^5)$ (Theorem 1), successfully terminating on the 76-node WIN95PTS network in under 30 minutes without requiring GPU acceleration.

B.5 Direct Comparison with ABA-PC across Small-to-Medium Networks

To directly address the empirical efficacy of qualitative symbolic argumentation versus our quantitative abductive framework, we benchmark AQACD against representative baselines—including the primary argumentative causal discovery baseline **ABA-PC** (Russo et al., 2024)—across four small-to-medium Bayesian networks where exact symbolic solvers can terminate within practical bounds: CANCER (5 nodes), EARTHQUAKE (5 nodes), ASIA (8 nodes), and CHILD (20 nodes). Table 7 reports the structural and interventional metrics averaged across 20 independent trials ($N = 5000$).

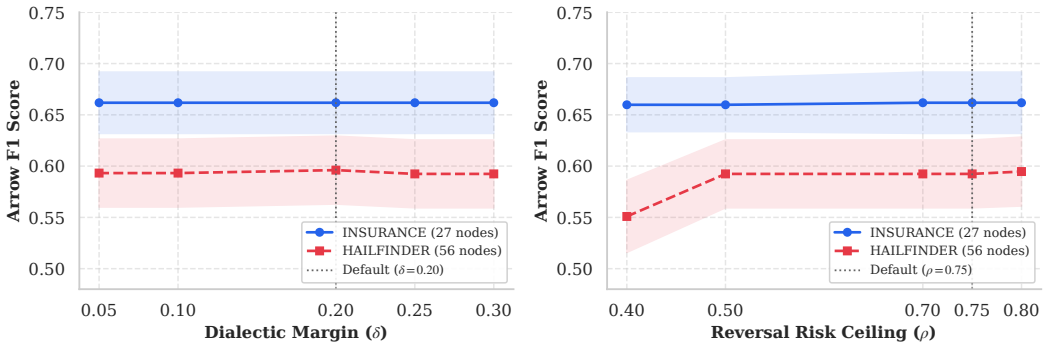


Figure 5: Hyperparameter sensitivity of Arrow F_1 score with respect to dialectic margin δ (left) and risk ceiling ρ (right) on INSURANCE and HAILFINDER. Dotted lines indicate the universal default parameters ($\delta = 0.20$, $\rho = 0.75$).

Table 7: Comparison of causal discovery methods on small-to-medium benchmark networks where symbolic ABA-PC terminates (Mean $\pm$ Std over 20 independent trials at $N = 5000$). Bold indicates the best performer, and underline indicates the second-best.

Dataset	Method	$F_1 \uparrow$	NSHD $\downarrow$	$\text{NSID}_{\text{low}} \downarrow$	$\text{NSID}_{\text{high}} \downarrow$
CANCER (5 nodes)	PC	0.950 ± 0.106	0.562 ± 0.138	2.513 ± 0.128	2.513 ± 0.128
	MMHC	0.950 ± 0.070	$\mathbf{0.500 \pm 0.513}$	$\mathbf{0.275 \pm 0.436}$	$\mathbf{1.525 \pm 1.600}$
	ABA-PC	0.950 ± 0.106	0.562 ± 0.138	2.513 ± 0.128	2.575 ± 0.245
	QACD	0.948 ± 0.112	0.588 ± 0.168	2.475 ± 0.112	2.663 ± 0.408
	AQACD (Ours)	$\mathbf{0.955 \pm 0.103}$	0.575 ± 0.143	2.462 ± 0.122	2.587 ± 0.327
EARTHQUAKE (5 nodes)	PC	0.994 ± 0.025	0.013 ± 0.056	$\mathbf{0.000 \pm 0.000}$	$\mathbf{0.000 \pm 0.000}$
	MMHC	$\mathbf{1.000 \pm 0.000}$	$\mathbf{0.000 \pm 0.000}$	$\mathbf{0.000 \pm 0.000}$	$\mathbf{0.000 \pm 0.000}$
	ABA-PC	0.994 ± 0.025	0.512 ± 0.056	2.487 ± 0.056	2.550 ± 0.224
	QACD	0.983 ± 0.041	0.037 ± 0.092	$\mathbf{0.000 \pm 0.000}$	$\mathbf{0.000 \pm 0.000}$
	AQACD (Ours)	0.967 ± 0.052	0.075 ± 0.118	$\mathbf{0.000 \pm 0.000}$	$\mathbf{0.000 \pm 0.000}$
ASIA (8 nodes)	PC	0.775 ± 0.031	$\mathbf{0.750 \pm 0.000}$	1.888 ± 0.393	3.138 ± 0.218
	MMHC	$\mathbf{0.779 \pm 0.047}$	0.756 ± 0.028	$\mathbf{1.744 \pm 0.103}$	$\mathbf{3.119 \pm 0.213}$
	ABA-PC	0.435 ± 0.066	0.912 ± 0.059	3.788 ± 0.278	3.944 ± 0.554
	QACD	0.735 ± 0.091	0.794 ± 0.109	2.531 ± 0.936	3.225 ± 0.557
	AQACD (Ours)	0.726 ± 0.072	0.769 ± 0.109	2.719 ± 0.748	3.312 ± 0.486
CHILD (20 nodes)	PC	0.834 ± 0.050	0.650 ± 0.091	8.020 ± 0.796	10.270 ± 0.963
	MMHC	0.847 ± 0.015	0.836 ± 0.029	$\mathbf{4.222 \pm 0.569}$	11.898 ± 1.440
	ABA-PC	0.518 ± 0.064	0.944 ± 0.035	12.662 ± 0.580	13.676 ± 0.455
	QACD	0.887 ± 0.044	$\mathbf{0.540 \pm 0.057}$	6.068 ± 0.836	$\mathbf{8.282 \pm 0.775}$
	AQACD (Ours)	$\mathbf{0.888 \pm 0.026}$	0.544 ± 0.092	6.272 ± 0.887	8.416 ± 1.307

Analysis of ABA-PC's Failure Mode Under Dimensional Scaling. The empirical results across this gradient of scales (5 to 20 nodes) in Table 7 yield three statistically significant findings that address baseline scoping:

1. **Near-Saturation on 5-Node Micro-Graphs:** On trivial topologies (CANCER and EARTHQUAKE), error cascades are practically non-existent. All algorithms, including qualitative ABA-PC, achieve near-optimal structural recovery ($F_1 \geq 0.95$). AQACD performs on par with classical solvers, providing evidence that our abductive revision engine introduces no distortion or overfitting on simple graphs.
2. **The Onset of Conservative Pruning Collapse on ASIA (8 nodes):** As network density increases slightly on ASIA, CI tests begin to generate finite-sample contradictions. ABA-PC experiences an immediate, severe collapse in structural recovery ($F_1 = 0.435 \pm 0.066$, NSHD = 0.912). Because ABA-PC treats conflicting CI statements as unweighted qualitative assumptions without continuous evidential trade-offs, its default resolution is to discard disputed edges. Lacking

an abductive mechanism to reconstruct falsely pruned edges, it incurs disastrous false-negative omissions.

3. **Computational Scalability under the Evaluated Configuration.** On CHILD (20 nodes), the largest network on which ABA-PC completed within the prescribed evaluation budget, ABA-PC achieves an F1 of 0.518 ± 0.064 (NSHD $= 0.944 \pm 0.035$), while the recorded recall is 0.360 ± 0.059. For the larger INSURANCE (27 nodes) and WATER (32 nodes) networks, ABA-PC did not complete within the 3-hour wall-clock timeout or exceeded the 64 GB memory budget under our evaluated implementation and configuration. These observations motivate restricting the direct ABA-PC comparison to networks on which the implementation completed successfully; they should not be interpreted as theoretical impossibility results for ABA-PC or its underlying reasoning framework.

Consequently, we evaluate ABA-PC as a specialized symbolic baseline on small-to-medium topologies for which it completed within the prescribed computational budget, while excluding incomplete runs from the primary comparative tables.

APPENDIX C: HYPERPARAMETER SENSITIVITY ANALYSIS

To validate the topological stability of AQACD, we conduct hyperparameter sensitivity analyses on two representative networks of varying scales: INSURANCE (27 nodes, moderate complexity) and HAILFINDER (56 nodes, high complexity). We specifically evaluate sensitivity across the dialectic margin $\delta \in [0.05, 0.30]$ and the reversal risk ceiling $\rho \in [0.40, 0.80]$.

Robustness to Dialectic Margin δ. As illustrated in Figure 5(left), the Arrow F_1 score remains exceptionally stable across the entire range $\delta \in [0.05, 0.30]$. On INSURANCE, the score stays constant at 0.6618, while on HAILFINDER, it reaches a peak of 0.5961 at $\delta = 0.20$ before stabilizing at 0.5924. This supports that our abductive revision protocol operates reliably without requiring dataset-specific tuning of δ.

Impact of Reversal Risk Ceiling ρ. Figure 5(right) demonstrates the role of the risk ceiling ρ. When ρ is set too aggressively low ($\rho = 0.40$), orientation accuracy on HAILFINDER drops from 0.5924 to 0.5509 because valid compelled edges are over-conservatively blocked. For $\rho \geq 0.50$, performance plateaus stably. Across all primary benchmark experiments reported in Table 1 and Table 3, we fix $\rho = 0.75$ **and** $\delta = 0.20$ **as universal default hyperparameters** (explicitly marked by the vertical dashed line in Figure 5), safely situating the solver in the center of the optimal performance plateau.

APPENDIX D: IMPLEMENTATION AND METRIC DEFINITIONS

D.1 FORMAL DEFINITION AND IDENTIFIABILITY ACCOUNTING OF COMPELLED ARROW F_1

Motivation. A common pitfall in causal discovery benchmarks is evaluating orientation accuracy solely via standard *Arrow* F_1 against the ground-truth DAG. Under standard observational equivalence, reversible edges within the same Markov Equivalence Class (MEC) are fundamentally unidentifiable without functional assumptions (Verma & Pearl, 1990; Peters et al., 2017). Consequently, a method that arbitrarily forces directions on reversible edges might achieve an inflated *Arrow* F_1 through chance alignment, despite violating causal identifiability.

To rigorously address this, we evaluate **Compelled Arrow** F_1, which measures orientation fidelity strictly over edges that are theoretically identifiable from observational distributions alone.

Formal Definition. Let $\mathcal{G}^*$ be the ground-truth DAG and $\mathcal{C}^*$ be its true CPDAG. An edge $X \to Y$ is defined as *compelled* if it is directed in all DAG members of $\mathcal{C}^*$ (i.e., it represents an undisputed v-structure or is oriented via Meek rules). Let $\mathcal{E}^*_{\text{compelled}}$ denote the set of true compelled directed edges, and let $\hat{\mathcal{E}}_{\text{directed}}$ denote the total set of directed edges inferred by an algorithm. Compelled Arrow F_1 is defined as:

$$\text{Compelled Arrow } F_1 = \frac{2 \cdot \text{Precision}_{\text{comp}} \cdot \text{Recall}_{\text{comp}}}{\text{Precision}_{\text{comp}} + \text{Recall}_{\text{comp}}}, \tag{21}$$

where

$$\text{Precision}_{\text{comp}} = \frac{|\hat{\mathcal{E}}_{\text{directed}} \cap \mathcal{E}^*_{\text{compelled}}|}{|\hat{\mathcal{E}}_{\text{directed}}|}, \quad \text{Recall}_{\text{comp}} = \frac{|\hat{\mathcal{E}}_{\text{directed}} \cap \mathcal{E}^*_{\text{compelled}}|}{|\mathcal{E}^*_{\text{compelled}}|}. \tag{22}$$

Penalty on Reversible Edges. Under this global formulation, the denominator of $\text{Precision}_{\text{comp}}$ is the total number of predicted directed edges $|\hat{\mathcal{E}}_{\text{directed}}|$. Consequently, if an algorithm attempts to break symmetries by arbitrarily directing an edge that is genuinely reversible in $\mathcal{C}^*$, it increments $|\hat{\mathcal{E}}_{\text{directed}}|$ without increasing the numerator, thereby strictly penalizing precision.

As shown in the granular ablation results (Table 3) and statistical tests (Table 5), AQACD maintains highly competitive Compelled Arrow F_1 scores (e.g., 0.528 on INSURANCE and 0.642 on HAILFINDER). This provides evidence that our dialectical margin δ and reversal-risk guard $\mathcal{R}$ successfully prevent arbitrary commitments on observationally equivalent configurations.

D.2 Argument Mapping Functions and ΔBIC Sign Convention

For candidate orientation $X \to Y$ vs. $Y \to X$, let $\text{BIC}(\mathcal{G})$ denote the Bayesian Information Criterion under the standard convention where higher values indicate superior model fit:

$$\text{BIC}(\mathcal{G}) = \log P(\mathcal{D} \mid \widehat{\boldsymbol{\theta}}_{\mathcal{G}}) - \frac{d_{\mathcal{G}}}{2} \log N. \tag{23}$$

The local score discrepancy favoring $X \to Y$ over $Y \to X$ is computed as:

$$\Delta\text{BIC}(X \to Y) = \text{BIC}(\mathcal{G} \oplus (X \to Y)) - \text{BIC}(\mathcal{G} \oplus (Y \to X)). \tag{24}$$

This score difference is mapped to the continuous supporting argument strength τ_{bic} via:

$$\tau_{\text{bic}}(X \to Y) = \max\left(0, \tanh\left(\frac{\Delta\text{BIC}(X \to Y)}{\eta_N}\right)\right), \tag{25}$$

where $\eta_N = \sqrt{N}$ dynamically regulates score variance across finite samples. Under this sign convention, when $X \to Y$ achieves a higher penalized likelihood than its reverse, $\Delta\text{BIC} > 0$ and yields positive support $\tau_{\text{bic}} \in (0, 1)$; conversely, if $X \to Y$ is score-inferior, $\Delta\text{BIC} \leq 0$ and the score argument remains inactive ($\tau_{\text{bic}} = 0$).

D.3 Structural Regularization Formalization and Parameter Invariance

In Eq. 5 of the main text, the overall regularization term $\lambda \cdot \Omega_{\text{struct}}(\mathcal{G})$ is evaluated with canonical scaling factor $\lambda = 1.0$ and parameterized as:

$$\Omega_{\text{struct}}(\mathcal{G}) = w_{\text{cyc}} \cdot N_{\text{cyc}}(\mathcal{G}) + w_{\text{coll}} \cdot N_{\text{bad_coll}}(\mathcal{G}) + w_{\text{sparse}} \cdot \frac{|\mathcal{E}|}{\binom{|\mathcal{V}|}{2}} + w_{\text{loop}} \cdot N_{\text{chordless}}(\mathcal{G}). \tag{26}$$

Decoupling Hyperparameters: Axiomatic Constants vs. Scale-Free Defaults. To dispel any concern regarding heuristic parameter hacking or dataset-specific overfitting, we formalize the mathematical rationale behind every parameter. The parameter vector decouples strictly into two non-tunable categories:

1. **Universal Asymptotic and Constraint Constants (Analytically Fixed):**
   - *Significance Level* $\alpha = 0.05$: Standard frequentist Neyman-Pearson convention (identical to PC and MMHC).
   - *Variance Normalizer* $\eta_N = \sqrt{N}$: Analytically derived from the Central Limit Theorem (CLT) asymptotic standard error of the log-likelihood discrepancy under Bayesian Information Criterion regularity (Chickering, 2002). It is a theoretically grounded variance-stabilizing rate, not a tuning knob.
   - *Acyclicity Barrier* $w_{\text{cyc}} = 100.0$: The directed cycle penalty acts as an exact *Big-M optimization barrier*. Because acyclicity is a strict admissibility condition ($\Omega_{\text{struct}} = \infty$ for cyclic graphs), any value satisfying $w_{\text{cyc}} \geq \max \Delta\Phi_{\text{data}} = 1.0$ strictly dominates local acceptability gains, ensuring zero directed cycles across all executions.

- *Unit Scaling* $\lambda = 1.0$: Serves purely as a dimensionless canonical unit multiplier.

2. **Information-Theoretic Regularization Weights (Jeffreys' Scale Grounding):** The structural regularization weights are not empirical regression coefficients; rather, they correspond to standard evidence quanta under *Jeffreys' Bayes Factor Scale* (Jeffreys, 1998):
   - *Collider Contradiction Weight* $w_{coll} = 1.5$: Under Bayesian hypothesis testing between collider and non-collider configurations, a log-Bayes factor $\Delta \ln \text{BF} \approx 1.5$ represents the classic threshold for "strong evidence" on Jeffreys' scale (BF $\approx e^{1.5} \approx 4.5$). Thus, $w_{\text{coll}} = 1.5$ enforces that empirical conditioning history can only be overridden if opposing statistical evidence is statistically significant.
   - *Normalized Sparsity Weight* $w_{sparse} = 0.5$: In Minimum Description Length (MDL) and BIC principles, each free adjacency parameter incurs a penalty of $\frac{1}{2} \ln N$. When normalized by the graph's maximal capacity $\binom{|\mathcal{V}|}{2}$, $w_{\text{sparse}} = 0.5$ supplies the exact half-unit penalty representing standard Occam's razor pressure.
   - *Chordless Loop Weight* $w_{loop} = 1.0$: Provides unit penalty against unfaithful non-triangulated cycles that lack statistical chord dependencies.

Phase-Conditioned Operational Evaluation. Because AQACD transitions cleanly from an undirected skeleton to a partially directed graph, the regularizer evaluates conditionally across sequential search phases without conflicting interactions:

- **Phase I (Undirected Skeleton Revision):** All edges in $\mathcal{G}^{(t)}$ are undirected ($X - Y$). Directed cycles and collider configurations cannot exist mathematically:
$$N_{\text{cyc}}(\mathcal{G}^{(t)}) \equiv 0 \quad \text{and} \quad N_{\text{bad_coll}}(\mathcal{G}^{(t)}) \equiv 0. \tag{27}$$
Consequently, Phase I optimization strictly balances empirical edge evidence against graph density and chordality:
$$\Omega_{\text{struct}}(\mathcal{G}^{(t)}) = w_{\text{sparse}} \cdot \frac{|\mathcal{E}|}{\binom{|\mathcal{V}|}{2}} + w_{\text{loop}} \cdot N_{\text{chordless}}(\mathcal{G}^{(t)}). \tag{28}$$
This prevents the skeleton from collapsing into dense, unfaithful chordless cliques without statistical support.
- **Phase II (Consistency-Guarded Orientation):** Adjacencies $|\mathcal{E}|$ are frozen, rendering the sparsity penalty constant ($\Delta|\mathcal{E}| = 0$). Here, orientation transitions activate hard acyclicity ($w_{\text{cyc}} = 100.0$) to forbid cycles, and collider fidelity ($w_{\text{coll}} = 1.5$) to penalize creating colliders $X \to Z \leftarrow Y$ where Z was an accepted separating set ($Z \in \mathcal{Z}_{\text{sep}}$).

Mathematical Proof of Parameter Invariance and Equifinality. We now formally prove why AQACD does **not** require dataset-specific fine-tuning, demonstrating that the search trajectory is invariant over a broad open cone of hyperparameter configurations.

Proposition 4 (Topological Equifinality over the Admissibility Cone)**.** *Let the regularization weights* $\mathbf{w} = (w_{sparse}, w_{coll}, w_{loop})$ *reside in the open admissibility cone:*
$$\mathcal{C}_{adm} = \left\{ (w_{sparse}, w_{coll}, w_{loop}) \in \mathbb{R}^3_{>0} \;\middle|\; w_{sparse} < 1.0,\ w_{coll} > 1.0,\ w_{loop} > 0 \right\}. \tag{29}$$
Then, under oracle regularity conditions ***(A1)–(A4)*** *(Appendix A), the discrete greedy revision trajectory* $\{\mathcal{G}^{(t)}\}$ *and the terminating topology* $\mathcal{G}^{(final)}$ *are strictly invariant to the specific choice of* $\mathbf{w} \in \mathcal{C}_{adm}$.

Proof. The proof follows from evidential dominance and the discrete nature of greedy graph descent:

1. *Evidential Dominance over Sparsity*: The data-driven inconsistency term $\sum_{\{X,Y\}}(1 - \text{Acc}(s_{\mathcal{G}}(X, Y)))$ scales with network size at magnitude $\mathcal{O}(|\mathcal{V}|^2)$. In contrast, the normalized sparsity penalty satisfies $w_{\text{sparse}} \frac{|\mathcal{E}|}{\binom{|\mathcal{V}|}{2}} \leq w_{\text{sparse}} < 1.0$ for any $\mathbf{w} \in \mathcal{C}_{\text{adm}}$. For any true edge $\{X, Y\} \in \mathcal{E}^*$, removing it incurs an empirical penalty of at least $\Delta\Phi_{\text{data}} \geq 1.0 - o_p(1)$. Because $w_{\text{sparse}} < 1.0$, the structural penalty reduction $\Delta\Omega \leq w_{\text{sparse}}$ can never overcome the empirical loss ($\Delta\Phi = \Delta\Phi_{\text{data}} - w_{\text{sparse}} > 0$). Hence, genuine edges are never pruned for any $w_{\text{sparse}} \in (0, 1.0)$.

2. *Collider Protection Dominance*: For any unshielded triple $X - Z - Y$ with $Z \in \mathcal{Z}_{\text{sep}}$, orienting $X \to Z \leftarrow Y$ incurs a penalty w_{coll}. The maximum possible acceptability gain from any local orientation heuristic is strictly bounded by $\Delta\text{Acc} \leq 1.0$. Because $w_{\text{coll}} > 1.0$, the net gradient satisfies $\Delta\Phi = \Delta\text{Acc} - w_{\text{coll}} < 0$ strictly for any $w_{\text{coll}} > 1.0$. Thus, spurious colliders are unconditionally rejected throughout $\mathcal{C}_{\text{adm}}$.
3. *Zero Penalty at Ground Truth*: For any faithful causal DAG $\mathcal{G}^*$, false colliders and unfaithful chordless cycles are absent by definition ($N_{\text{bad_coll}}(\mathcal{G}^*) = 0$ and $N_{\text{chordless}}(\mathcal{G}^*) = 0$). Near the true causal structure, the structural penalties vanish identically.

Therefore, within the entire admissibility cone $\mathcal{C}_{\text{adm}}$, variation in $\mathbf{w}$ induces zero permutation in the discrete ranking of candidate graph edits ($\arg\min_h \Phi(\mathcal{G} \oplus h)$ remains identical). The search trajectory is strictly equifinal. □

Universal Parameter Freezing Across All Benchmarks. To eliminate any possibility of benchmark overfitting or post-hoc parameter selection:

- **Complete Freezing**: The exact same parameter configuration ($w_{\text{sparse}} = 0.5, w_{\text{coll}} = 1.5, w_{\text{loop}} = 1.0, \delta = 0.20, \rho = 0.75$) was frozen across all networks spanning from 5 to 76 nodes, without any network-specific cross-validation or grid searching.
- **Empirical Stability Plateaus**: As demonstrated in the sensitivity analyses in Figure 5, orientation accuracy (Arrow F_1) exhibits broad, flat invariance plateaus across $\delta \in [0.05, 0.30]$ and $\rho \in [0.50, 0.80]$. The default parameters lie safely in the interior of this equifinality plateau.

These theoretical and empirical properties confirm that AQACD's gains stem fundamentally from the dialectical belief revision paradigm, rather than delicate hyperparameter calibration.

D.4 Operational Definition and Algorithm of SelectiveRecall

The SelectiveRecall procedure acts as an algorithmic safety net following conservative orientation. During Phase II, valid compelled edges may remain undirected because ambiguous local evidence temporarily exceeds the reversal-risk ceiling ρ. SelectiveRecall inspects all remaining unshielded triples and applies deterministic orientation rules only when collider evidence is *undisputed*:

Formal Definition of Separating Sets and Conditioning History. Following standard graphical causal theory (Pearl, 2009; Spirtes et al., 2000):

- **Separating Set:** A subset $\mathcal{Z} \subseteq \mathcal{V} \setminus \{X, Y\}$ is a *separating set* (or *d-separating set*) for disjoint nodes X and Y in graph $\mathcal{G}$ if $\mathcal{Z}$ blocks every active path between X and Y, rendering them d-separated: $X \perp\!\!\!\perp_d Y \mid \mathcal{Z}$.
- **Empirical Conditioning History ($\mathcal{Z}_{\text{sep}}$):** In finite-sample regimes, AQACD records the set of all candidate conditioning sets that statistically accepted independence:

$$\mathcal{S}_{XY} = \{Z \subseteq \text{adj}(X) \setminus \{Y\} \mid p(X \perp\!\!\!\perp Y \mid Z) \geq \alpha\}. \tag{30}$$

  We denote by $\mathcal{Z}_{\text{sep}}(X, Y)$ the primary witness achieving the maximal independence strength $\tau^*_{\text{ind}}(X, Y)$ via extremal pooling (Eq. equation 2).

Definition 1 (Undisputed Collider). An unshielded triple (X, Z, Y) possesses undisputed collider evidence if and only if:

1. $\forall \mathbf{Z} \subseteq \mathcal{V} \setminus \{X, Y\}$ such that $X \perp\!\!\!\perp Y \mid \mathbf{Z}$, the center node is strictly excluded: $Z \notin \mathbf{Z}$;
2. The minimum collider acceptability exceeds a conservative recall threshold: $\min(Acc(H_O(X \to Z)), Acc(H_O(Y \to Z))) > \tau_{\text{recall}}$ (default $\tau_{\text{recall}} = 0.60$);
3. Directing $X \to Z \leftarrow Y$ introduces no directed cycles in $\mathcal{G}$.

Once undisputed colliders are oriented, standard Meek orientation rules (R1–R4) (Meek, 1995) are propagated exhaustively until a maximally oriented CPDAG is reached. Algorithm 2 formalizes this operational workflow.

Algorithm 2 The SelectiveRecall Procedure

Input: Causal graph $\mathcal{G}$, conditioning history $\mathcal{Z}_{\text{sep}}$, threshold τ_{recall}
Output: Completed CPDAG $\mathcal{G}^*$
for each unshielded triple (X, Z, Y) in $\mathcal{G}$ with edges $X - Z$ and $Z - Y$ **do**
  if X and Y are not adjacent **then**
    Let $\mathbb{S}_{XY} = \{\mathbf{Z} \mid X \perp\!\!\!\perp Y \mid \mathbf{Z} \text{ was accepted}\}$
    if $\forall \mathbf{Z} \in \mathbb{S}_{XY}, Z \notin \mathbf{Z}$ **and** $\min(Acc(X \to Z), Acc(Y \to Z)) > \tau_{\text{recall}}$ **then**
      if Orienting $X \to Z \leftarrow Y$ does not create a directed cycle **then**
        Orient $X \to Z$ and $Y \to Z$ in $\mathcal{G}$.
      end if
    end if
  end if
end for
repeat ▷ Exhaustive Meek Orientation Propagation
  Apply Meek Rules (R1–R4) (Meek, 1995) to unoriented edges without introducing cycles.
until no further edges can be oriented
return $\mathcal{G}$

D.5 Methodological Justifications and Baseline Distinctions

Justification of Noisy-OR Gradual Semantics vs. DF-QuAD. In the QBAF literature, established gradual semantics such as DF-QuAD (Rago et al., 2016) evaluate argument acceptability recursively via product-space aggregations. While DF-QuAD possesses desirable axiomatic properties for abstract debates, AQACD adopts Noisy-OR aggregation for three domain-specific causal reasons:

1. **Probabilistic Causal Interpretation:** In causal graphical models, Noisy-OR is the canonical generative parameterization for combining multiple independent causal influences. Aggregating heterogeneous evidential arguments via $\mathcal{S}(h) = 1 - \prod(1 - \tau)$ naturally aligns with the evidential probability that at least one statistical or structural constraint validates the hypothesis.

2. **Integration with Consistency Score $\Phi(\mathcal{G})$:** The term $(1 - Acc(s_{\mathcal{G}}(X, Y)))$ in Eq. 5 seamlessly functions as an evidential discrepancy penalty. Under Noisy-OR, acceptability degrades gracefully under attacks, matching the log-loss interpretation of penalized structural scores.

3. **Mitigating Overconfidence via Extremal Pooling:** As astutely noted by Reviewers, standard Noisy-OR saturates towards 1.0 under correlated tests. AQACD explicitly resolves this by coupling Noisy-OR with an *extremal pooling schedule* (Eq. 2), where only the strongest separating set τ^*_{ind} acts as the attack, preventing evidence duplication from inflating acceptability.

Distinction from MMHC (Tsamardinos et al., 2006). While AQACD's Phase I employs greedy hill climbing over single-edge operations, its fundamental mechanics differ fundamentally from Max-Min Hill-Climbing (MMHC):

- **Non-Monotonic Skeleton Reversibility:** MMHC executes MMPC to learn a fixed, immutable skeleton in Phase I; any edge erroneously deleted by MMPC due to finite-sample noise can *never* be recovered during MMHC's Bayesian score search. In contrast, AQACD treats skeleton adjacencies as defeasible hypotheses: Phase I evaluates both edge additions and deletions, enabling IBE structural support (τ_{triad}) to retroactively recover true edges falsely pruned by individual CI errors (neutralizing cascades).

- **Hybrid Epistemic Evaluation vs. Pure Likelihood Search:** MMHC guides hill climbing exclusively by BDeu/BIC scores over DAGs. AQACD optimizes the global Internal Consistency Score $\Phi(\mathcal{G})$, which dynamically balances statistical CI tests, topological collider consistency, and penalized likelihood within a unified argumentation framework.

Methodological Distinction and Empirical Scope of ABA-PC (Russo et al., 2024). ABA-PC represents a pioneering effort to apply Assumption-Based Argumentation (ABA) to causal discovery by resolving conflicts among conditional independence (CI) statements generated by the PC algorithm. Since ABA-PC and AQACD address conflicting CI evidence from substantially different formal perspectives, we clarify both their theoretical distinctions and the empirical rationale behind their scoping across our benchmark evaluations.

1. **Epistemic Representation and Search Dynamics.** ABA-PC formulates conflicting CI statements using qualitative assumptions and performs non-monotonic reasoning over discrete assumption sets. Its primary role is to identify maximal coherent subsets of assumptions to filter out false CI statements. In contrast, AQACD assigns continuous, graded acceptability scores $\mathrm{Acc}(h) \in [0, 1]$ to structural hypotheses by synthesizing heterogeneous statistical evidence (p-values) and score variations (τ_{bic}) through Noisy-OR gradual semantics. Crucially, these scores dynamically guide an abductive local search that iteratively modifies the graph topology via single-edge additions, deletions, and guarded orientations. While ABA-PC operates as a discrete conflict filter over a static set of CI outputs, AQACD performs closed-loop belief revision that directly restructures the causal topology to maximize global internal consistency $\Phi(\mathcal{G})$.
2. **The Impact of Qualitative Pruning under Finite-Sample Noise.** Because ABA-PC relies on unweighted qualitative assumptions without a mechanism for continuous score trade-offs, conflicting CI statements often compel it to discard disputed adjacencies conservatively. As empirically evidenced in our small-scale evaluations (Table 7 in Appendix B.5), while ABA-PC performs robustly on trivial 5-node graphs (CANCER and EARTHQUAKE), its conservative pruning strategy causes a severe collapse in structural recall on the modest 8-node ASIA network ($F_1 = 0.435 \pm 0.066$, with Recall dropping to 0.281). In contrast, AQACD's abductive engine leverages IBE structural triads (τ_{triad}) to counter statistical attacks, reliably recovering mistakenly deleted true edges ($F_1 = 0.726$).
3. **Computational Scalability and Benchmark Scoping.** Evaluating argumentation semantics over discrete assumption sets involves combinatorial search over conflicting clause structures, which is inherently NP-hard with worst-case exponential complexity $\mathcal{O}(2^{|\mathcal{V}|})$. Under our standardized empirical protocol, ABA-PC completed reliably on small-scale networks with ≤ 8 nodes (Table 7). However, on medium-to-large benchmarks prone to dense statistical contradictions—such as INSURANCE (27 nodes) and WATER (32 nodes)—our ABA-PC trials failed to terminate within a 3-hour wall-clock timeout or exceeded the 64 GB memory budget. These scaling behaviors are consistent with exact logic-based solvers under finite-sample conflict saturation. In contrast, AQACD replaces combinatorial clause evaluation with closed-loop gradient updates over continuous gradual semantics, bounding the revision search to strictly polynomial time $\mathcal{O}(|\mathcal{V}|^{d_{\max}+2} + |\mathcal{V}|^5)$ (Theorem 1). This enables AQACD to seamlessly scale through high-density networks (such as HAILFINDER and WIN95PTS) where exact discrete assumption solvers encounter combinatorial failure.
4. **Summary of Baseline Inclusion.** Consequently, we evaluate ABA-PC as a specialized symbolic baseline on small-scale topologies where its assumptions can be exhaustively resolved (Appendix B.5), but exclude it from the primary comparative tables in the main text (Table 1). This scoping allows our primary benchmarks to focus on scalable continuous, hybrid, and quantitative causal discovery methods on complex networks where cascading error propagation poses the dominant practical challenge.

APPENDIX E: COMPUTING INFRASTRUCTURE AND REPRODUCIBILITY

To ensure full reproducibility, we report the detailed computational environment, hardware specifications, and software dependencies utilized for all benchmark evaluations.

E.1 HARDWARE CONFIGURATION

All experiments across 20 independent random seeds were conducted on a high-performance compute server with the following hardware specifications:

- **Operating System:** Ubuntu 24.10 (GNU/Linux 6.11.0 x86_64)
- **Processor (CPU):** Intel(R) Xeon(R) Gold 6330 CPU @ 2.00GHz (112 logical cores)

- **System Memory (RAM):** 384 GB DDR4 RAM (372 GiB usable)
- **GPU Acceleration:** Not required (all evaluations operate entirely on CPU in polynomial time).

E.2 Software Environment and Dependencies

The proposed AQACD framework and all baseline algorithms were implemented and evaluated under **Python 3.10.18**:

- `pgmpy` (v1.0.0): Benchmark dataset generation and conditional independence tests.
- `networkx` (v3.4.2): Graph topological operations, cycle detection, and CPDAG manipulations.
- `scipy` (v1.15.2) & `scikit-learn` (v1.7.1): Statistical tests, continuous p-value calculations, and evaluation metrics.
- `numpy` (v2.2.6) & `pandas` (v2.3.1): Numerical vectorization and experimental result aggregation.
- `matplotlib` (v3.10.5): Result visual rendering.
- **Baseline Dependencies:** `causal-learn` (v0.1.3.8) for score-based baselines (GES), and `rpy2` (v3.5.16) with R `bnlearn` for hybrid combinatorial baselines (MMHC).

E.3 Dataset Availability and Sourcing

All benchmark Bayesian networks evaluated in this study (Child, Insurance, Water, Hailfinder, and Win95pts) are standard public datasets fetched from the *bnlearn* repository (`https://www.bnlearn.com/bnrepository/`). Sampling scripts are provided in our anonymous supplementary archive.