

Causal Discovery as Dialectical Aggregation: A Quantitative Argumentation Framework

Sheng Wei^{1,2,4}, Yulin Chen¹, Beishui Liao^{2,3,4*}

¹The College of Computer Science and Technology, Zhejiang University

²ZLAIRE, Zhejiang University

³School of Philosophy, Zhejiang University

⁴The State Key Lab of Brain-Machine Intelligence
{shengwei, yulinchen, baiseliao}@zju.edu.cn

Abstract

Constraint-based causal discovery is brittle in finite-sample regimes because erroneous conditional-independence (CI) decisions can cascade into substantial structural errors. We propose Quantitative Argumentation for Causal Discovery (QACD), a semantics-driven framework that represents CI outcomes as graded, defeasible arguments rather than irreversible constraints. QACD maps statistical test outcomes to argument strengths and aggregates conflicting evidence through connectivity-mediated witness propagation, producing a fixed-point acceptability labeling over candidate adjacencies. Experiments on standard benchmark Bayesian networks suggest that QACD improves structural coherence and interventional reliability in several noisy or inconsistent CI regimes, while remaining competitive with classical constraint-based, hybrid, and prior argumentation-based baselines.

1 Introduction

Causal discovery, the task of inferring causal structure from observational data, is a central problem in Artificial Intelligence, as it enables reasoning systems to move beyond statistical association toward interventional and counterfactual reasoning (Neuberg 2003; Yang et al. 2025). The objective is to reconstruct the underlying Directed Acyclic Graph (DAG), or more practically its Markov equivalence class—typically represented as a Completed Partially Directed Acyclic Graph (CPDAG)—from independent and identically distributed (i.i.d.) samples drawn from an unknown joint distribution. Such causal models underpin applications ranging from gene regulatory network analysis (Sachs et al. 2005) to economic and social policy evaluation (Athey and Imbens 2017).

Two paradigms dominate the literature. Score-based methods search for a graph maximizing a penalized likelihood score (Chickering 2002; Zheng et al. 2018), while constraint-based methods infer structure by testing conditional-independence (CI) relations (Spirtes, Glymour, and Scheines 2000; Colombo, Maathuis, and others 2014). Within the latter family, the Peter–Clark (PC) algorithm remains a canonical reference due to its efficiency and asymptotic guarantees: under standard assumptions, PC is sound

and consistent when CI queries are answered by an oracle.

In practice, however, CI evidence obtained from statistical testing is inherently uncertain. Finite samples, measurement error, and mismatches between test assumptions and data routinely produce both false independences and false dependences. Constraint-based pipelines nevertheless treat accepted CI statements monotonically: once an edge is removed or a separating set is fixed, later evidence cannot overturn that commitment. Under finite samples, such monotonic reasoning is brittle, as early CI errors can accumulate into globally inconsistent commitments (Ling et al. 2024) and cascade through subsequent orientation steps (Ding et al. 2024; Uhler et al. 2013).

We argue that this brittleness stems from a fundamental conceptual mismatch: uncertain statistical inferences are treated as hard logical constraints rather than defeasible evidence (Uhler et al. 2013). Consequently, causal discovery should be reframed as a problem of defeasible reasoning rather than pure constraint satisfaction.

This motivates an argumentation-theoretic perspective, in which CI outcomes are represented as graded pieces of evidence whose conflicts must be resolved dialectically (Dung 1995; Pollock 1987; Wang and Shen 2025). Existing uses of argumentation in causal discovery, however, have mostly focused on later stages such as edge orientation or on qualitative rather than quantitative evidence (Russo, Rapberger, and Toni 2024). What therefore remains largely unexplored is how to integrate graded statistical evidence directly into the skeleton discovery stage through an explicit argumentation semantics.

To address this gap, we propose *Quantitative Argumentation for Causal Discovery* (QACD), which evaluates candidate adjacencies by dialectically aggregating uncertain CI evidence. Its main contributions are:

- **A quantitative argumentation formalization of CI-based causal discovery.** CI outcomes are interpreted as graded, defeasible arguments rather than irreversible constraints, enabling reasoning under inconsistent statistical evidence.
- **Connectivity-mediated propagated attenuation.** We introduce a structure-dependent conflict mechanism through which global graph context modulates the strength of local CI claims.

*The corresponding author is Beishui Liao (e-mail: baiseliao@zju.edu.cn).

- **A fixed-point dialectical semantics for skeleton revision.** Edge acceptability is revised through finite-horizon or convergent updates, yielding a stable skeleton that reflects aggregated evidential conflicts rather than greedy deletions.

The remainder of this paper is organised as follows. Section 2 reviews causal graphs and background on quantitative argumentation. Section 3 presents QACD and its dialectical semantics. Section 4 evaluates the approach empirically, followed by related work and conclusions.

A long version of this paper, including supplementary material and additional implementation details, can be found in (Wei, Chen, and Liao 2026).

2 Preliminaries

This section reviews the causal-graphical and argumentation-theoretic notions used throughout the paper and fixes notations.

2.1 Causal Graphs

We assume that the underlying causal structure over observed variables can be represented by a DAG (Pearl 1988). Let $G = (\mathbf{V}, \mathbf{E})$ be a DAG, where $\mathbf{V} = \{V_1, \dots, V_n\}$ is a set of random variables and $\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$ is a set of directed edges. An edge $X \rightarrow Y$ indicates that X is a (direct) cause of Y in G .

For a node $X \in \mathbf{V}$, a (possibly non-simple) path between two nodes is a sequence of adjacent nodes, where edge directions are ignored. A node M on a path is a *collider* on that path if the two incident edges are oriented $X \rightarrow M \leftarrow Y$; otherwise M is a *non-collider*.

Definition 1 (d-separation). Let $X, Y \in \mathbf{V}$ and $\mathbf{Z} \subseteq \mathbf{V} \setminus \{X, Y\}$. A path p between X and Y is *active* (or *d-connecting*) given $\mathbf{Z}$ if:

1. every non-collider on p is not in $\mathbf{Z}$; and
2. every collider on p is in $\mathbf{Z}$ or has a descendant in $\mathbf{Z}$.

If there exists no active path between X and Y given $\mathbf{Z}$, then X and Y are *d-separated* by $\mathbf{Z}$ in G , denoted $X \perp\!\!\!\perp_G Y \mid \mathbf{Z}$.

Graphical separation is linked to probabilistic conditional independence through standard assumptions.

Definition 2 (Causal Markov Assumption). Let P be the joint probability distribution over $\mathbf{V}$ and G be a causal DAG. The distribution P satisfies the *Causal Markov Assumption* with respect to G if every d-separation relation in G implies a conditional independence in P . That is, for all disjoint sets of variables $X, Y \in \mathbf{V}$ and $\mathbf{Z} \subseteq \mathbf{V} \setminus \{X, Y\}$:

$$X \perp\!\!\!\perp_G Y \mid \mathbf{Z} \implies X \perp\!\!\!\perp_P Y \mid \mathbf{Z}. \quad (1)$$

While the Markov assumption allows us to read independencies from the graph, we also require the converse to infer the graph structure from observed independencies.

Definition 3 (Faithfulness Assumption). The distribution P is said to be *faithful* to the DAG G if every conditional independence relation in P implies a d-separation in G . That is, for all distinct variables $X, Y \in \mathbf{V}$ and all $\mathbf{Z} \subseteq \mathbf{V} \setminus \{X, Y\}$:

$$X \perp\!\!\!\perp_P Y \mid \mathbf{Z} \implies X \perp\!\!\!\perp_G Y \mid \mathbf{Z}. \quad (2)$$

While we assume faithfulness as a background idealization, our framework is explicitly designed to operate under finite-sample regimes where CI test outcomes may violate faithfulness due to statistical error.

Definition 4 (Markov Equivalence). Two DAGs G_1 and G_2 over $\mathbf{V}$ are *Markov equivalent* if they entail the same set of d-separation relations. A necessary and sufficient condition for Markov equivalence is that G_1 and G_2 possess the same *skeleton* (adjacency structure ignoring direction) and the same set of *v-structures* (unshielded colliders of the form $X \rightarrow Z \leftarrow Y$ where X and Y are not adjacent). The set of all DAGs Markov equivalent to a graph G constitutes its *Markov Equivalence Class* (MEC), denoted by $[G]$.

Since observational data (under the Markov and Faithfulness assumptions) only allows us to identify the MEC rather than a unique DAG, we utilise a graphical object to represent the entire class.

Definition 5 (Completed Partially Directed Acyclic Graph). A Markov equivalence class $[G]$ can be uniquely represented by a CPDAG G^* , which has the same skeleton as any DAG in $[G]$, and whose edges are directed if and only if they are compelled across the class.

Our goal is to recover the target CPDAG G^* from finite samples (Anand et al. 2025). Constraint-based methods use CI tests to remove adjacencies and separating sets to orient edges; under finite samples, CI errors motivate the non-monotonic treatment developed in the next section.

Orientation Rules (Meek). Given an undirected skeleton and a collection of separating sets, standard constraint-based causal discovery procedures apply a set of sound graphical orientation rules, known as *Meek rules* (Meek 1995), to orient all compelled edges and obtain the unique CPDAG representing the corresponding MEC. These rules are deterministic, purely graphical, and do not introduce additional statistical assumptions.

Conditional Independence Tests and p -values. Throughout this paper, CI relations are assessed using standard statistical tests, each returning a p -value. We interpret the p -value not as a hard decision criterion, but as a graded indicator of evidential strength for a CI claim: larger p -values correspond to stronger statistical support for conditional independence (Jabbari et al. 2017).

2.2 Quantitative Argumentation Frameworks

To situate our method more clearly within formal argumentation, we briefly recall the attack-only quantitative setting used in this paper. We require only a minimal set of notions: arguments, attacks, weights, and gradual acceptability.

Definition 6 (Quantitative Argumentation Framework). In this paper, we use a lightweight attack-only quantitative argumentation framework (QAF) of the form

$$\mathcal{Q} = \langle A, R, w, \rho \rangle,$$

where A is a finite set of arguments, $R \subseteq A \times A$ is an attack relation, $w : A \rightarrow [0, 1]$ assigns intrinsic strengths to arguments, and $\rho : R \rightarrow [0, 1]$ assigns strengths to attacks. If $(a, b) \in R$, then a attacks b with strength $\rho(a, b)$.

Definition 7 (Gradual Semantics). Let $\mathcal{Q} = \langle A, R, w, \rho \rangle$ be a QAF. A gradual semantics assigns to each argument $a \in A$ an acceptability degree $s(a) \in [0, 1]$, rather than a binary status such as accepted or rejected. Such a semantics may be specified by an operator $\Gamma : [0, 1]^{|A|} \rightarrow [0, 1]^{|A|}$, whose iterates determine the acceptability degrees of arguments under the influence of weighted attacks. Intuitively, an argument remains acceptable to the extent that it withstands the strength of its attackers.

The framework proposed in this paper instantiates this attack-only and gradual setting for causal discovery. Section 3 treats CI statements as graded attackers of adjacency hypotheses and defines a state-dependent acceptability update over candidate edges. In particular, the effective attack strengths in QACD may depend on the current graph state through witness-based propagation.

3 Methodology

We now present QACD, which learns causal skeletons by dialectically aggregating uncertain CI evidence. As shown in Figure 1, Phase I constructs a permissive candidate graph and collects CI evidence, while Phase II iteratively revises edge acceptability under direct and witness-propagated attacks before thresholding and orientation. Algorithm 1 implements the QACD semantics.

3.1 Problem Setup and Notation

Let $\mathcal{V}$ be a finite set of observed variables generated by an unknown causal DAG $\mathcal{G}^*$ satisfying the Causal Markov and Faithfulness assumptions. Given i.i.d. samples, a CI testing procedure returns a statistic (e.g., a p -value) for statements of the form $X \perp\!\!\!\perp Y \mid \mathbf{Z}$.

Our goal is to recover the MEC of $\mathcal{G}^*$, represented as a CPDAG, while avoiding brittle monotonic decisions caused by finite-sample CI errors.

Argumentative Representation. We write $X-Y$ for an undirected adjacency and $X \rightarrow Y$ for a directed edge. QACD maintains a symmetric *acceptability matrix* $\mathbf{S} \in [0, 1]^{|V| \times |V|}$, where S_{XY} denotes how acceptable the hypothesis “ X is adjacent to Y ” currently is. Values close to 1 indicate strong support; values near 0 indicate rejection. We refer to S_{XY} simply as the acceptability of the edge hypothesis e_{XY} .

QACD can be viewed as a two-sorted quantitative argumentative system whose state space consists of edge hypotheses and representative CI-derived independence arguments. Edge hypotheses e_{XY} are the claims whose acceptability is iteratively updated, whereas CI arguments τ_{XY} carry fixed initial strengths $\sigma_0(\tau_{XY})$ and act as graded attackers. Witness-based propagation is not introduced as a separate argument node, but as a state-dependent operator that modulates how CI-derived attack pressure extends to adjacent edge hypotheses through the current graph structure. We distinguish three classes of argumentative objects:

1. **Edge Hypotheses** (e_{XY}): The target claims to be justified or defeated.

2. **Independence Arguments** (τ): Statistical findings that attack edge hypotheses. An independence argument τ_{XY} is a tuple defined by a conditioning set $\mathbf{Z}_\tau \subseteq \mathcal{V} \setminus \{X, Y\}$ and a test outcome asserting $X \perp\!\!\!\perp Y \mid \mathbf{Z}_\tau$. We assign each τ_{XY} a *base strength* $\sigma_0(\tau_{XY}) \in [0, 1]$ derived from its test statistic, treating it as a defeasible reason to reject e_{XY} .
3. **Structural Modifiers**: Context-dependent modifiers induced by the evolving graph structure, specifically by witness paths. These modifiers do not constitute separate argument nodes; instead, they modulate how CI-derived attack pressure is propagated to structurally implicated edge hypotheses when independence evidence conflicts with strong local connectivity.

In the subsequent methodology, we detail how specific arguments are instantiated (Phase I) and dialectically aggregated (Phase II).

3.2 Phase I: Candidate Graph Generation

Phase I constructs a permissive candidate skeleton to avoid premature deletions.

For each pair $\{X, Y\}$, we compute

$$p_{XY}^{\min} = \min_{\mathbf{Z} \subseteq \mathcal{V} \setminus \{X, Y\}, |\mathbf{Z}| \leq K} p(X \perp\!\!\!\perp Y \mid \mathbf{Z}), \quad (3)$$

where K is the maximum conditioning-set size used in Phase I, and include $X-Y$ in the candidate graph G_{cand} whenever $p_{XY}^{\min} < \delta_{\text{cand}}$ (candidate threshold). This rule is intentionally conservative: if any conditioning set suggests dependence, the edge is retained as a viable hypothesis.

Each variable pair additionally yields a single representative CI statement τ_{XY} , corresponding to the conditioning set that provides the strongest empirical support for independence, as quantified by the largest associated p -value. In Phase II, τ_{XY} is treated as an *independence argument*, whose evidential strength is determined by the value of its corresponding test statistic.

For each unordered pair (X, Y) , QACD retains a single representative CI argument, namely the statement associated with the largest observed p -value among the tested conditioning sets. Accordingly, the global CI pool is denoted by $\mathcal{T} = \{\tau_{XY} : X < Y\}$. Thus, the update uses one representative independence argument per unordered pair.

3.3 Phase II: Dialectical Aggregation via Quantitative Argumentation

Having constructed the candidate skeleton, we now resolve conflicts among edges by instantiating a state-dependent QAF. In this framework, the set of arguments is fixed, but the effective attack relations depend on the evolving acceptability state of the graph. Consequently, the influence of evidence is inherently context-sensitive: the same CI statement may have different effects depending on the surrounding structure.

Phase II uses a gradual semantics in which edge hypotheses are revised under weighted CI attacks and context-dependent witness-based propagated attenuation.

Figure 2 illustrates this interaction between adjacency hypotheses, CI-based rebuttals, and witness-induced propagated attenuation.




Figure 1: Overview of QACD. Phase I builds a permissive candidate graph. Phase II aggregates CI evidence via quantitative argumentation with direct attacks and witness-based propagation, iteratively updating edge acceptability before thresholding and orientation.




Figure 2: A representative CI argument τ_{XY} directly attenuates the target edge e_{XY} and, when a witness connection $X-W-Y$ with $W \notin Z_\tau$ is present, also propagates attenuation to the bridging edges e_{XW} and e_{WY} in proportion to the witness connectivity c_W .

Example 1 (Running example). Suppose the current candidate graph contains the triangle $X-Y$, $X-W$, and $W-Y$. Assume that Phase I retains a representative CI statement $\tau_{XY} : X \perp\!\!\!\perp Y \mid Z_\tau$ with a large p -value, and that $W \notin Z_\tau$. Then τ_{XY} provides strong direct pressure against the target edge hypothesis e_{XY} . If the current acceptability scores of e_{XW} and e_{WY} are still high, W becomes an eligible witness, so part of the same pressure is propagated to the two bridging edges. In this way, QACD does not only revise the disputed edge $X-Y$ in isolation; it also weakens the local configuration that sustains the conflicting pattern.

Let τ_{XY} denote the representative CI statement retained for (X, Y) , with base strength $\sigma_0(\tau_{XY}) \in [0, 1]$. This representative independence statement is treated as an attack argument against e_{XY} , weakening its acceptability. Beyond this direct attack, QACD introduces witness-based propagated attenuation (formalized in Definition 8), which allows strong independence evidence to weaken structurally implicated bridging edges under inconsistent CI evidence.

Formally, the graded semantics is computed via an iterative dialectical update:

$$\mathbf{S}^{(t+1)} = \mathbf{S}^{(t)} \odot \mathbf{A}(\mathbf{S}^{(t)}), \quad (4)$$

where $\odot$ is elementwise multiplication, $\mathbf{S}^{(t)}$ denotes the current acceptability and $\mathbf{A}(\mathbf{S}^{(t)}) \in [0, 1]^{|\mathcal{V}| \times |\mathcal{V}|}$ is an atten-

uation matrix determined by (i) direct attacks induced by independence arguments and (ii) connectivity-mediated propagated attenuation triggered by short witness connections. We stop after $T_{\max}$ iterations or when $\|\mathbf{S}^{(t+1)} - \mathbf{S}^{(t)}\|_\infty$ is below a small tolerance. That is, an edge hypothesis retains its current acceptability unless it is weakened by an applicable defeasible influence. The iteration continues until the acceptability state stabilizes up to the prescribed stopping criterion.

We now describe the quantitative argumentation semantics used to revise $\mathbf{S}$ in the presence of conflicting CI evidence.

For any fixed acceptability state $S \in [0, 1]^{|\mathcal{V}| \times |\mathcal{V}|}$, QACD induces a state-dependent QAF

$$\mathcal{Q}(S) = \langle \mathcal{A}, R_S, w, \rho_S \rangle,$$

where $\mathcal{A} = \mathcal{E} \cup \mathcal{T}$, $\mathcal{E} = \{e_{XY} : X < Y\}$, and $\mathcal{T} = \{\tau_{XY} : X < Y\}$. Here e_{XY} is the edge-hypothesis argument for the adjacency claim between X and Y , and τ_{XY} is the representative CI-derived independence argument for the same unordered pair retained in Phase I. The weight function is defined by $w(e_{XY}) = S_{XY}^{(0)}$ and $w(\tau_{XY}) = \sigma_0(\tau_{XY})$.

The attack relation R_S contains each direct attack (τ_{XY}, e_{XY}) . Moreover, whenever $W \notin Z_\tau$ is an eligible witness for (X, Y) at state S , R_S also contains the propagated attacks (τ_{XY}, e_{XW}) and (τ_{XY}, e_{WY}) . The corresponding effective attack strengths are

$$\rho_S(\tau_{XY}, e_{XY}) = \lambda \sigma_0(\tau_{XY})$$

for direct attacks, and

$$\rho_S(\tau_{XY}, e_{XW}) = \rho_S(\tau_{XY}, e_{WY}) = \frac{\lambda}{2} \sigma_0(\tau_{XY}) S_{XW} S_{WY}$$

for propagated attacks.

Independence arguments and direct attacks. Each independence argument τ_{XY} provides a weighted direct attack on the edge hypothesis e_{XY} . In the current implementation, its CI-test p -value is mapped to a base strength

$$\sigma_0(\tau_{XY}) = \begin{cases} 0, & p(\tau_{XY}) < \alpha, \\ \frac{p(\tau_{XY}) - \alpha}{1 - \alpha}, & p(\tau_{XY}) \geq \alpha, \end{cases} \quad (5)$$

so that only non-rejected CI claims contribute positive independence strength, and among them larger p -values yield stronger defeasible support for independence.

This strength is interpreted as a direct rebuttal on the edge hypothesis e_{XY} . Since the current implementation retains a single representative CI argument for each unordered pair (X, Y) , the direct attenuation term is

$$A_{XY}^{\text{dir}} = 1 - \lambda \sigma_0(\tau_{XY}). \quad (6)$$

The parameter λ controls how aggressively defeasible evidence propagates through the structure. Importantly, this aggregation does not represent a vote or a probability, but a cumulative weakening of justification. This form yields a bounded, order-independent attenuation semantics. When multiple applicable direct or propagated factors affect the same edge hypothesis, they combine multiplicatively to weaken its acceptability.

Aggregation rationale. Eq. (6) implements a bounded, proportional attenuation of the edge hypothesis by its representative CI argument. This form preserves acceptability within $[0, 1]$ and composes naturally with the iterative update in Eq. (4).

In argumentation theory, a rebuttal directly attacks a claim, whereas an undercut challenges the applicability of such an attack without asserting the opposite conclusion (Amgoud, Maudet, and Parsons 2000). In the causal discovery setting, strong CI evidence may therefore undercut not the edge hypothesis itself, but the structural relations that sustain conflicting connectivity patterns.

Definition 8 (Witness-Based Propagated Attenuation via Witness Connectivity). Fix an independence argument τ_{XY} with conditioning set $\mathbf{Z}_\tau$ and base strength $\sigma_0(\tau_{XY})$. At iteration t , a node $W \in \mathcal{V} \setminus \{X, Y\}$ is an *eligible witness* for (X, Y) with respect to τ if (i) it forms a length-2 connection $X-W-Y$ in the current candidate structure (equivalently, $S_{XW}^{(t)} > 0$ and $S_{WY}^{(t)} > 0$), and (ii) $W \notin \mathbf{Z}_\tau$. Such a witness induces a *connectivity support* for $X-Y$ measured by

$$u_\tau^{(t)}(W) = S_{XW}^{(t)} \cdot S_{WY}^{(t)}. \quad (7)$$

Intuitively, when a strong witness connection $X-W-Y$ remains acceptable, the CI argument τ_{XY} is not interpreted in isolation. Instead, its influence is allowed to propagate to the bridging edge hypotheses e_{XW} and e_{WY} in proportion to both the argument strength and the current witness connectivity. This state-dependent propagation weakens the local structural support underlying the competing witness pattern, as illustrated in Figure 2.

Design Rationale for Witness-Based Propagated Attenuation. Restricting witnesses to length-2 connections is a principled and computationally grounded design choice. Such paths are the minimal non-trivial connectivity units in causal graphs and align with the local reasoning patterns exploited by constraint-based methods, where CI errors can cascade through short neighborhoods. When strong independence evidence conflicts with a strongly supported length-2 connectivity pattern outside its conditioning set, the inconsistency can often be resolved by weakening one of the bridging relations, thereby restoring local coherence under

Algorithm 1 QACD: Quantitative Argumentation for Causal Discovery

```

1: Input: data  $\mathcal{D}$  over variables  $\mathcal{V}$ ; CI test procedure; max
   conditioning size  $K$ ; candidate threshold  $\delta_{\text{cand}}$ ; final
   threshold  $\delta_0$ ; aggressiveness  $\lambda$ ; max iterations  $T_{\text{max}}$ ;
   tolerance  $\varepsilon$ 
2: Output: estimated CPDAG  $\hat{G}^*$ 
3: Phase I: Candidate Generation
4: Initialize  $G_{\text{cand}} \leftarrow$  empty graph
5: Initialize CI-statement pool  $\mathcal{T} \leftarrow \emptyset$ 
6: for all unordered pairs  $\{X, Y\} \subseteq \mathcal{V}$  do
7:   Compute  $p_{XY}^{\min}$  using Eq. (3)
8:   Let  $\mathbf{Z}_{XY}^* \in \arg \max p(X \perp\!\!\!\perp Y \mid \mathbf{Z})$  and add  $\tau_{XY}$  to
      $\mathcal{T}$ 
9:   if  $p_{XY}^{\min} < \delta_{\text{cand}}$  then
10:     Add undirected edge  $X-Y$  to  $G_{\text{cand}}$ 
11:   end if
12: end for
13: Initialize acceptability matrix  $\mathbf{S}^{(0)}$  from  $G_{\text{cand}}$ 
14: Phase II: Dialectical Aggregation
15: for  $t = 0$  to  $T_{\text{max}} - 1$  do
16:   Initialize attenuation matrix  $\mathbf{A}^{(t)} \leftarrow \mathbf{1}$ 
17:   for all CI statements  $\tau_{XY}$  in the pool  $\mathcal{T}$  do
18:     Compute base strength  $\sigma_0(\tau_{XY})$  using Eq. (5)
19:     Update direct attenuation on  $(X, Y)$  using Eq. (6)
20:     for all witness nodes  $W$  such that  $X-W-Y$  in
        $G_{\text{cand}}$  and  $W \notin \mathbf{Z}_\tau$  do
21:       Update propagated attenuation on  $(X, W)$  and
          $(W, Y)$  using Eq. (8)
22:     end for
23:   end for
24:   Update acceptability:  $\mathbf{S}^{(t+1)} \leftarrow \mathbf{S}^{(t)} \odot \mathbf{A}^{(t)}$ 
25:   if  $\|\mathbf{S}^{(t+1)} - \mathbf{S}^{(t)}\|_\infty < \varepsilon$  then
26:     break
27:   end if
28: end for
29: Skeleton and CPDAG Recovery
30: Threshold  $\mathbf{S}^{(T)}$  at  $\delta_0$  to obtain skeleton  $\hat{G}_{\text{skel}}$ 
31: Apply standard orientation rules to obtain CPDAG  $\hat{G}^*$ 
32: return  $\hat{G}^*$ 

```

inconsistent CI evidence. Longer paths can be viewed as compositions of length-2 units and are therefore indirectly captured through iterative propagation; extending propagation to longer paths is possible but computationally more expensive.

Although inspired by context-sensitive undercut-style reasoning, witness connectivity is used here as a state-dependent structural modifier that propagates CI-derived attack pressure to adjacent bridging edges. In the present formulation, witness connectivity serves as a state-dependent propagation signal: when a strong witness pattern remains acceptable, part of the influence of the CI-derived attack on e_{XY} is redistributed to the adjacent bridging edges.

Propagation operator. The strength of a witness W under the current acceptability is $u_\tau^{(t)}(W)$. The propagated atten-

uation applied to each bridging edge induced by τ_{XY} and witness W is defined as

$$\begin{aligned} A_{XW \leftarrow \tau_{XY}}^{\text{prop}}(W) &= A_{WY \leftarrow \tau_{XY}}^{\text{prop}}(W) \\ &= 1 - \frac{\lambda}{2} \sigma_0(\tau_{XY}) u_{\tau}^{(t)}(W), \end{aligned} \quad (8)$$

where the factor $\frac{1}{2}$ is a symmetric normalization choice that splits the propagated attenuation equally across the two bridging edges. It is a simple heuristic design choice that preserves symmetry between e_{XW} and e_{WY} . We set $\lambda \in [0, 1]$ to ensure $A_{e \leftarrow \tau}^{\text{prop}}(W) \in [0, 1]$. Direct and propagated attenuation factors are combined multiplicatively in the attenuation matrix $\mathbf{A}(\mathbf{S}^{(t)})$; if no factor applies to an edge, its attenuation is set to 1.

For clarity, the overall attenuation applied to an edge hypothesis e at iteration t is the product of all direct and propagated factors that apply to that edge:

$$A_e(\mathbf{S}^{(t)}) = \prod_{\tau \in \mathcal{T}_{\text{dir}}(e)} A_{e \leftarrow \tau}^{\text{dir}} \cdot \prod_{(\tau, W) \in \mathcal{T}_{\text{prop}}(e)} A_{e \leftarrow \tau}^{\text{prop}}(W),$$

with the convention that an empty product equals 1. Here, $\mathcal{T}_{\text{dir}}(e)$ denotes the set of representative CI arguments that directly attenuate e , and $\mathcal{T}_{\text{prop}}(e)$ denotes the set of witness-indexed propagated attenuation terms that apply to e at iteration t . Equation (4) is therefore understood entrywise as $S_e^{(t+1)} = S_e^{(t)} \cdot A_e(\mathbf{S}^{(t)})$ for each edge hypothesis e .

Basic properties of the iterative semantics follow directly from the form of the attenuation factors.

Theorem 1 (Boundedness). If $\mathbf{S}^{(0)} \in [0, 1]^{|V| \times |V|}$, then $\mathbf{S}^{(t)} \in [0, 1]^{|V| \times |V|}$ for all $t \geq 0$.

Proof. By construction, $\mathbf{A}(\mathbf{S}^{(t)})$ is obtained by multiplying factors of the form $(1 - \lambda \sigma_0(\tau))$ and $(1 - \frac{\lambda}{2} \sigma_0(\tau) u_{\tau}^{(t)}(W))$, each of which lies in $[0, 1]$ when $\lambda, \sigma_0(\tau), u_{\tau}^{(t)}(W) \in [0, 1]$. Equation (4) updates $\mathbf{S}^{(t+1)}$ by elementwise multiplication of $\mathbf{S}^{(t)}$ with $\mathbf{A}(\mathbf{S}^{(t)})$, so every entry remains within $[0, 1]$. $\square$

Theorem 2 (Monotonicity). For every unordered pair $\{X, Y\}$ and every $t \geq 0$, we have $S_{XY}^{(t+1)} \leq S_{XY}^{(t)}$.

Proof. Each entry of $\mathbf{A}(\mathbf{S}^{(t)})$ lies in $[0, 1]$, hence Eq. (4) gives $S_{XY}^{(t+1)} = S_{XY}^{(t)} \cdot A_{XY}(\mathbf{S}^{(t)}) \leq S_{XY}^{(t)}$. $\square$

Theorem 3 (Pointwise convergence). For every unordered pair $\{X, Y\}$, the sequence $\{S_{XY}^{(t)}\}_{t \geq 0}$ converges to a limit $S_{XY}^{(*)} \in [0, 1]$.

Proof. By Theorem 1 and Theorem 2, each scalar sequence $\{S_{XY}^{(t)}\}_{t \geq 0}$ is bounded below by 0 and monotone non-increasing, hence it converges in $\mathbb{R}$. Collecting the limits entrywise yields $\mathbf{S}^{(*)} \in [0, 1]^{|V| \times |V|}$. $\square$

Proposition 1 (Fixed-point characterization). Let $F : [0, 1]^{|V| \times |V|} \rightarrow [0, 1]^{|V| \times |V|}$ be defined by $F(S) = S \odot A(S)$. If witness propagation is evaluated over the fixed set of length-2 witness triples in the candidate graph, with attenuation factor $1 - \frac{\lambda}{2} \sigma_0(\tau_{XY}) S_{XW} S_{WY}$, then each entry of

$A(S)$ is a finite product of continuous factors in entries of S . Hence F is continuous on $[0, 1]^{|V| \times |V|}$. By Theorem 3, $S^{(t)} \rightarrow S^{(*)}$ entrywise, and therefore

$$S^{(*)} = \lim_{t \rightarrow \infty} S^{(t+1)} = \lim_{t \rightarrow \infty} F(S^{(t)}) = F(S^{(*)}).$$

Thus the limit state $S^{(*)}$ is a fixed point of the QACD revision operator.

Stability of the Iterative Revision. By Theorem 3 and Proposition 1, the infinite revision process converges pointwise to a fixed point of the QACD revision operator under the chosen initialization. In practice, QACD uses a finite iteration budget T_{max} and returns the last iterate once the stopping criterion is met. Thus, the output should be understood as a finite-horizon stable revision of edge acceptability under the proposed quantitative argumentation semantics, rather than as a claim that all evidential conflicts are resolved by a non-trivial equilibrium.

Unlike greedy constraint-based deletion, this iterative semantics allows the influence of each CI statement to be continuously re-evaluated as the surrounding graph structure evolves, enabling non-greedy conflict resolution under inconsistent statistical evidence.

Complexity. Let $n = |V|$ denote the number of variables and let K be the maximum conditioning-set size used in Phase I. Phase I performs CI tests for each unordered pair (X, Y) over conditioning sets of size at most K , yielding a worst-case cost of $O(n^2 \sum_{k=0}^K \binom{n}{k})$ CI tests, as in standard constraint-based pipelines. In addition, Phase I stores at most one representative independence statement τ_{XY} per variable pair, resulting in a CI-statement pool $\mathcal{T}$ of size $O(n^2)$.

Phase II iterates a fixed number T_{max} of dialectical updates. For each $\tau_{XY} \in \mathcal{T}$, eligible witnesses are common neighbors of X and Y , at most $\min\{\deg(X), \deg(Y)\}$ in G_{cand} . With adjacency-list intersection, each iteration costs $O(\sum_{X < Y} \min\{\deg(X), \deg(Y)\})$, bounded by $O(n^3)$ in the worst case and close to $O(n^2 \bar{d})$ in sparse graphs, where $\bar{d}$ is the average degree of G_{cand} . Thus, Phase II adds only a controlled polynomial-time aggregation step.

3.4 Graph Recovery

After T_{max} iterations (or convergence), we threshold the acceptability matrix to produce a skeleton:

$$X - Y \in \hat{G}_{\text{skel}} \iff S_{XY}^{(T_{\text{max}})} \geq \delta_0. \quad (9)$$

This step converts the graded acceptability semantics into a crisp adjacency decision, separating highly supported edges from those dialectically rejected.

We then orient edges by identifying unshielded colliders using the separating sets obtained during Phase I CI testing. The resulting partially oriented graph is closed under the complete set of Meek orientation rules (Meek 1995), yielding the estimated CPDAG $\hat{G}^*$. This recovery procedure mirrors standard constraint-based pipelines, ensuring comparability with classical methods.

Category	Dataset	$ \mathcal{V} $	$ \mathcal{E} $	P
Small	EARTHQUAKE	5	4	10
	SURVEY	6	6	21
	ASIA	8	8	18
Medium	CHILD	20	25	230
	INSURANCE	27	52	1008
	WATER	32	66	10083
Large	HAILFINDER	56	66	2656
	WIN95PTS	76	112	574

Table 1: Benchmark dataset statistics. $|\mathcal{V}|$, $|\mathcal{E}|$, and P denote the numbers of nodes, edges, and parameters, respectively.

4 Experiments

From an argumentation perspective, the experiments evaluate whether dialectical aggregation of CI evidence improves structural coherence under finite-sample uncertainty. In particular, we examine whether treating CI outcomes as defeasible arguments mitigates error propagation in skeleton discovery and leads to more causally valid graphs.

We evaluate QACD on a suite of benchmark Bayesian networks to assess (i) skeleton recovery accuracy and (ii) downstream causal validity. We compare against representative constraint-based, hybrid, and argumentation-based baselines.

4.1 Experimental Setup

We evaluate QACD on standard benchmark networks and compare it against representative causal discovery methods under a unified experimental protocol.

Datasets. We use eight standard Bayesian networks categorized by size. Statistics are summarized in Table 1.¹

Baselines. We compare QACD against representative methods from three families: **constraint-based** (PC-stable (Spirtes and Glymour 1991) and MPC (Colombo, Maathuis, and others 2014)), **hybrid** (MMHC (Tsamardinos, Brown, and Aliferis 2006)), and **argumentation-based** (ABAPC (Russo, Rapberger, and Toni 2024)). ABAPC is evaluated only on small networks due to runtime constraints. For the WATER dataset, MMHC is omitted due to deterministic relationships in the data.

Protocol. All benchmarks are discrete Bayesian networks. CI testing uses the Pearson χ^2 test with significance level $\alpha = 0.05$. For each network, we sample $N = 5000$ i.i.d. observations from the ground-truth model and repeat experiments over 50 random seeds, reporting Mean $\pm$ Std. All methods are evaluated under the same underlying CI-testing protocol and significance level. The difference is that classical constraint-based baselines use thresholded CI decisions, whereas QACD additionally exploits the full p-value output as graded defeasible evidence; this is part of the modeling design of QACD rather than a difference in the underlying statistical test.

QACD Configuration. Unless stated otherwise, QACD uses fixed hyperparameters across datasets: conditioning-set bound $K = 3$, $\delta_{\text{cand}} = 0.05$, $\delta_0 = 0.05$, $\lambda = 0.5$, and

$T_{\text{max}} = 20$ dialectical iterations. These hyperparameters are fixed across all datasets and are not tuned separately for individual networks.

Metrics. We report (1) Skeleton F_1 score, (2) Normalized SHD (NSHD), and (3) Normalized SID (NSID) (Peters and Bühlmann 2015). NSHD and NSID are normalized by the number of true edges and therefore need not lie in $[0, 1]$. Lower NSHD and NSID indicate better causal validity.

4.2 Structural Accuracy and Causal Validity Across Benchmarks

Tables 2 and 3 summarize results across small, medium, and large networks. Overall, QACD consistently produces structures with strong causal validity (low NSHD/NSID) while maintaining competitive skeleton accuracy across datasets of varying complexity.

QACD shows the greatest efficacy on networks where CI evidence is noisy or inconsistent. On INSURANCE, QACD achieves the highest F_1 and the lowest NSHD while substantially tightening NSID bounds relative to PC, MPC, and MMHC. Similar behavior appears on SURVEY and HAILFINDER, where QACD improves NSHD and consistently reduces NSID_{high}.

Across several datasets (e.g., ASIA and CHILD), QACD exhibits a recurring trade-off between recall and structural coherence. While aggressive baselines such as MPC or MMHC sometimes achieve slightly higher skeleton F_1 , QACD produces lower NSHD and tighter NSID bounds, indicating superior interventional reliability.

On datasets with deterministic relationships (e.g., WATER), QACD remains stable and improves the lower SID estimate while remaining comparable on the upper SID estimate, even when greedy methods become sensitive to statistical degeneracy. Because argument strengths are attenuated rather than globally optimized, the method avoids instability caused by degenerate sufficient statistics.

On low-complexity networks (e.g., EARTHQUAKE), PC, MMHC, and QACD all recover the ground-truth structure, confirming that dialectical aggregation maintains robustness when CI evidence is already consistent. On WIN95PTS, QACD remains competitive but does not dominate, reflecting the expected trade-off of conservative filtering in regimes where CI evidence is highly reliable.

Overall, the results indicate that QACD is most beneficial when CI evidence is noisy or inconsistent. By revising edge acceptability through connectivity-mediated propagation, dialectical aggregation improves global structural coherence and reduces interventional error, while remaining competitive in skeleton recovery across datasets of varying sizes.

4.3 Ablation: Effect of Dialectical Aggregation

To isolate the impact of the dialectical aggregation phase, we compare full QACD with the initial candidate skeleton before aggregation, i.e., the variant with $T_{\text{max}} = 0$, as shown in Table 4. Across datasets, removing the iterative aggregation phase substantially degrades both skeleton quality and causal validity. The effect is most pronounced on networks where short-range witness connectivity provides reliable structural

¹<https://www.bnlearn.com/bnrepository/>.

Dataset	Method	$F_1 \uparrow$	NSHD $\downarrow$	NSID _{low} $\downarrow$	NSID _{high} $\downarrow$
ASIA	PC	0.8724 ± 0.0341	0.5750 ± 0.1118	1.5500 ± 0.7984	2.4250 ± 0.7984
	MPC	0.9340 ± 0.0344	0.3875 ± 0.0379	0.7875 ± 0.7939	1.4050 ± 0.3733
	MMHC	0.7487 ± 0.0459	0.7500 ± 0.0000	1.8250 ± 0.1677	3.0750 ± 0.1677
	ABAPC	0.5000 ± 0.1800	0.6550 ± 0.2825	1.4650 ± 0.8488	4.1900 ± 0.9900
	QACD(ours)	0.9181 ± 0.0341	0.3750 ± 0.0000	0.2250 ± 0.1630	1.1000 ± 0.1630
EARTHQUAKE	PC	1.0000 ± 0.0000	0.0000 ± 0.0000	0.0000 ± 0.0000	0.0000 ± 0.0000
	MPC	0.9978 ± 0.0157	0.0050 ± 0.0354	0.0000 ± 0.0000	0.0000 ± 0.0000
	MMHC	1.0000 ± 0.0000	0.0000 ± 0.0000	0.0000 ± 0.0000	0.0000 ± 0.0000
	ABAPC	0.9400 ± 0.0600	0.1400 ± 0.1450	0.0000 ± 0.0000	2.3300 ± 1.1350
	QACD(ours)	1.0000 ± 0.0000	0.0000 ± 0.0000	0.0000 ± 0.0000	0.0000 ± 0.0000
SURVEY	PC	0.6489 ± 0.1280	0.5667 ± 0.1491	2.3000 ± 0.0745	2.3000 ± 0.0745
	MPC	0.6196 ± 0.1491	0.6167 ± 0.1438	2.4500 ± 0.2189	2.4833 ± 0.2901
	MMHC	0.8218 ± 0.0488	0.9000 ± 0.2236	1.3000 ± 0.4314	3.2000 ± 0.9531
	ABAPC	0.4200 ± 0.1100	0.8030 ± 0.1830	2.3700 ± 0.6080	2.7360 ± 0.4610
	QACD(ours)	0.7588 ± 0.0604	0.4333 ± 0.1491	1.9333 ± 0.3651	1.9333 ± 0.3651

Table 2: **Small Networks Results.** Comparison of Skeleton F_1 , Normalized SHD (NSHD), and Normalized SID (NSID). All values are reported as **Mean $\pm$ Std.** **Bold** indicates the best performance (higher is better for F_1 , lower is better for NSHD/NSID).

cues, confirming the role of dialectical aggregation in stabilizing skeleton discovery under noisy CI evidence.

4.4 Impact of Sample Size

We analyze the sensitivity of QACD and PC to sample size on the INSURANCE network ($N \in \{500, 1000, 2000, 5000\}$; Figure 3), reporting Mean $\pm$ Std over 5 runs. Using PC as the canonical constraint-based baseline isolates the contribution of dialectical aggregation under identical CI inputs. QACD outperforms PC at $N = 500$ (0.5443 ± 0.0172 vs. 0.5114 ± 0.0520), $N = 1000$ (0.6383 ± 0.0275 vs. 0.5414 ± 0.0536), and $N = 2000$ (0.6968 ± 0.0388 vs. 0.6420 ± 0.0362); at $N = 5000$, the gap narrows (0.7361 ± 0.0301 vs. 0.7267 ± 0.0362). This suggests that dialectical aggregation is most helpful in low-sample regimes and becomes less advantageous as CI decisions stabilize with increasing N .

QACD is intentionally conservative: dialectical aggregation suppresses weakly supported edges, which can reduce recall in dense graphs. In settings where CI evidence is already highly consistent or sample sizes are very large, the advantage of non-monotonic conflict resolution diminishes. These cases reflect a design trade-off rather than instability of the proposed method.

4.5 Discussion

The preceding results reveal a consistent pattern rather than a uniform dominance claim. The results suggest that QACD is most beneficial when CI evidence is noisy or mutually inconsistent. On datasets such as INSURANCE, SURVEY, and HAILFINDER, dialectical aggregation improves NSHD and NSID while maintaining competitive skeleton F_1 . This is consistent with treating CI outcomes as defeasible arguments: weak or conflicting CI claims are attenuated through connectivity-mediated revision, reducing the propagation of local CI errors into global structural mistakes. As sample size increases and CI decisions become more stable, the benefit of dialectical conflict resolution correspondingly diminishes.




Figure 3: **Sample Size Efficiency on INSURANCE.** QACD (orange) outperforms PC (blue) in finite-sample regimes and maintains competitive asymptotic performance as N increases.

Across datasets such as ASIA and CHILD, QACD exhibits a recurring trade-off between skeleton recall and causal validity. While aggressive baselines such as MPC or MMHC sometimes achieve slightly higher F_1 , QACD often yields lower NSHD and tighter NSID bounds, indicating improved structural coherence and interventional reliability. This reflects its conservative revision strategy: rather than making irreversible early commitments, QACD iteratively suppresses weakly supported adjacencies.

On datasets with deterministic relations, such as WATER, QACD remains stable and improves some intervention-oriented metrics even when greedy or likelihood-based methods become sensitive to degeneracy; conversely, when CI evidence is already highly consistent or informative, as in EARTHQUAKE or WIN95PTS, dialectical aggregation provides more limited additional benefit.

Overall, the experiments support the usefulness of viewing

Dataset	Method	$F_1 \uparrow$	NSHD $\downarrow$	NSID _{low} $\downarrow$	NSID _{high} $\downarrow$
INSURANCE	PC	0.6664 $\pm$ 0.0300	0.6846 $\pm$ 0.0617	10.8000 $\pm$ 0.7295	10.8154 $\pm$ 0.7054
	MPC	0.6748 $\pm$ 0.0208	0.6500 $\pm$ 0.0211	10.3769 $\pm$ 0.4361	10.3885 $\pm$ 0.4496
	MMHC	0.6335 $\pm$ 0.0265	0.8000 $\pm$ 0.0258	9.5923 $\pm$ 0.2073	10.6923 $\pm$ 0.2159
	QACD(ours)	0.7361 $\pm$ 0.0301	0.5692 $\pm$ 0.0554	8.7462 $\pm$ 0.4713	8.8577 $\pm$ 0.5437
CHILD	PC	0.8445 $\pm$ 0.0449	0.6400 $\pm$ 0.1020	8.3040 $\pm$ 1.0196	10.5440 $\pm$ 0.9926
	MPC	0.9008 $\pm$ 0.0263	0.5440 $\pm$ 0.0607	7.2080 $\pm$ 0.2252	9.1520 $\pm$ 0.7124
	MMHC	0.8425 $\pm$ 0.0118	0.8400 $\pm$ 0.0000	4.2160 $\pm$ 0.3220	12.4560 $\pm$ 0.0358
	QACD(ours)	0.8861 $\pm$ 0.0232	0.4400 $\pm$ 0.1020	4.2400 $\pm$ 0.9104	5.4320 $\pm$ 1.8156
WATER	PC	0.4550 $\pm$ 0.0244	0.9000 $\pm$ 0.0275	7.1303 $\pm$ 0.5925	8.1545 $\pm$ 0.2966
	MPC	0.4898 $\pm$ 0.0201	0.8758 $\pm$ 0.0127	6.7091 $\pm$ 0.5812	8.2030 $\pm$ 0.3152
	QACD(ours)	0.4931 $\pm$ 0.0094	0.9152 $\pm$ 0.0380	6.1242 $\pm$ 0.2905	8.2818 $\pm$ 0.5038
HAILFINDER	PC	0.5489 $\pm$ 0.0131	1.3788 $\pm$ 0.0283	9.1455 $\pm$ 0.6605	10.1000 $\pm$ 0.4164
	MPC	0.5536 $\pm$ 0.0107	1.4061 $\pm$ 0.0392	9.8788 $\pm$ 0.6359	10.8970 $\pm$ 0.6722
	MMHC	0.6490 $\pm$ 0.0091	0.6576 $\pm$ 0.0136	9.8455 $\pm$ 0.0166	13.4061 $\pm$ 0.0166
	QACD(ours)	0.7193 $\pm$ 0.0293	0.6788 $\pm$ 0.0984	7.0545 $\pm$ 1.3054	9.6333 $\pm$ 2.4174
WIN95PTS	PC	0.6668 $\pm$ 0.0367	0.6589 $\pm$ 0.0487	4.8964 $\pm$ 0.6593	5.7964 $\pm$ 0.5919
	MPC	0.6353 $\pm$ 0.0311	0.7000 $\pm$ 0.0436	5.4089 $\pm$ 0.4750	6.3339 $\pm$ 0.3662
	MMHC	0.7147 $\pm$ 0.0136	0.6339 $\pm$ 0.0089	6.0446 $\pm$ 0.4247	7.8006 $\pm$ 0.2481
	QACD(ours)	0.6870 $\pm$ 0.0321	0.6429 $\pm$ 0.0296	4.7464 $\pm$ 0.5152	6.4054 $\pm$ 0.3719

Table 3: Medium and Large Networks Results.

Dataset	QACD ($T_{\max} = 20$)		Ablation ($T_{\max} = 0$)	
	F_1	NSID _{low}	F_1	NSID _{low}
ASIA	0.9181	0.2250	0.7733	2.2500
EARTHQUAKE	1.0000	0.0000	0.8756	0.0000
SURVEY	0.7588	1.9333	0.5127	2.4667

Table 4: **Ablation Study.** Comparison of the full dialectical aggregation model ($T_{\max} = 20$) against the initial candidate skeleton before aggregation ($T_{\max} = 0$).

finite-sample causal discovery as reasoning over conflicting statistical arguments. In the evaluated noisy or inconsistent CI regimes, dialectical aggregation often improves global coherence and reduces interventional error while remaining competitive in skeleton recovery.

5 Related Work

This work relates to both causal discovery and computational argumentation, especially methods for handling uncertainty and conflict in structure learning.

5.1 Causal Discovery from Observational Data

Constraint-based approaches, exemplified by the PC algorithm (Spirtes, Glymour, and Scheines 2000) and Fast Causal Inference (Spirtes, Meek, and Richardson 1995), infer causal structure by testing CI relations and compiling accepted independences into graph edits. Under the Markov and Faithfulness assumptions, these methods are asymptotically consistent, but in finite-sample regimes they are sensitive to CI errors: a single false independence or dependence may trigger incorrect edge deletions and propagate to later orientation steps (Uhler et al. 2013). Variants such as PC-

stable reduce order dependence (Colombo, Maathuis, and others 2014), and multiple-testing corrections can improve statistical robustness, but most constraint-based pipelines still rely on hard CI decisions that become irreversible once incorporated into the graph. QACD instead delays irreversible commitments by aggregating CI outcomes as graded defeasible evidence through an iterative acceptability semantics.

Score-based methods, such as Greedy Equivalence Search and its extensions (Chickering 2002; Ramsey et al. 2017), formulate causal discovery as search over graph structures using penalized likelihood criteria. Recent continuous optimization approaches, including NOTEARS and subsequent variants (Zheng et al. 2018; Ng, Huang, and Zhang 2024), further cast structure learning as differentiable optimization with acyclicity constraints. These methods provide important alternatives to purely constraint-based discovery, but they do not explicitly model conflicts among CI outcomes as argumentative interactions. QACD occupies a complementary position: it preserves the local efficiency of CI testing while introducing a quantitative argumentation layer for resolving evidential conflicts during skeleton discovery.

5.2 Argumentation in Artificial Intelligence

Abstract and quantitative argumentation. Since Dung’s seminal formulation of AFs (Dung 1995), argumentation has become a foundational paradigm for non-monotonic reasoning under inconsistency (Li, Liao, and Chen 2025; Liao and van der Torre 2024). QAFs extend this setting by assigning intrinsic strengths to arguments and defining semantics based on graded acceptability (Chi and Liao 2022; Rago et al. 2016). Unlike (Amgoud and Ben-Naim 2018), which models arguments in a bipolar framework, QACD treats conditional independence tests as defeasible evidence. This distinction allows QACD to better handle noisy or un-

certain data, providing a more robust framework for causal discovery.

Argumentation for causal inquiry. The use of argumentation in causal reasoning is a growing but still relatively specialized area. Early work primarily focused on explanation, employing argumentative structures to justify or communicate causal claims to users (Cyras et al. 2021). More recent studies have explored argumentation-based mechanisms for causal discovery and conflict handling (Russo, Rapberger, and Toni 2024; Li et al. 2026). Related probabilistic approaches connect uncertainty over arguments with graded acceptability (Hunter and Thimm 2017).

The work most closely related to ours is ABAPC (Russo, Rapberger, and Toni 2024), which augments PC with argumentation mechanisms mainly for resolving orientation-level conflicts such as competing v-structures. In contrast, QACD integrates quantitative argumentation directly into skeleton discovery. By treating CI statements as graded attackers of adjacency hypotheses and propagating their influence through witness-based structural context, QACD enables non-monotonic revision of edge hypotheses before the orientation phase.

6 Conclusion

We presented QACD, a QAF for causal discovery that treats CI outcomes as defeasible evidence and aggregates them under structural context. Rather than compiling finite-sample CI decisions into irreversible constraints, QACD revises edge acceptability through direct CI attacks and witness-based propagated attenuation. This provides an iterative, semantics-driven mechanism for resolving conflicting statistical evidence during skeleton discovery.

Empirically, QACD yields more coherent structures and improved interventional validity in several regimes where CI evidence is noisy or inconsistent, while remaining competitive in skeleton recovery. These results support the view that finite-sample causal discovery can be fruitfully formulated as reasoning over graded and potentially conflicting statistical arguments.

Several directions remain open. First, while this paper focuses on the causally sufficient setting, the framework could be extended to latent confounding by introducing arguments for bi-directed edges in ancestral graphs. Second, the iterative update could be accelerated for very large graphs using approximate message-passing or sparsity-aware implementations. Third, although we expect finite-sample conflicts to diminish under standard Markov and faithfulness assumptions with consistent CI testing, a complete oracle-level consistency analysis for QACD remains future work. Finally, interventional data could be incorporated as a new class of intervention arguments, enabling argumentation-based discovery from mixed observational and experimental evidence.

Acknowledgments

This research was partially supported by the National Natural Science Foundation of China (62576309).

AI Declaration

During the preparation of this work, the authors used ChatGPT for language polishing, grammar correction, and editorial refinement. After using this tool, the authors thoroughly reviewed and edited all generated content and take full responsibility for the content of this publication.

References

- Amgoud, L., and Ben-Naim, J. 2018. Weighted bipolar argumentation graphs: Axioms and semantics. In *Twenty-Seventh International Joint Conference on Artificial Intelligence-IJCAI 2018*, 5194–5198.
- Amgoud, L.; Maudet, N.; and Parsons, S. 2000. Modelling dialogues using argumentation. In *Proceedings Fourth International Conference on MultiAgent Systems*, 31–38. IEEE.
- Anand, T. V.; Ribeiro, A. H.; Tian, J.; Hripcsak, G.; and Bareinboim, E. 2025. Causal discovery over clusters of variables in markovian systems. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Athey, S., and Imbens, G. W. 2017. The state of applied econometrics: Causality and policy evaluation. *Journal of Economic perspectives* 31(2):3–32.
- Chi, H., and Liao, B. 2022. A quantitative argumentation-based automated explainable decision system for fake news detection on social media. *Knowledge-Based Systems* 242:108378.
- Chickering, D. M. 2002. Optimal structure identification with greedy search. *Journal of machine learning research* 3(Nov):507–554.
- Colombo, D.; Maathuis, M. H.; et al. 2014. Order-independent constraint-based causal structure learning. *J. Mach. Learn. Res.* 15(1):3741–3782.
- Cyras, K.; Rago, A.; Albini, E.; Baroni, P.; and Toni, F. 2021. Argumentative XAI: A survey. In Zhou, Z., ed., *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, 4392–4399. ijcai.org.
- Ding, X.; Wu, H.; Wang, J.; Xu, J.; and Xin, M. 2024. Causal discovery for rolling bearing fault under missing data: From the perspective of causal effect and information flow. *IEEE Transactions on Instrumentation and Measurement*.
- Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence* 77(2):321–357.
- Hunter, A., and Thimm, M. 2017. Probabilistic reasoning with abstract argumentation frameworks. *Journal of Artificial Intelligence Research* 59:565–611.
- Jabbari, F.; Ramsey, J.; Spirtes, P.; and Cooper, G. 2017. Discovery of causal models that contain latent variables through bayesian scoring of independence constraints. In *Joint European conference on machine learning and knowledge discovery in databases*, 142–157. Springer.
- Li, Z.; Fang, X.; Chen, C.; Li, M.; and Liao, B. 2026. Enhancing conflict resolution in language models via abstract argumentation. *Neurocomputing* 664:132093.

- Li, Z.; Liao, B.; and Chen, C. 2025. Investigation of semantic behavior in probabilistic argumentation. *International Journal of Approximate Reasoning* 187:109551.
- Liao, B., and van der Torre, L. 2024. Attack-defense semantics of argumentation. In *Computational Models of Argument*. IOS Press. 133–144.
- Ling, Z.; Li, B.; Zhang, Y.; Zhou, P.; Wu, X.; Huang, Y.; Yu, K.; and Wu, X. 2024. Causal discovery using weight-based conditional independence test. *ACM Transactions on Knowledge Discovery from Data* 19(1):1–24.
- Meek, C. 1995. Causal inference and causal explanation with background knowledge. In Besnard, P., and Hanks, S., eds., *UAI '95: Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence, Montreal, Quebec, Canada, August 18-20, 1995*, 403–410. Morgan Kaufmann.
- Neuberg, L. G. 2003. Causality: models, reasoning, and inference, by judea pearl, cambridge university press, 2000. *Econometric Theory* 19(4):675–685.
- Ng, I.; Huang, B.; and Zhang, K. 2024. Structure learning with continuous optimization: A sober look and beyond. In *Causal Learning and Reasoning*, 71–105. PMLR.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann.
- Peters, J., and Bühlmann, P. 2015. Structural intervention distance for evaluating causal graphs. *Neural Comput.* 27(3):771–799.
- Pollock, J. L. 1987. Defeasible reasoning. *Cognitive science* 11(4):481–518.
- Rago, A.; Toni, F.; Aurisicchio, M.; Baroni, P.; et al. 2016. Discontinuity-free decision support with quantitative argumentation debates. *KR* 16:63–73.
- Ramsey, J.; Glymour, M.; Sanchez-Romero, R.; and Glymour, C. 2017. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International journal of data science and analytics* 3(2):121–129.
- Russo, F.; Rapberger, A.; and Toni, F. 2024. Argumentative Causal Discovery. In *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning*, 938–949.
- Sachs, K.; Perez, O.; Pe’er, D.; Lauffenburger, D. A.; and Nolan, G. P. 2005. Causal protein-signaling networks derived from multiparameter single-cell data. *Science* 308(5721):523–529.
- Spirtes, P., and Glymour, C. 1991. An algorithm for fast recovery of sparse causal graphs. *Social science computer review* 9(1):62–72.
- Spirtes, P.; Glymour, C. N.; and Scheines, R. 2000. *Causation, prediction, and search*. MIT press.
- Spirtes, P.; Meek, C.; and Richardson, T. 1995. Causal inference in the presence of latent variables and selection bias. In *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence*, 499–506.
- Tsamardinos, I.; Brown, L. E.; and Aliferis, C. F. 2006. The max-min hill-climbing bayesian network structure learning algorithm. *Machine learning* 65(1):31–78.
- Uhler, C.; Raskutti, G.; Bühlmann, P.; and Yu, B. 2013. Geometry of the faithfulness assumption in causal inference. *The Annals of Statistics* 436–463.
- Wang, Z., and Shen, Y. 2025. Fuzzy labelling semantics for quantitative argumentation. *Journal of Logic and Computation* 35(4):exaf009.
- Wei, S.; Chen, Y.; and Liao, B. 2026. Causal discovery as dialectical aggregation: A quantitative argumentation framework.
- Yang, X.; Lan, T.; Qiu, H.; and Zhang, C. 2025. Non-linear causal discovery via dynamic latent variables. *IEEE Transactions on Automation Science and Engineering*.
- Zheng, X.; Aragam, B.; Ravikumar, P. K.; and Xing, E. P. 2018. Dags with no tears: Continuous optimization for structure learning. *Advances in neural information processing systems* 31.