Bayesian Network Structure Learning through Large Language Models

Extended Abstract

Jiebin Cai
ShenZhen University
ShenZhen, China
caijiebin2023@email.szu.edu.cn

Yinghui Pan
ShenZhen University
ShenZhen, China
panyinghui@szu.edu.cn

Yifeng Zeng
Northumbria University
Newcastle, UK
yifeng.zeng@northumbria.ac.uk

Han Liu
ShenZhen University
ShenZhen, China
han.liu@szu.edu.cn

ABSTRACT

Constructing Bayesian Networks in data-scarce scenarios typically relies on costly expert knowledge. While Large Language Models (LLMs) offer a promising data-free alternative, they often suffer from hallucinations and generate structurally invalid networks containing cycles. To address these challenges, we propose a novel multi-agent framework comprising a Decider, Critic, and Arbiter (DCA) for automated BN structure learning. By integrating triplet-based causal reasoning with a confidence-driven network refinement strategy, our approach effectively eliminates redundant edges and ensures the generation of valid Directed Acyclic Graphs (DAGs). Experimental results on standard benchmarks demonstrate that our method significantly outperforms existing LLM-based baselines in terms of F1-score and Structural Hamming Distance (SHD), while successfully avoiding cycles and isolated nodes.

KEYWORDS

Large Language Models; Bayesian Networks; Causal Discovery; Multi-agent Systems

ACM Reference Format:
Jiebin Cai, Yinghui Pan, Yifeng Zeng, and Han Liu. 2026. Bayesian Network Structure Learning through Large Language Models: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/SGHT1267

1 INTRODUCTION

Bayesian Networks represent a robust probabilistic graphical model that integrates probability theory with graph theory, intuitively representing variables and their conditional dependencies through a directed acyclic graph [2]. These models have demonstrated powerful capabilities in uncertainty reasoning and decision support across numerous fields, such as medical diagnosis and risk assessment. However, the construction of traditional Bayesian Networks has heavily relied on large-scale labeled data. Data-driven structure learning methods, including constraint-based algorithms like the PC algorithm [3] and score-based searches, often see their performance deteriorate sharply in real-world scenarios characterized by data scarcity, high-dimensional sparsity, or selection bias. Consequently, exploring methods for constructing Bayesian Networks that rely solely on prior knowledge or text descriptions rather than data has become a crucial research direction.

In recent years, Large Language Models have achieved breakthrough progress in causal inference and text understanding tasks. These models acquire rich world knowledge and causal reasoning abilities through extensive pre-training, enabling them to directly infer associations between variables based on natural language descriptions. Despite this promise, existing methods based on Large Language Models still face significant challenges. The reasoning process of these models often exhibits randomness and semantic inconsistency, which results in incorrect or redundant edge directions. More critically, the generated networks frequently contain cycles or unreasonable edges that violate the strict structural constraints of Bayesian Networks, specifically the requirement to be a directed acyclic graph.

To address these issues, this paper proposes a multi-agent adjudication framework for the construction of Bayesian Networks using Large Language Models. Our approach integrates triplet reasoning and network refinement to achieve the automatic generation of highly consistent and acyclic networks. We introduce a collaborative mechanism involving three distinct agent roles: the Decider, the Critic, and the Arbiter. This framework allows for a joint reasoning process over node triplets to enhance the accuracy of causal judgments. Following the initial inference, a network refinement stage is employed to identify and remove redundant edges. Furthermore, based on confidence scores derived from edge occurrence ratios, a cycle removal procedure is applied to ensure the resulting structure is topologically valid. Experimental results demonstrate that the proposed method outperforms existing baseline approaches on multiple benchmark datasets, effectively reducing redundant edges and cyclic structures while maintaining high semantic consistency.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/SGHT1267

2 METHODOLOGY

This paper proposes a large language model-based method for constructing Bayesian Networks that integrates multi-agent collaboration with triplet reasoning. The overall process consists of two main stages: the causal relationship inference stage and the network refinement stage.

In the causal relationship inference stage, we employ a multi-agent adjudication framework comprising three distinct roles: the Decider, the Critic, and the Arbiter. Instead of querying simple node pairs, the model processes variables in triplets to capture more contextual dependency information. The Decider agent provides an initial causal hypothesis based on the semantic descriptions of the variables. The Critic agent then evaluates this hypothesis, checking for logical consistency and potentially offering a counter-argument. If a disagreement arises, the Arbiter agent intervenes to make the final judgment by weighing the reasoning from both sides. This adversarial interaction significantly reduces the hallucination often seen in single-agent models.

Following the generation of initial causal edges, we calculate a confidence score for each edge based on its statistical frequency of occurrence across different triplets. This statistical foundation is crucial for the subsequent network refinement stage.

The second stage focuses on topological optimization to ensure the final structure is a valid directed acyclic graph. First, the large language model is prompted to identify and remove redundant edges where an indirect causal path already explains the relationship. Second, to eliminate cyclic structures, we apply a confidence-based cycle removal strategy. When a cycle is detected, the algorithm identifies the edge with the lowest statistical confidence score derived from the inference stage and removes it. This process preserves the most distinct causal pathways while enforcing the necessary structural constraints of Bayesian Networks.

3 EXPERIMENTAL RESULTS

We evaluated our DCA framework on a comprehensive set of benchmarks, including standard networks from bnlearn and custom domain-specific datasets. Due to space constraints, we present results on three representative datasets against four baselines: zero-shot, Chain-of-Thought (CoT) [5], ROT [4], and Multi-expert [1].

Table 1 presents the SHD and F1-score. Our method achieves the highest F1-score across all datasets. On Cancer and Alarm, we significantly outperform all baselines in SHD.

For the large-scale Hailfinder network, while simple baselines (Zero-shot/CoT) achieve lower SHD values, this metric is misleading without considering structural validity. As shown in Table 2, these methods generate a large number of isolated nodes (e.g., 21.8 for Zero-shot), failing to construct a connected graph.

In contrast, our method ensures topological correctness. As detailed in Table 2, advanced baselines like ROT and Multi-expert suffer from generating invalid cyclic structures (e.g., 252 cycles for ROT on Hailfinder). Our framework, leveraging the DCA mechanism and confidence-based refinement, achieves 0 cycles and 0 isolated nodes (on most datasets), demonstrating a superior balance between causal accuracy and structural legality.

Table 1: Performance Comparison (SHD & F1).

Dataset	Method	SHD ↓	F1 ↑
Cancer	Zero-shot	2.6 ± 0.55	0.72
	CoT	2.8 ± 0.45	0.70
	ROT	5.0 ± 0.0	0.62
	Multi-expert	4.0 ± 0.0	0.67
	Ours	**0.0 ± 0.0**	**1.00**
Alarm	Zero-shot	133.2 ± 3.11	0.24
	CoT	136.4 ± 4.72	0.24
	ROT	411.67 ± 2.31	0.13
	Multi-expert	354.0 ± 4.9	0.14
	Ours	**92.8 ± 6.38**	**0.25**
Hailfinder	Zero-shot	**103.2 ± 3.83**	0.11
	CoT	104.0 ± 2.92	0.13
	ROT	1027.33 ± 10.2	0.10
	Multi-expert	753.7 ± 19.6	0.12
	Ours	245.2 ± 11.44	**0.14**

Table 2: Structural Validity (Cycles & Isolated Nodes).

Dataset	Method	Cycles ↓	Iso Nodes ↓
Cancer	Zero-shot	0.0	0.6
	CoT	0.0	0.8
	ROT	0.0	0.0
	Multi-expert	0.0	0.0
	Ours	**0.0**	**0.0**
Alarm	Zero-shot	0.0	1.0
	CoT	0.0	1.2
	ROT	114.3	**0.0**
	Multi-expert	43.8	0.2
	Ours	**0.0**	1.0
Hailfinder	Zero-shot	0.0	21.8
	CoT	0.0	21.2
	ROT	252.67	**0.0**
	Multi-expert	33.67	**0.0**
	Ours	**0.0**	**0.0**

4 CONCLUSION

We propose a multi-agent framework (DCA) for data-free Bayesian network learning. By integrating Decider, Critic, and Arbiter agents with triplet reasoning and confidence-based refinement, our method mitigates hallucinations and ensures valid DAGs. Experiments show that our approach outperforms SOTA baselines in F1-score and structural correctness, achieving zero cycles and minimal isolated nodes.

REFERENCES

[1] Do Xuan Long, Duong Ngoc Yen, Anh Tuan Luu, Kenji Kawaguchi, Min-Yen Kan, and Nancy F Chen. 2024. Multi-expert prompting improves reliability, safety, and usefulness of large language models. *arXiv preprint arXiv:2411.00492* (2024).

[2] Judea Pearl. 2014. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Elsevier.

[3] Peter Spirtes, Clark Glymour, and Richard Scheines. 1991. From probability to causality. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 64, 1 (1991), 1–36.

[4] Li Wang, Xi Chen, XiangWen Deng, Hao Wen, MingKe You, WeiZhi Liu, Qi Li, and Jian Li. 2024. Prompt engineering in consistency and reliability with the evidence-based guideline for LLMs. *npj Digital Medicine* 7 (2024). https://doi.org/10.1038/s41591-024-00717-4

[5] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.