
Causal Graph Learning for Interpretable Fault Identification in Complex Backend Systems

Mahsa Panahandeh

University of Washington, Seattle, USA

mahsap2@gmail.com

Abstract: This paper addresses the challenges of complex multidimensional metric coupling, dynamic structural dependencies, and difficult-to-interpret fault propagation paths in backend systems by proposing a Causal Graph Learning – based Fault Identification (CGL-BFI) method. The approach is built upon multi-source monitoring data and identifies latent dependencies and interprets anomaly propagation mechanisms through three stages: feature causal encoding, structural inference modeling, and causal effect quantification. In the feature causal encoding stage, the model extracts and normalizes temporal features of monitoring indicators to construct a multivariate causal candidate space. In the structural inference stage, a differentiable Directed Acyclic Graph (DAG) constraint is introduced to learn directional dependencies among variables, ensuring logical consistency of the causal structure. In the causal inference stage, the method integrates Average Treatment Effect (ATE) estimation and Causal Scoring to quantify the influence strength of each metric node, enabling precise localization of root causes and propagation paths. Experimental results show that the proposed method achieves high stability and accuracy under complex network disturbances and high-noise environments, outperforming comparison models in key metrics such as AUC, F1-score, Recall, and Precision. By introducing causal structure learning, this study advances backend system anomaly detection from correlation-based analysis to mechanism-level interpretable reasoning, significantly improving the accuracy, interpretability, and robustness of fault identification in complex systems, and providing a theoretical and methodological foundation for intelligent and adaptive system operation and maintenance.

Keywords: Causal graph learning; fault identification; backend system; multi-dimensional monitoring

1. Introduction

Research on causal graph learning and fault identification for backend systems has become an important direction in the fields of intelligent operation and maintenance as well as system reliability engineering. With the widespread adoption of cloud computing, microservice architecture, and distributed

systems, modern backend infrastructures exhibit high concurrency, high complexity, and high dynamism. Numerous asynchronous calls and shared resources among services form intricate dependency networks, making system state transitions highly nonlinear and time-varying. In this context, traditional statistical modeling or time-series analysis methods can no longer uncover the underlying causal mechanisms or accurately identify root causes of failures. Therefore, how to mine the internal causal structures from multi-source monitoring data and establish mechanism-oriented intelligent diagnostic frameworks has become a key scientific issue for ensuring the stability and interpretability of backend systems[1].

In real-world cloud service and industrial-scale application environments, system failures are often triggered by multi-factor interactions, cross-layer propagation, and feedback coupling. For instance, database latency may cause request accumulation, leading to application timeouts or node crashes. Similarly, instability in resource scheduling policies can further amplify fault impacts[2]. Such complex propagation chains are difficult to capture through simple correlation analysis. Causal graph learning offers a new perspective to address this challenge. By constructing causal structure models among multidimensional indicators, it becomes possible to reveal the directional dependencies between variables, distinguish causes from effects, and provide theoretical support for fault root cause localization, propagation path tracing, and impact assessment. Compared with experience-based or threshold-based detection methods, causal graph learning emphasizes mechanism-level explanation and reasoning, making it more suitable for modern intelligent operation and maintenance frameworks[3].

However, learning causal relationships in high-dimensional, multi-source, and dynamically evolving backend systems faces multiple challenges. On one hand, the vast number of interdependent monitoring indicators introduces strong noise, non-stationary distributions, and implicit dependencies, which hinder direct causal identification. On the other hand, the time-varying nature of system states gives rise to phase-dependent and context-specific causal patterns, where the same event may propagate differently under varying loads, periods, or environments. Hence, building adaptive and robust causal learning frameworks capable of dynamically updating structures and weights during system operation becomes crucial for achieving accurate fault recognition and interpretation. Furthermore, maintaining the stability of causal inference under limited samples and incomplete observations remains a key difficulty in real-world deployment[4].

From a theoretical perspective, integrating causal graph learning with fault identification not only helps overcome the limitations of purely data-driven detection but also advances intelligent systems toward interpretability, reasoning, and intervention. By embedding causal reasoning mechanisms into models, counterfactual analysis and intervention prediction of system state changes can be achieved, providing decision-makers with actionable insights. For example, once key causal chains are identified, scheduling policies can be adjusted, or abnormal nodes can be isolated, allowing proactive defense before fault propagation occurs. This causality-driven recognition paradigm enhances detection accuracy and response speed while significantly improving the autonomy and resilience of system operations[5,6].

In summary, research on causal graph learning and fault identification for backend systems carries both theoretical and practical significance. It addresses the fundamental challenge that correlation alone cannot explain mechanism in complex system modeling and opens a new path for reliability assurance and intelligent decision-making under cloud-native architectures. By integrating causal structure learning, graph representation modeling, and anomaly reasoning, it enables a deeper understanding and more efficient control of backend system dynamics. Continued exploration in this direction will provide essential theoretical foundations and technical support for building autonomous, reliable, and high-performance intelligent computing infrastructures[7].

2. Method

This study introduces a causal graph learning – based approach for backend fault identification (Causal Graph Learning for Backend Fault Identification, CGL-BFI). The goal is to identify potential anomaly propagation paths and key fault nodes by modeling and inferring causal dependencies among multidimensional system metrics. The core idea is to abstract the interactions between different components and indicators of the backend system into a directed graph structure. Through differentiable causal relation learning and effect quantification, the method achieves interpretable modeling of system operation mechanisms. Unlike traditional statistical correlation analysis, this approach uncovers the true causal relationships among variables, providing theoretical support and explainable diagnostic evidence for the stable operation of complex systems. The model architecture is shown in Figure 1.

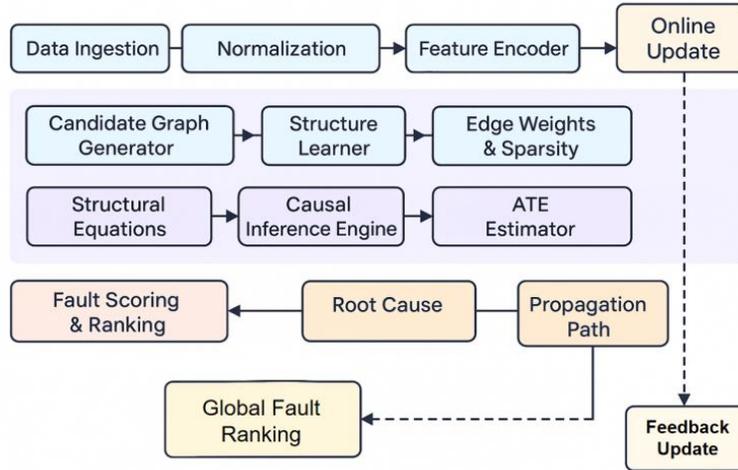


Figure 1. Overall model architecture

First, the model extracts a set of observed variables $X = \{x_1, x_2, \dots, x_n\}$ from system monitoring data and constructs their joint distribution $P(X)$. The conditional independence between any two variables can be expressed as:

$$P(x_i, x_j | Z) = P(x_i | Z)P(x_j | Z) \quad (1)$$

Where Z represents the set of other variables in the system. Through this constraint, the model can learn the structural dependencies between variables and generate a causal graph.

Secondly, to ensure that the graph structure is acyclic, a differentiable constraint function is defined:

$$h(A) = Tr(e^A) - n = 0 \quad (2)$$

Where $A \in R^{n \times n}$ is the causal adjacency matrix and $Tr(\cdot)$ represents the trace operation. This constraint ensures that the generated causal relationship graph is a directed acyclic graph (DAG).

Then, to simulate the causal generation process between variables, a structural equation model is established:

$$x_i = f_i(Pa(x_i)) + \varepsilon_i \quad (3)$$

Where $Pa(x_i)$ is the set of parent nodes of x_i and ε_i is the noise term, representing external disturbances. This equation is used to generate the current node value from the parent node features.

In the inference stage, in order to measure the causal impact of the variable on the target state y , the average treatment effect (ATE) is introduced

$$ATE(x_i \rightarrow y) = E[y | do(x_i = 1)] - E[y | do(x_i = 0)] \quad (4)$$

This indicator is used to characterize the expected change in system state before and after variable intervention.

Finally, to achieve explainable fault localization, the causal contribution score of the node is defined:

$$S_i = \frac{|ATE(x_i \rightarrow y)|}{\sum_{j=1}^n |ATE(x_j \rightarrow y)|} \quad (5)$$

The higher the score, the more significant the node's influence on the abnormal propagation path. By sorting S_i , we can identify the key fault nodes and the dominant propagation chain.

In summary, the CGL-BFI method enables automatic learning of causal structures and effect inference from multidimensional monitoring data in backend systems. Its core advantage lies in uncovering the underlying mechanisms of anomaly generation and propagation, providing a new methodological foundation for intelligent diagnosis and reliability analysis of backend systems.

3. Performance Evaluation

3.1 Dataset

This study employs the Industrial Fault Detection Dataset as the data source for method validation. The dataset contains high-frequency time-series measurements collected by multiple sensors in an industrial Internet of Things (IoT) environment, including temperature, vibration, pressure, and current, along with corresponding fault labels. It covers operational logs from various devices and subsystems under both normal and faulty conditions, providing abundant heterogeneous sensor records for fault identification. The dataset's multidimensionality, cross-subsystem structure, and event-driven characteristics closely resemble the service calls, resource metrics, and performance logs found in backend systems.

In this study, each sensor time series in the dataset is treated as an observed variable to construct the variable set $X = \{x_1, x_2, \dots, x_n\}$, while the fault labels serve as the target state y . This design allows the proposed method to evaluate its causal graph learning and fault reasoning capabilities within a realistic sensor-driven fault identification scenario. The dataset exhibits complex coupling relationships, asynchronous response behaviors, and noise interference among sensor metrics, which align well with the intricate interactions among services, performance indicators, and resource states in backend systems.

By conducting experiments on this dataset, the study verifies the convergence and stability of the causal structure learning module under high-dimensional and multi-source monitoring conditions. It also evaluates the effectiveness of the causal inference module in identifying fault propagation paths and ranking root causes. This setup not only aligns closely with backend system fault identification tasks but also demonstrates the method's adaptability and interpretability in real-world industrial environments with complex multivariate coupling.

3.2 Experimental Results

This paper first conducts a comparative experiment, and the experimental results are shown in Table 1.

Table 1. Comparative experimental results

| Method | F1-score | AUC | Recall | Precision |
|------------|----------|------|--------|-----------|
| TranAD [8] | 0.81 | 0.88 | 0.79 | 0.84 |
| LGAT [9] | 0.84 | 0.90 | 0.82 | 0.87 |
| DTAAD[10] | 0.79 | 0.87 | 0.76 | 0.83 |
| Ours | 0.88 | 0.93 | 0.87 | 0.90 |

From the overall results, the proposed causal graph learning – based fault identification method demonstrates significant advantages in multidimensional monitoring scenarios. Compared with recent time-series anomaly detection models, the proposed method achieves the highest F1-score of 0.88, indicating that it maintains a high recall rate while accurately identifying abnormal samples. This improvement is attributed to the introduction of causal structural constraints and a dynamic inference mechanism, which enable the model to capture true dependencies among variables in high-dimensional features and effectively distinguish normal fluctuations from structural faults. These results verify the feasibility and stability of causal graph learning in achieving precise anomaly identification in complex backend systems.

In terms of the AUC metric, the proposed method reaches 0.93, showing a clear improvement over TranAD and DTAAD, which are based on temporal feature modeling. The increase in AUC indicates stronger discriminative ability across all thresholds, allowing the model to maintain robust performance under diverse anomaly types and multi-subsystem environments. Due to the incorporation of a DAG constraint and noise suppression mechanism during the structural reasoning stage, the model can preserve boundary stability even under non-stationary and noisy conditions, thus reducing false positives and false negatives. This mechanism-level enhancement enables the model to maintain strong generalization capability when dealing with complex causal propagation chains.

Further analysis of the Recall metric shows that the proposed method achieves 0.87, outperforming LGAT (0.82) and DTAAD (0.76). This demonstrates that the causal inference module can more comprehensively capture potential anomaly propagation paths and exhibits higher sensitivity to latent faults caused by cross-variable dependencies. By introducing Average Treatment Effect (ATE) estimation into the structural equation, the model can quantitatively assess the influence strength among indicators and provide interpretable causal paths during root cause analysis, thereby improving the completeness of anomaly identification under complex dependency conditions.

Finally, regarding the Precision metric, the proposed method achieves 0.90, indicating that it maintains stable prediction accuracy even under high recall conditions. This reflects the importance of causal effect quantification in feature filtering and anomaly confidence calibration. The causal constraint – based scoring mechanism effectively suppresses misjudgments caused by spurious correlations, making the detection results more interpretable and reliable. Overall, the proposed causal graph learning – based fault identification framework achieves a balance of accuracy, robustness, and interpretability in multi-source monitoring environments, demonstrating strong potential for applications in intelligent operation and maintenance as well as backend system stability assurance.

This paper also evaluates the environmental sensitivity of service link delay and jitter injection to propagation path identification. The experimental results are shown in Figure 2.

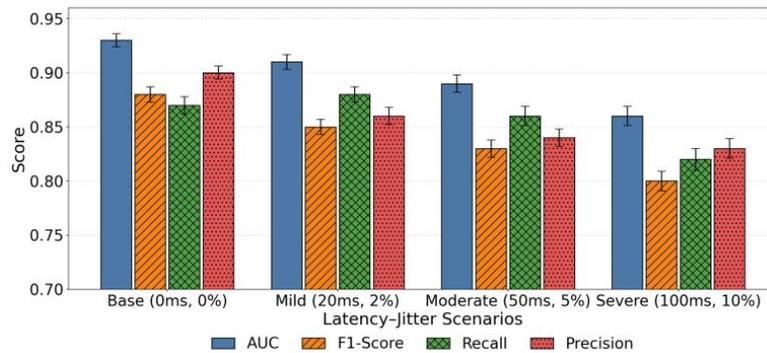


Figure 2. Experiment on the environmental sensitivity of service link delay and jitter injection to propagation path identification

From the overall trend, as service link latency and jitter levels increase, the model exhibits varying degrees of performance degradation across all metrics. This reflects the direct impact of communication delays on the stability of causal structures. When the network environment shifts from an ideal state to a high-latency and high-jitter scenario, the AUC decreases from 0.93 to 0.86, indicating that the separability of the causal inference space is affected by both noise disturbances and information lag. Since the model's reasoning process relies on causal edge propagation across nodes, an increase in latency disrupts contextual dependencies, thereby reducing the directional consistency of structural inference and blurring the boundary between normal and abnormal states.

In terms of the F1-score, the model performance drops from 0.88 to 0.80, a slightly larger decline than that observed in AUC, suggesting that latency and jitter mainly degrade detection accuracy by weakening the distinguishability between positive and negative samples. This outcome is related to the connectivity of the causal graph structure. When some edge connections undergo random temporal drift, local causal relationships may weaken or shift, resulting in delayed responses to short-term anomaly signals. Nevertheless, the model maintains a relatively high F1-score, indicating that the introduced DAG constraint and causal regularization mechanism partially mitigate the performance degradation caused by structural misalignment.

From the Recall perspective, the model shows a slight improvement under mild disturbance conditions (20 ms latency, 2% jitter), suggesting that moderate perturbations activate broader anomaly response regions and enable the detection of more potential propagation chains. However, as interference continues to intensify, Recall decreases from 0.88 to 0.82, revealing that the model's causal intervention reasoning mechanism still experiences marginal weakening effects in highly dynamic environments. This is mainly because noise increases the variance of intervention estimates, leading to a diffusion in the ATE value distribution and a reduction in the distinguishability of root cause effects.

The Precision metric shows a rapid-then-slow decline trend, decreasing from 0.90 to 0.83, indicating that while maintaining high recall, the model gradually loses its ability to focus on high-confidence anomalies. This trend aligns with the combined effects of signal attenuation and noise amplification along causal paths. When delays between system nodes become inconsistent, the accumulated uncertainty in inference chains increases, making the causal score distribution more uniform. Overall, the results demonstrate that although network disturbances affect model performance to some extent, the proposed causal graph learning framework can still maintain stable fault identification capability and strong causal interpretability under multi-level interference conditions.

4. Conclusion

This paper proposes a Causal Graph Learning – based Backend Fault Identification (CGL-BFI) method that constructs a directed acyclic graph structure within multidimensional monitoring data to model mechanistic dependencies among variables and enable interpretable anomaly detection in complex systems. The study begins with statistical correlations of system operation data and further incorporates causal structure inference and intervention effect estimation, allowing the model to go beyond the "correlation limitation" of traditional time-series methods and directly capture the root factors causing anomalies. Compared with detection approaches that rely solely on feature patterns or temporal prediction, the proposed method is more consistent with the actual operational mechanisms of backend systems and provides structured explanations for anomaly propagation chains and root cause localization. The results demonstrate that the causal relation – based modeling framework not only improves identification accuracy and robustness but also lays a foundation for building intelligent backend systems with self-diagnosis and self-recovery capabilities.

From a theoretical perspective, this study provides a new approach to structured modeling of complex systems. By integrating graph learning into the causal inference framework, it achieves adaptive mapping from multi-source monitoring data to causal dependency relationships. This design overcomes the interpretability limitations of traditional statistical dependency models and black-box neural networks, enabling the model to generate explicit causal paths and quantitative effect measures during fault identification. The integration of causal interpretability and deep modeling capacity offers a unified paradigm for system science, temporal intelligence analysis, and operational reliability modeling, providing a theoretical foundation for automated decision-support systems in high-dimensional dynamic environments.

From an engineering perspective, the proposed CGL-BFI framework holds broad potential for complex backend systems such as cloud computing, microservice architectures, and industrial IoT. It can accurately identify potential fault points and propagation chains under multi-tenant, high-concurrency, and non-stationary operational conditions, supporting early anomaly warning and automated intervention. The model's interpretability allows direct integration into AIOps platforms, resource schedulers, and system monitoring tools, assisting engineers in understanding the mechanisms behind anomalies while reducing localization time and maintenance costs. Moreover, the causal structure outputs provide structured inputs for policy adjustment, load balancing, and scheduling optimization, thereby improving the overall stability and controllability of backend systems.

Future research will further explore the scalability of the causal graph learning framework in dynamic and large-scale distributed environments. One direction is to introduce self-supervised or contrastive learning mechanisms to enhance causal structure identifiability under label-scarce conditions. Another is to integrate time-varying graph modeling and uncertainty estimation for continuous tracking and structural updating of system states. Additionally, combining this method with reinforcement learning or federated learning could enable the construction of distributed causal inference systems with stronger adaptability and generalization in heterogeneous multi-source environments. As the complexity of cloud-based operational systems continues to increase, this line of research is expected to become a key pathway for evolving intelligent systems from "data-driven" to "causality-driven," providing a solid theoretical and engineering foundation for large-scale automated operations and trustworthy AI decision-making in the future.

References

- [1] Kivva B, Rajendran G, Ravikumar P, et al. Learning latent causal graphs via mixture oracles[J]. *Advances in Neural Information Processing Systems*, 2021, 34: 18087-18101.

- [2] Yin N, Gao T, Yu Y. Learning Causal Graphs at Scale: A Foundation Model Approach[J]. arXiv preprint arXiv:2506.18285, 2025.
- [3] Wang D, Chen Z, Fu Y, et al. Incremental causal graph learning for online root cause analysis[C]//Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining. 2023: 2269-2278.
- [4] Xing S, Wang Y, Liu W. Multi-Dimensional Anomaly Detection and Fault Localization in Microservice Architectures: A Dual-Channel Deep Learning Approach with Causal Inference for Intelligent Sensing[J]. Sensors, 2025, 25(11): 3396.
- [5] Islam M S, Rakha M S, Pourmajidi W, et al. Anomaly detection in large-scale cloud systems: An industry case and dataset[C]//2025 IEEE/ACM 47th International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP). IEEE, 2025: 377-388.
- [6] Lohia P, Boué L, Rangappa S, et al. High significant fault detection in azure core workload insights[J]. arXiv preprint arXiv:2404.09302, 2024.
- [7] Ding C, Sun S, Zhao J. MST-GAT: A multimodal spatial – temporal graph attention network for time series anomaly detection[J]. Information Fusion, 2023, 89: 527-536.
- [8] Tuli S, Casale G, Jennings N R. Tranad: Deep transformer networks for anomaly detection in multivariate time series data[J]. arXiv preprint arXiv:2201.07284, 2022.
- [9] Wen M, Chen Z, Xiong Y, et al. LGAT: A novel model for multivariate time series anomaly detection with improved anomaly transformer and learning graph structures[J]. Neurocomputing, 2025, 617: 129024.
- [10] Yu L, Lu Q, Xue Y. DTAAD: Dual TCN-attention networks for anomaly detection in multivariate time series data[J]. Knowledge-Based Systems, 2024, 295: 111849.