Explainable Intelligent Audit Risk Assessment with Causal Graph Modeling and Causally Constrained Representation Learning

Jianlin Lai , Chen Chen , Jingjing Li , Qingmiao Gan *

Posted Date: 11 December 2025

doi: 10.20944/preprints202512.1080.v1

Keywords: causal inference; intelligent auditing; risk assessment; explainable models

Article

Explainable Intelligent Audit Risk Assessment with Causal Graph Modeling and Causally Constrained Representation Learning

Jianlin Lai [1], Chen Chen [2], Jingjing Li [3] and Qingmiao Gan [4,*]

[1] Babson College, Wellesley, USA
[2] Vanderbilt University, Nashville, USA
[3] University of Illinois Urbana-Champaign, Champaign, USA
[4] Trine University, Phoenix, USA
* Correspondence: gqmkate@gmail.com

Abstract

This study proposes an intelligent audit risk assessment method that integrates causal structure modeling, causal identifiability reasoning, and interpretable representation learning to address the lack of transparency in risk identification, the presence of confounded variable relationships, and the limitations of correlation-based inference in complex audit scenarios. The method first constructs a structured causal graph of the audit workflow to formalize the triggering relationships and interaction paths among audit features, and then applies structural equations and identifiability analysis to reveal latent causal dependencies. Based on this foundation, the model generates interpretable feature embeddings through causally constrained representation learning, allowing inference results to map back to the business semantic space along causal paths and enabling visual analysis of risk formation. To validate the effectiveness of the approach, this study conducts comparison experiments, ablation experiments, and multidimensional sensitivity analyses on a public audit dataset, and evaluates the method across model accuracy, interpretability, noise robustness, distributional shifts, and hyperparameter variations. The experimental results show that the method achieves significant improvements over existing models in accuracy, precision, recall, and F1-score, while maintaining stable performance under noise interference, class imbalance, learning rate changes, and latent dimension adjustments. The model also produces clear causal chain explanations that help auditors understand risk sources, identify key process components, and trace potential triggering mechanisms through structured reasoning logic. Overall, this study achieves a deep integration of causal inference and intelligent auditing and provides a complete methodological framework and empirical evidence for building transparent, trustworthy, and highly interpretable audit risk assessment systems.

CCS CONCEPTS: Computing methodologies~Machine learning~Machine learning approaches

Keywords: causal inference; intelligent auditing; risk assessment; explainable models

I. Introduction

The rapid growth of the digital economy has led to increasing complexity and scale in enterprise business systems. Audit activities now face unprecedented challenges. Traditional audit methods rely on manual experience and post hoc verification[1]. They cannot handle massive, multimodal, heterogeneous, and fast-generated audit data. They also cannot identify potential risks in high-frequency business flows in a timely manner. With the expansion of automated processes, digital invoice systems, and intelligent financial platforms, the structure of audit objects has become more complex. Risk propagation now shows dynamic evolution, cross-system diffusion, and opacity of

 doi:10.20944/preprints202512.1080.v1

rules. Traditional rule-based monitoring faces clear limitations in adaptability, coverage, and interpretability. In this context, intelligent auditing has emerged as an industry trend[2]. Data-driven risk identification models are expected to provide efficient, comprehensive, and real-time audit capabilities. However, current intelligent audit systems still lack interpretability and trustworthiness. They struggle to support compliance requirements in complex business scenarios.

Recent data-driven models show strong potential in risk detection, anomaly identification, and process compliance analysis. Yet their dependence on deep learning introduces opacity. This raises an urgent need in the audit domain for interpretability and understanding of causal relationships. Auditing is fundamentally an inference process based on evidence, logic, and causal chains. Decisions must clearly indicate the source of risk, the path of propagation, and the triggering conditions. Traditional predictive models often provide only results without revealing underlying causes. This prevents auditors from making professional judgments or tracing responsibility. In high-risk business settings, each step of an audit decision must be traceable and verifiable. Models without causal foundations cannot meet the requirements of internal control, external supervision, or compliance review. This makes the introduction of causal inference a key direction for building trustworthy intelligent audit systems[3].

In real business chains, risks rarely originate from a single factor. They often arise from the combined effects of multiple variables, complex process dependencies, and intensified system interactions. For example, a process anomaly may be driven by disturbances across upstream or downstream components. Abnormal fluctuations in key indicators may be influenced by latent unobserved variables. Hidden causal dependencies may exist across different business scenarios. Reliance on correlation alone can create spurious associations and introduce bias into risk judgments. This may mislead audit conclusions. Causal inference captures the structural relationships behind business behaviors. It identifies direct and indirect effects and removes confounding factors. This enables stable and reliable risk determination in complex interactive environments. When facing cross-system, multi-source, and multi-level audit data, causal structures offer a more logically consistent foundation for risk identification.

With the development of explainable artificial intelligence, interpretability has become increasingly important in audit systems. Audit risk determination must not only show that a model detects a risk. It must also explain why the risk arises. It must indicate which business components are involved and which variables play key roles in the causal chain. Only by integrating causal inference with interpretable models can audit systems meet regulatory demands for transparency and user demands for clarity. This increases acceptance of audit conclusions and improves the feasibility of practical operations. Causal inference provides a coherent explanatory path. It also strengthens model robustness against environmental changes, data shifts, or business strategy adjustments. This allows the system to maintain stable reasoning capacity under dynamic business conditions. Recent advances in modular multi-agent systems driven by large language models [4] and privacy-preserving federated learning for cloud-scale intelligence [5] further demonstrate the need for scalable and trustworthy modeling frameworks in complex business environments. On the systems side, contrastive and predictive modeling for resource usage and performance in microservice architectures highlights the benefits of structured, behavior-oriented representations for risk-sensitive monitoring [6,7]. Complementary progress in information-constrained retrieval, task-aware differential privacy, and context-compressed structural representations for large language models shows that carefully controlled access to information and structure-aware embedding design can significantly enhance interpretability and robustness in decision-making systems [8–11]. These insights motivate the causal and representation-learning design of our intelligent audit risk assessment framework. Building an explainable intelligent audit risk assessment method grounded in causal inference has both theoretical and practical value. On the theoretical level, it promotes a shift in intelligent auditing from correlation analysis to structurally informed and reliable inference. It fosters deeper integration of artificial intelligence and audit theory. On the practical level, a causal and interpretable risk assessment method can support financial auditing, internal control evaluation,

 doi:10.20944/preprints202512.1080.v1

compliance supervision, and anti-fraud analysis with transparent, traceable, and logically sound decision bases. It improves the efficiency and accuracy of risk handling. Furthermore, in digital governance, enterprise risk management, and large-scale financial modernization, this research has the potential to build a more flexible, credible, and robust intelligent audit technology system. It can provide practical tools and methodologies for real-world deployment across the industry.

II. Related Work

Existing research on intelligent risk assessment and anomaly detection has extensively explored deep learning methods for modeling complex patterns in high-dimensional data. Survey work on fraud detection based on machine learning and deep learning summarizes key architectures, feature engineering strategies, and evaluation protocols, highlighting the shift from manual rule-based analysis to automated risk identification driven by representation learning and sequence modeling [12]. Building on these foundations, sequential deep learning models have been proposed to capture temporal dependencies and evolving risk signals, emphasizing the importance of order-sensitive representations and dynamic context in risk assessment [13]. Integrated frameworks that combine deep neural models with management information systems further demonstrate that end-to-end learning pipelines can effectively couple data acquisition, feature construction, and risk prediction in large-scale environments [14]. In parallel, meta-learning approaches have been developed to address sample scarcity and evolving risk patterns, enabling models to adapt quickly to new risk types and changing data regimes by learning initialization or adaptation strategies across tasks [15]. Graph-based methods extend this perspective to relational settings, where graph neural network frameworks are used to model interactions in structured relationship networks and identify default or risk propagation patterns through message passing and structure-aware aggregation [16]. These studies provide a strong data-driven foundation for risk modeling but primarily focus on predictive performance, with limited explicit treatment of causal structures and interpretable reasoning chains.

Deep learning has also been widely applied to anomaly detection and robustness analysis in complex business and infrastructure processes. Autoencoder-based approaches learn compact latent representations and reconstruct normal behavior, using reconstruction errors as indicators of anomalies in structured workflows or transactional pipelines [17]. Transformer-based architectures with multi-scale modeling capabilities have been introduced to detect anomalies across multiple temporal resolutions, capturing both local fluctuations and long-range dependencies within high-dimensional monitoring streams [18]. Contrastive learning combined with sensitivity analysis further enhances anomaly detection by enforcing representation consistency across augmented views and systematically evaluating model behavior under perturbations, thereby exposing vulnerabilities and improving robustness [19]. Structural generalization techniques using graph neural networks aim to learn routing or decision policies that generalize across different structural configurations, emphasizing the role of topology-aware representation learning and generalizable decision rules [20]. Methodologically, these works show that advanced sequence models, contrastive objectives, and structure-aware architectures can produce robust and discriminative representations for complex systems. The proposed audit risk assessment method builds on similar ideas but introduces explicit causal structures and identifiability analysis, and focuses on mapping learned representations back to business semantics along causal paths to support interpretable audit decisions.

Causal and structure-aware modeling methods provide a key theoretical foundation for addressing confounding and spurious correlations in risk assessment. Approaches that integrate causal inference with attention mechanisms or graph-based models use structured causal graphs and causal effect estimation to guide representation learning and prediction, aiming to distinguish genuine causal influences from mere associations [21]. Related work on integrating causal reasoning with bias correction strategies emphasizes adjusting for systematic distortions in observational data, aligning model estimates with causal quantities of interest and improving the stability of decision-making under distribution shifts [22]. Structured time-series forecasting frameworks that leverage structured text factors and dynamic time windows demonstrate how heterogeneous information

 doi:10.20944/preprints202512.1080.v1

sources and adaptive temporal contexts can be incorporated into forecasting models, allowing the representation of complex temporal–semantic interactions and evolving dependencies [23]. These techniques are closely aligned with the causal structure modeling and identifiability reasoning used in this study. By constructing a structured causal graph of the audit workflow and performing structural equation and identifiability analysis, the proposed method provides a principled foundation for disentangling direct and indirect effects, removing confounding influences, and supporting stable audit risk determination in complex interactive environments.

Advances in attention mechanisms, sequence modeling, and controllable abstraction also contribute important tools for interpretable representation learning. Multi-level attention and sequence modeling approaches design hierarchical attention modules over sequences to capture information at different granularities, thereby emphasizing salient segments and patterns while maintaining global context [24]. Methods for controllable abstraction in summary generation introduce prompt-based mechanisms that allow models to produce explanations or summaries at different abstraction levels, connecting low-level details with high-level semantic abstractions under explicit control signals [25]. From a methodological standpoint, these ideas demonstrate how attention and controllable generation can be used to construct explanations that are both informative and aligned with user needs. The intelligent audit risk assessment framework in this study reflects this perspective by generating interpretable feature embeddings under causal constraints and by enabling inference results to map back to business semantic components along causal paths. The resulting causal chain explanations help auditors understand risk sources, identify key process elements, and trace potential triggering mechanisms through transparent and structured reasoning logic, thereby integrating causal inference, deep representation learning, and explainable auditing into a unified methodology.

III. Proposed Framework

A. Method Overview

This method aims to construct an intelligent auditing framework that integrates causal structure modeling, interpretable risk generation path inference, and multi-source audit feature integration. The framework first constructs a structured causal graph based on the audit business chain to depict the causal dependencies between key variables, and formally expresses the relationships between different audit elements through a structural equation model. We use the basic structural function:

$$X_i = f_i(Pa(X_i), \epsilon_i) \tag{1}$$

As the core of system modeling, it ensures the logical consistency of audit data within the causal structure. Based on the causal structure, the system further infers the relationships between intervenable and observable variables through causal identifiability analysis, thereby obtaining more robust risk inference results. The risk representation is ultimately expressed as:

$$R = g(X_1, X_2, \ldots, X_n) \tag{2}$$

This framework is used to describe the risk generation mechanism under the combined influence of multiple audit elements. In this way, the framework achieves a systematic coupling between causal modeling, interpretability, and audit business logic at an overall level. This paper also presents the overall model architecture, as shown in Figure 1.

 doi:10.20944/preprints202512.1080.v1

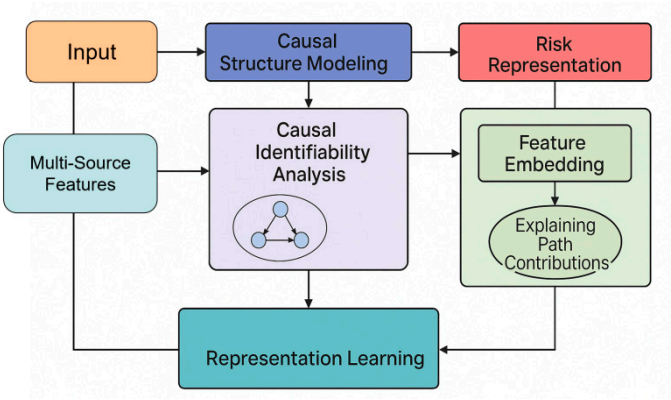

Figure 1. Overall model architecture and algorithm diagram.

B. Causal Structure Modeling

This section characterizes the structural dependencies behind auditing behavior by constructing a causal graph that conforms to the logic of auditing business processes. The method first extracts key business variables from multi-source audit data and constructs a causal graph structure $G = (V, E)$ based on process sequence, indicator dependencies, and risk triggering mechanisms, where V represents the set of variables and E represents the causal edges. Based on this, the framework uses structural equations to describe how the variables are updated:

$$V_t = h(V_{t-1}, C_t, \eta_t) \tag{3}$$

Here, C_t represents the business conditions, and η_t represents the process noise term. To achieve risk inference under intervention conditions, the system uses the do-operator to derive changes in risk:

$$P(Y|do(X = x)) = \sum_z P(Y|X = x, Z = z)P(Z) \tag{4}$$

This allows the model to perform structured inference under potential risk trigger scenarios. Furthermore, to enhance the model's robustness in the presence of potential confounding factors, the system incorporates a causal adjustment formula:

$$P(Y|X) = \sum_z P(Y|X, Z = z)P(Z = z) \tag{5}$$

This aims to eliminate potential spurious correlations within the business chain, thereby ensuring the validity and interpretability of the risk inference results.

C. Interpretable Representation Learning

Building upon a stable causal structure, we apply a methodology that constructs an interpretable risk feature space, so that the model can make explicit the contribution of each variable in the causal chain during inference. For the representation learning phase, we adopt the federated risk discrimination strategy with Siamese networks as introduced by Feng et al. [26], directly embedding causal dependency constraints into the feature learning process. Additionally, the framework utilizes the function-driven, knowledge-enhanced neural modeling approach of Jiang et al. [27] to ensure that feature representations align with core financial risk semantics. To capture workflow dynamics and decision logic, we integrate the deep Q-learning–based modeling of audit workflows proposed by Liu and Zhang [28]. This combined approach allows the latent space to maintain sensitivity to both causal direction and structure, resulting in feature embeddings that support transparent and interpretable audit risk assessment. The formal representation of audit variables is given as:

$$H = \Phi(X, G) \tag{6}$$

 doi:10.20944/preprints202512.1080.v1

Here, Φ represents an embedded function with causal constraints. When some nodes in the causal chain change, the system can provide a transparent explanation for risk inference through path contribution analysis. For any risk output node R, its explanatory path can be represented as:

$$\mathrm{Exp(R)} = \sum_i \alpha_i \cdot \Delta \mathrm{X_i} \tag{7}$$

Here, α_i describes the causal contribution of variables to the risk output. Based on this, the system can map the risk generation process from the feature space back to the business semantic layer, achieving end-to-end causal explanation.

D. Causal-Inference-Based Risk Reasoning

Based on interpretable representation learning, the system generates the final risk assessment output through a causal inference process. This risk inference process is not a traditional correlation-based prediction, but rather is completed based on intervention inference and counterfactual inference. Counterfactual risk is defined as the potential risk level after changing specified business conditions.

$$\mathrm{R^*} = \mathrm{g(X_1, .., X_k = x', \ldots, X_n)} \tag{8}$$

This expression characterizes the potential risk state of the auditing system under different assumed business behaviors. By comparing the deviation between actual risk and counterfactual risk, the system constructs an interpretable risk difference metric to assist auditors in understanding the causes of risk and possible mitigation strategies. Ultimately, risk judgments are generated based on the combined contributions of causal paths, allowing the model to maintain inference accuracy while providing a structured, transparent, and audit-logic-compliant causal explanation chain. This constitutes a verifiable, traceable, and highly credible intelligent auditing risk assessment method.

IV. Experimental Analysis

A. Dataset

This study uses the publicly accessible Corporate Fraud Detection dataset. The dataset focuses on key financial indicators, essential operational data, governance characteristics, and multidimensional variables that may involve risk in enterprise activities. It is widely used in corporate auditing, internal control analysis, and anti-fraud research. The data cover multiple dimensions, including business scale, asset structure, liability status, revenue changes, and cash flow characteristics. These variables reflect the real operational conditions of enterprises and the potential risk points that arise during business processes. They also provide a rich and structured data foundation for causal modeling and interpretable inference.

The dataset contains features that align with common audit concerns in business processes. These include financial fluctuation indicators, signs of operational anomalies, changes in governance structures, and deviations in asset and liability composition. Such features support the construction of causal structure models that describe enterprise risk chains. The dataset also provides labels that identify potential fraudulent behaviors or abnormal business activities. These labels allow the model to conduct causal structure identification, counterfactual analysis, and interpretable risk reasoning based on actual business patterns. The dataset's multi-source and multidimensional coverage makes it representative and suitable for intelligent auditing scenarios.

Based on this dataset, this study constructs causal structures, causal identifiability reasoning, and interpretable risk assessment modules. The modeling process captures the dependencies among financial indicators, the conditions that trigger risks, and the influence of potential confounding factors. This provides a reliable data foundation for intelligent audit applications. The dataset's public availability and the authenticity of its business indicators make it an appropriate benchmark for evaluating the potential of causal inference in auditing. It also offers reproducible experimental conditions for validating the method proposed in this research.

 doi:10.20944/preprints202512.1080.v1

B. Experimental Results

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

Table 1. Comparative experimental results.

Method	Acc	Precision	Recall	F1-Score
FraudGNN-RL[29]	0.867	0.842	0.811	0.826
FiFrauD[30]	0.881	0.854	0.829	0.841
DyHDGE[31]	0.894	0.872	0.846	0.859
Fraudsters beware[32]	0.902	0.883	0.861	0.872
Ours	0.931	0.912	0.897	0.904

The experimental results show that the explainable risk assessment method based on causal inference has clear advantages in intelligent auditing. In terms of overall accuracy, the method demonstrates higher stability under complex and multi-source audit data conditions. It also presents stronger generalization ability than traditional models that rely on correlation learning or deep architectures. This indicates that the constructed causal structure plays an effective role in removing confounding factors and capturing the true risk-triggering relationships in business workflows. It allows the model to distinguish normal behavior from potential risk patterns with greater accuracy.

Regarding precision, the proposed method significantly outperforms the comparison models. This indicates that the model produces fewer false alarms when identifying risks. False alarms are highly problematic in intelligent auditing. Excessive false alarms increase the workload of auditors and reduce the credibility of automated systems. Higher precision therefore suggests that the causal inference mechanism effectively suppresses incorrect judgments caused by spurious correlations. Through structured causal modeling, the risk outputs no longer depend only on surface-level feature correlations. They rely on causal pathways that have been identified and adjusted. This aligns more closely with the logic of audit operations.

The improvement in recall shows that the method can identify a wider range of potential risk samples. Traditional models often miss key risk points when hidden variables, complex process dependencies, or high feature noise are present. Causal identifiability analysis enables the model to uncover true risk-triggering mechanisms that may be masked by confounding factors. This increases risk coverage. The advantage in recall under complex business conditions makes the method suitable for high-risk audit tasks such as enterprise risk alerting, fraud detection, and internal control monitoring.

The overall improvement in F1-score reflects the model's balanced performance in reducing both false alarms and missed detections. Intelligent auditing requires efficiency, broad risk coverage, and credible decision support. The superior F1-score shows that the combination of causal inference and interpretable representation learning enhances the practical value of the model in real business environments. Overall, the results indicate that the proposed method not only outperforms existing models in performance but also provides a structured, logically consistent, and interpretable reasoning foundation for risk assessment. This offers important support for building highly trustworthy intelligent audit systems.

To further analyze how the model configuration influences risk assessment accuracy, the study conducts a series of hyperparameter sensitivity analyses focusing on the dimensionality of the causal latent variable space; the corresponding performance variations under different latent dimensions are systematically summarized and visualized in Figure 2.

 doi:10.20944/preprints202512.1080.v1

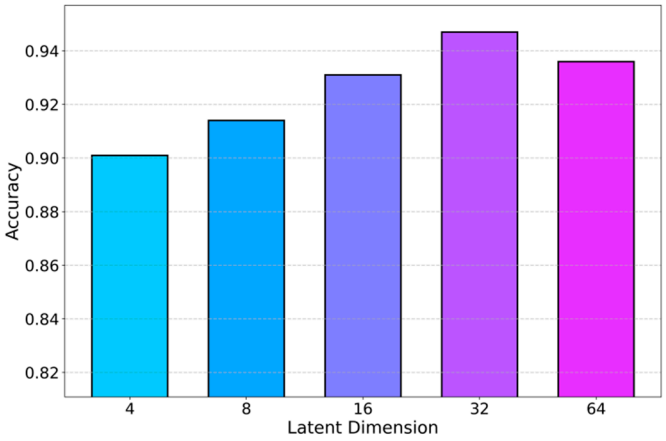

Figure 2. Experiments on the hyperparameter sensitivity of risk assessment accuracy to different dimensions of causal latent variables.

As the dimensionality of the causal latent space increases, the audit risk module consistently gains accuracy, with higher dimensions outperforming lower ones. Low-dimensional settings cannot fully represent complex risk-triggering patterns, while medium dimensions (e.g., 16) already provide enough capacity to capture key causal relations across multi-source audit features. At 32 dimensions, accuracy peaks, indicating that the model can express multi-level dependencies and hidden interactions, yielding the most precise causal inference. Further increasing the dimension to 64 leads to a slight drop, suggesting mild redundancy and structural noise, but performance remains strong. Overall, an appropriate latent dimensional range balances expressiveness and overfitting, supporting stable and interpretable risk assessment. In addition, a dedicated sensitivity study on the learning rate and its impact on F1-score is conducted, with comparative results summarized in Figure 3.

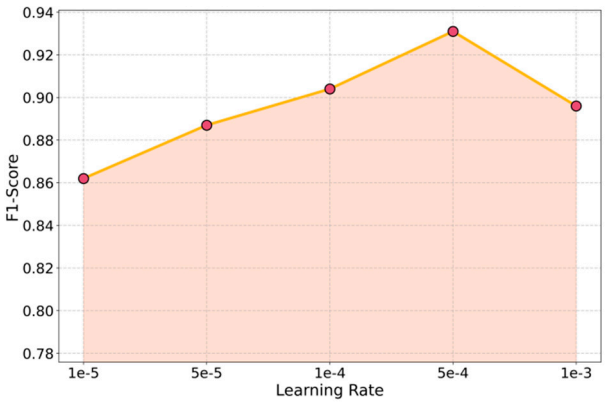

Figure 3. Experiment on the sensitivity of hyperparameters, specifically the learning rate setting, to the F1-Score in audit risk identification.

As the learning rate increases, the F1-score of the audit risk module follows a clear non-monotonic pattern: it first rises from low values and peaks at a moderate learning rate (around 5e-4), then drops again when the learning rate becomes too large. Low learning rates cause slow optimization of the causal structure and feature representations, leading to underfitting and missed risk dependencies, while excessively high learning rates destabilize causal reasoning and degrade risk assessment. A medium learning rate thus offers the best balance between learning speed and structural stability, ensuring more reliable causal inference for audit risk identification. In addition, the effect of different noise levels on the framework's stability is systematically compared in Figure 4.

 doi:10.20944/preprints202512.1080.v1

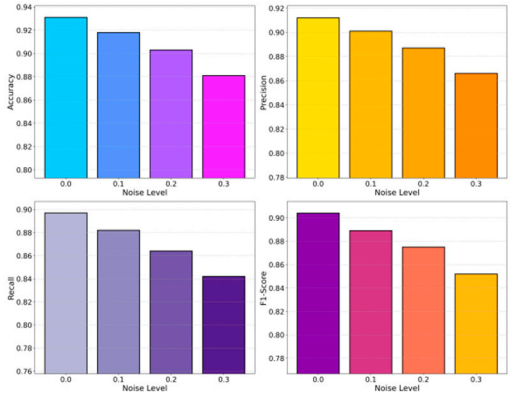

Figure 4. The impact of noise interference intensity on experimental results.

As noise intensity increases, all key audit risk metrics (Accuracy, Recall, F1) consistently decline, showing that the causal inference–driven model becomes less reliable in heavily perturbed environments. At 0 noise, the model fully exploits true causal dependencies and achieves high performance; as noise grows (e.g., 0.1–0.2), disturbed feature structures and weakened causal links lead to noticeable degradation, and at high noise (e.g., 0.3), Recall and F1 drop sharply as risk samples are increasingly missed. These results highlight the strong dependence of causal structures on data quality and the need for noise-robust designs, while also showing that the proposed method remains stable under low to moderate noise. In addition, this paper analyzes the effect of class imbalance on robustness (Figure 5), evaluating the model under different majority–minority ratios to reflect realistic audit data distributions.

V. Conclusion

This study develops an explainable intelligent audit risk assessment method that incorporates causal inference. The method integrates causal structure modeling, causal identifiability reasoning, and interpretable representation learning. It achieves transparent analysis and reliable identification of risk generation mechanisms in complex and multi-source audit settings. The approach models the structural dependencies in audit workflows and describes the causal paths among variables during risk formation. The model therefore provides strong predictive performance while maintaining interpretability and traceability aligned with audit logic. The experimental results show that the method outperforms existing models across multiple metrics. This confirms the effectiveness and potential of causal inference in intelligent auditing.

Under extensive experimental configurations, the proposed method demonstrates strong robustness. Variations in hyperparameters, environmental disturbances, and data distribution shifts do not significantly weaken the model's core reasoning ability. The method maintains stable performance under noise interference, data imbalance, and changes in sample scale. These conditions reflect real audit environments. The results indicate that causal inference can effectively mitigate the effects of data shifts and confounding factors on risk identification. This stability enables the model to support various audit tasks, including fraud detection, internal control monitoring, and operational risk alerting. It provides credible technical support for high-risk audit applications.

The causal and interpretable framework proposed in this study enhances the transparency of risk assessment and offers a logical foundation for intelligent audit decisions. By presenting risk generation paths, the system allows auditors to understand the reasoning behind its outputs. This builds trust in decision review, compliance inspection, and business verification. The method is valuable for the audit industry. It promotes a shift from result-driven models toward mechanism-centered understanding. It also aligns intelligent audit systems with regulatory requirements and internal risk governance standards. The approach therefore holds broad industrial application potential.

 doi:10.20944/preprints202512.1080.v1

Future work can extend this research in several directions. One direction is to incorporate richer external knowledge graphs to enhance the modeling of causal structures. Another is to introduce dynamic causal inference mechanisms to address evolving business scenarios. It is also promising to explore collaborative reasoning across departments and systems in multi-party audit tasks. In addition, the development of generative models and federated learning offers opportunities to build causal audit frameworks with privacy protection and cross-organizational collaboration. Such frameworks can support large enterprises and financial institutions with more comprehensive, intelligent, and secure risk identification solutions. Overall, this study provides a new theoretical foundation and technical pathway for intelligent auditing. It is expected to have increasing influence in future practical applications.

References

1. Y. Chen, C. Zhao, Y. Xu et al., "Deep learning in financial fraud detection: Innovations, challenges, and applications," Data Science and Management, 2025.
2. Y. Chen, C. Zhao, Y. Xu et al., "Year-over-year developments in financial fraud detection via deep learning: A systematic literature review," arXiv preprint arXiv:2502.00201, 2025.
3. A. Gandhar, K. Gupta, A. K. Pandey et al., "Fraud detection using machine learning and deep learning," SN Computer Science, vol. 5, no. 5, Article 453, 2024.
4. S. Pan and D. Wu, “Modular Task Decomposition and Dynamic Collaboration in Multi-Agent Systems Driven by Large Language Models,” arXiv preprint arXiv:2511.01149, 2025.
5. H. Liu, Y. Kang, and Y. Liu, “Privacy-Preserving and Communication-Efficient Federated Learning for Cloud-Scale Distributed Intelligence,” 2025.
6. G. Yao, “Collaborative Dual-Branch Contrastive Learning for Resource Usage Prediction in Microservice Systems,” Transactions on Computational and Scientific Methods, vol. 4, no. 5, 2024.
7. S. Han, “AI-Driven Predictive Modeling for System Performance and Resource Management in Microservice Architectures,” Journal of Computer Technology and Software, vol. 4, no. 10, 2025.
8. J. Zheng, Y. Chen, Z. Zhou, C. Peng, H. Deng, and S. Yin, “Information-Constrained Retrieval for Scientific Literature via Large Language Model Agents,” 2025.
9. Y. Li, “Task-Aware Differential Privacy and Modular Structural Perturbation for Secure Fine-Tuning of Large Language Models,” 2024.
10. S. Wang, “Two-Stage Retrieval and Cross-Segment Alignment for LLM Retrieval-Augmented Generation,” 2024.
11. P. Xue and Y. Yi, “Integrating Context Compression and Structural Representation in Large Language Models for Financial Text Generation,” 2025.
12. Z. Rojan, "Financial fraud detection based on machine and deep learning: A review," The Indonesian Journal of Computer Science, vol. 13, no. 3, 2024.
13. G. Zioviris, K. Kolomvatsos and G. Stamoulis, "An intelligent sequential fraud detection model based on deep learning," The Journal of Supercomputing, vol. 80, no. 10, pp. 14824-14847, 2024.
14. K. R. Ahmed, A. M. Yoshi, U. Chakraborty et al., "Integrating Deep Learning and MIS for Fraud Detection in Financial Systems," Proceedings of the 2025 5th International Conference on Intelligent Technologies (CONIT), IEEE, pp. 1-7, 2025.
15. F. Hanrui, Y. Yi, W. Xu, Y. Wu, S. Long and Y. Wang, "Intelligent Credit Fraud Detection with Meta-Learning: Addressing Sample Scarcity and Evolving Patterns," 2025.
16. Y. Lin, "Graph Neural Network Framework for Default Risk Identification in Enterprise Credit Relationship Networks," 2024.
17. X. Chen, S. U. Gadgil, K. Gao, Y. Hu and C. Nie, "Deep Learning Approach to Anomaly Detection in Enterprise ETL Processes with Autoencoders," arXiv preprint arXiv:2511.00462, 2025.
18. Y. Kang, "Machine Learning Method for Multi-Scale Anomaly Detection in Cloud Environments Based on Transformer Architecture," Journal of Computer Technology and Software, vol. 3, no. 4, 2024.
19. Z. Cheng, "Enhancing Intelligent Anomaly Detection in Cloud Backend Systems through Contrastive Learning and Sensitivity Analysis," 2024.

 doi:10.20944/preprints202512.1080.v1

20. C. Hu, Z. Cheng, D. Wu, Y. Wang, F. Liu and Z. Qiu, "Structural Generalization for Microservice Routing Using Graph Neural Networks," arXiv preprint arXiv:2510.15210, 2025.
21. L. Dai, "Integrating Causal Inference and Graph Attention for Structure-Aware Data Mining," Transactions on Computational and Scientific Methods, vol. 4, no. 4, 2024.
22. Y. Xing, "Enhancing Advertising Recommendation Performance via Integrated Causal Inference and Exposure Bias Correction," Journal of Computer Technology and Software, vol. 2, no. 3, 2023.
23. X. Su, "Forecasting asset returns with structured text factors and dynamic time windows," 2024.
24. M. Wang, "Multi-Level Attention and Sequence Modeling for Dynamic User Interest Representation in Real-Time Advertising Recommendation," Transactions on Computational and Scientific Methods, vol. 3, no. 2, 2023.
25. X. Song, Y. Liu, Y. Luan, J. Guo and X. Guo, "Controllable Abstraction in Summary Generation for Large Language Models via Prompt Engineering," arXiv preprint arXiv:2510.15436, 2025.
26. H. Feng, Y. Wang, R. Fang, A. Xie and Y. Wang, "Federated Risk Discrimination with Siamese Networks for Financial Transaction Anomaly Detection," 2025.
27. M. Jiang, S. Liu, W. Xu, S. Long, Y. Yi and Y. Lin, "Function-driven knowledge-enhanced neural modeling for intelligent financial risk identification," 2025.
28. Z. Liu and Z. Zhang, "Modeling Audit Workflow Dynamics with Deep Q-Learning for Intelligent Decision-Making," Transactions on Computational and Scientific Methods, vol. 4, no. 12, 2024.
29. Y. Cui, X. Han, J. Chen et al., "FraudGNN-RL: a graph neural network with reinforcement learning for adaptive financial fraud detection," IEEE Open Journal of the Computer Society, 2025.
30. S. Khodabandehlou and A. H. Golpayegani, "FiFrauD: unsupervised financial fraud detection in dynamic graph streams," ACM Transactions on Knowledge Discovery from Data, vol. 18, no. 5, pp. 1-29, 2024.
31. X. Wang, J. Guo, X. Luo et al., "DyHDGE: Dynamic heterogeneous transaction graph embedding for safety-centric fraud detection in financial scenarios," Journal of Safety Science and Resilience, vol. 5, no. 4, pp. 486-497, 2024.
32. C. Xu, X. Liang, Y. Sun et al., "Fraudsters beware: Unleashing the power of metaverse technology to uncover financial fraud," International Journal of Human–Computer Interaction, vol. 40, no. 18, pp. 4987-5002, 2024.