Journal of Computer Technology and Software
ISSN:2998-2383
Vol. 4, No. 9, 2025

Deep Graph Learning for Concealed Accounting Fraud Detection in Enterprise Networks

Xuhao Meng[1*]
[1]Columbia University, New York City, USA
*Corresponding author: Xuhao Meng; violetmeng33@gmail.com

Abstract: As corporate operations become increasingly networked and digitalized, fund flows, business interactions, and control relations intertwine across multiple entities to form complex transaction networks. Accounting fraud, therefore, exhibits more structured, collaborative, and concealed characteristics, which pose substantial challenges to traditional methods based on individual features or rule-driven detection. This study addresses accounting fraud detection in complex corporate transaction networks by modeling firms and their transaction relations within a unified graph structure. A graph neural network is introduced to jointly model node attributes and relational context. Fraudulent behavior is characterized from a structural perspective through its organizational patterns and risk propagation paths in the network. A multi-layer message passing mechanism integrates local transaction features with higher-order structural dependencies. The model can automatically learn latent anomaly patterns and relational cues without relying on manual rules. On this basis, a unified risk scoring framework maps structured representations to entity-level fraud risk intensity. Systematic identification of abnormal behavior in complex transaction environments is thus achieved. Comparative evaluation with representative methods shows that the proposed approach delivers more stable overall discrimination and risk differentiation. It effectively reduces misclassification and improves the detection of concealed fraudulent entities. Further analysis confirms the suitability of this structured modeling paradigm for complex relational scenarios. Integrating transaction network semantics with deep representation learning helps overcome the expressive limitations of traditional accounting analysis. This work provides a unified network-oriented modeling approach for intelligent auditing and financial risk analysis and offers a valuable reference for anomaly detection in complex economic systems.

Keywords: Enterprise transaction networks; accounting risk identification; graph structure modeling; deep representation learning

1. Introduction

In the context of deep integration between digital operations and information systems, corporate economic activities are expanding with unprecedented scale and complexity. Capital flows, business processes, contractual relationships, and control structures are intertwined across multiple platforms and multiple actors. They gradually form highly coupled corporate transaction networks. These networks cross organizational boundaries and exhibit strong heterogeneity, dynamics, and concealment. Traditional accounting analysis, centered on individual transactions or static financial indicators, can no longer fully capture economic substance. In such settings, accounting fraud rarely appears as an isolated event. It is often embedded in complex transactional structures and concealed through related party transactions, circular fund flows, and shell entity transfers, which substantially increase detection and regulatory difficulty[1].

From the perspective of risk governance and financial transparency, accounting fraud distorts the quality of financial information and may trigger distorted investment decisions, resource misallocation, and even systemic financial risk. As organizational forms and business models continue to evolve, fraudulent behavior increasingly shows networked, chained, and strategic characteristics. Methods relying solely on rule matching, ratio analysis, or human judgment usually detect only ex post or superficial anomalies. They struggle to uncover latent risk patterns hidden deep within transaction structures. As a result, modeling interactions among entities in high-dimensional and multi-relational environments and understanding how abnormal behaviors propagate and coordinate within networks has become a critical challenge in accounting, audit, and risk identification.

Corporate transaction networks provide a structured perspective for understanding accounting fraud. By representing firms, accounts, projects, or affiliated entities as nodes and modeling fund transfers, business transactions, or control relationships as edges, the overall organization of economic activities can be more faithfully captured. Under this representation, fraud is not merely reflected by abnormal values at individual nodes. It often manifests as anomalous combinations of local substructures, asymmetric relational patterns, or abrupt changes in network evolution. However, real-world transaction networks are typically large-scale, multi-typed, and highly imbalanced in structure. Traditional graph statistics or handcrafted feature methods face difficulties

in achieving both expressive power and generalization, which limits their effectiveness in complex scenarios[2].

Graph neural networks offer a new theoretical and methodological foundation for modeling complex transaction networks. By performing information propagation and representation learning directly on graph structures, these methods can automatically integrate node attributes with structural context without relying on explicit rules. They are capable of capturing multi-level association patterns and latent dependencies. For accounting fraud detection, this capability enables understanding the behavioral position of an enterprise or transaction within the overall network. It also supports identifying covert risk patterns formed through multi-actor coordination while jointly modeling structural and semantic anomalies. This paradigm moves beyond analyses focused on local numerical features or isolated samples and enables a more systematic characterization of complex fraudulent behavior.

Against this background, research on graph neural network-based accounting fraud detection for complex corporate transaction networks carries significant theoretical and practical value[3]. At the methodological level, it promotes a shift in accounting analysis from static and linear views toward structured and relational perspectives, which deepens understanding of the internal mechanisms of fraud. At the application level, it provides scalable and transferable technical support for intelligent auditing, risk early warning, and regulatory technology. It also enhances the ability to identify highly concealed and structurally complex fraud. By incorporating graph-based modeling and deep representation learning, this line of research is expected to offer a more robust and forward-looking analytical framework for financial risk governance in complex economic systems.

2. Related work

Early research on accounting fraud detection can largely be traced to methods based on financial indicators and statistical analysis. These studies typically rely on structured financial variables such as leverage ratios, profitability measures, and cash flow ratios. Threshold rules, discriminant analysis, or regression models are used to describe differences between normal firms and fraudulent firms. Such methods were practical when information environments were relatively simple, and fraud patterns were limited. Their core assumption is that fraud leads to clear, stable, and interpretable numerical deviations in financial statements[4]. As transaction structures become more complex, this assumption weakens. Fraudulent behavior is often dispersed across multiple transactions and multiple entities. As a result, financial indicators of a single firm may not appear abnormal, which reduces the effectiveness of traditional approaches.

With the development of machine learning, research has gradually introduced supervised learning and anomaly detection models. Fraud detection is framed as a classification or risk scoring problem[5]. These methods extract features from transaction records, account behaviors, or financial texts. Nonlinear models are then used to learn complex decision boundaries, which partially overcomes the limits of manual rules. However, most studies still adopt the assumption of sample independence. They focus on the attribute distribution of individual firms or single transactions. Widespread related party transactions, control relationships, and fund circulation among firms are largely ignored. By neglecting relational structure, these models struggle to identify collusive fraud, group-level manipulation, or cross-entity risk propagation. This limits their applicability in real-world networked environments.

To address these limitations, some studies have introduced network analysis ideas and explicitly modeled transaction relationships among firms or accounts. Graph statistics and network metrics are then used for fraud detection. Typical features include node degree, clustering coefficient, centrality, and community structure[6]. These methods can reveal abnormal relational patterns and structural imbalance to some extent. However, they usually depend on manually designed structural features and have limited expressive power. They also struggle to capture higher-order dependencies in multi-relational and multi-scale networks. In addition, network features are often tightly coupled with specific business settings. Model transferability and generalization are, therefore, weak. This makes it difficult to cope with evolving transaction structures and changing fraud strategies.

In recent years, graph representation learning and graph neural networks have provided a new direction for accounting fraud detection. Compared with traditional network analysis, these studies emphasize end-to-end representation learning on graph structures. Multi-layer information propagation mechanisms are used to integrate node attributes with neighborhood structure. Latent risk patterns can thus be learned automatically[7]. Existing work shows clear advantages in modeling complex relational dependencies, identifying structured anomalies, and reducing reliance on feature engineering. However, most studies focus on generic relational modeling or assume static networks. Limited attention is given to heterogeneous relations, structural noise, and strategic disguise of fraud that are common in corporate transaction networks. How to better exploit structural information in transaction networks under the accounting context and how to design graph neural network models tailored to complex economic systems remain open and important research questions.

3. Proposed Framework

In this paper, we model complex enterprise transaction systems as a directed heterogeneous graph to uniformly characterize the enterprise entities and their transaction relationships. Let the transaction network be represented as:

$$G = Attention(V, E, X)$$

In this modeling approach, $V = \{v_1, \ldots, v_n\}$ represents the set of economic entity nodes such as enterprises or accounts, $E \subseteq V \times V$ represents the set of relational edges

such as transactions, control, or capital flows, and $X \in R^{N \times d}$ is a node attribute matrix used to describe the basic accounting characteristics or transaction statistics of the entities. Each edge $e_{ij} \in E$ represents an economic connection from node v_i to node v_j, and its directionality reflects the flow of funds or business. This modeling approach can preserve the real interaction relationships between enterprises at the structural level, providing a unified input for subsequent structured representation learning. This paper also presents the overall model architecture, as shown in Figure 1.

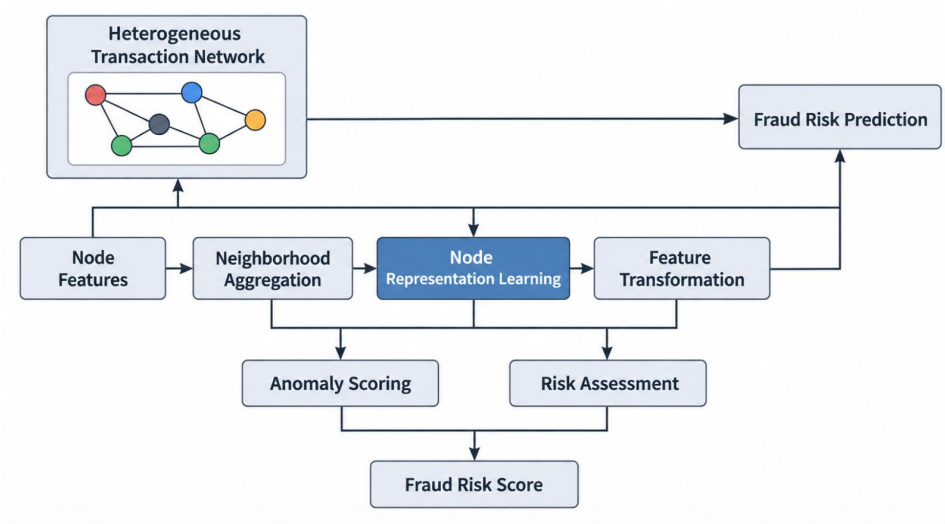

Figure 1. Overall model architecture

In the graph representation learning process, node representations evolve layer by layer through the aggregation and updating of neighborhood information. Let $h_i^{(l)}$ be the hidden representation of node v_i in layer l, and the initial state is given by the node attributes, i.e.:

$$h_i^{(0)} = x_i$$

In each layer, a node first collects information from its neighborhood $N(i)$ and forms an intermediate representation through weighted aggregation:

$$\tilde{h}_i^{(l)} = \sum_{j \in N(i)} \alpha_{ij}^{(l)} h_j^{(l)}$$

Here, $\alpha_{ij}^{(l)}$ represents the influence weight of node v_j on node v_i, used to characterize the relative importance of different transaction relationships in structure propagation. This weighting mechanism helps mitigate the interference of noisy edges and weak associations on representation learning.

After neighborhood information aggregation, node representations are updated through linear mapping and nonlinear transformations to enhance the representation of the feature space. Specifically, the node representation of layer $l+1$ is defined as follows:

$$h_i^{(l+1)} = \sigma(W^{(l)} \tilde{h}_i^{(l)} + b^{(l)})$$

Here, $W^{(l)}$ and $b^{(l)}$ are the learnable weight matrix and bias term, respectively, and $\sigma(\cdot)$ represents the element-wise nonlinear activation function. Through multi-layer stacking, this update process can gradually integrate higher-order structural dependencies, enabling node representations to simultaneously encode local transaction behavior and broader network context information, thereby characterizing potential structural anomaly patterns.

After obtaining the node representation $h_i^{(L)}$ of the final layer, the model maps it to an accounting fraud risk score, used to characterize the potential degree of abnormality of the entity. Specifically, the risk prediction function is defined as:

$$s_i = sigmoid(w^T h_i^{(L)})$$

Here, w is a learnable parameter vector, and s_i represents the fraud risk intensity of node v_i. This continuous scoring format helps reflect the relative risk positions of different entities in the transaction network, rather than a simple binary division. By unifying the modeling of structural information, node attributes, and relationship propagation mechanisms, the proposed method can characterize the distribution of accounting risks in complex transaction environments at the overall network level, providing structured support for subsequent risk analysis and regulatory decisions.

4. Experimental Analysis

4.1 Dataset

This study adopts the SEC Accounting and Auditing Enforcement Releases dataset as the data source for accounting fraud detection, including financial misstatements and disclosure violations. AAER consists of enforcement releases issued by the United States regulator for violations related to financial reporting, auditing, and accounting

practices. The releases describe violation types, involved entities, and event details. In academic research, AAER is widely regarded as an authoritative proxy label for accounting fraud or financial misreporting. It is commonly used to construct supervision signals and risk event samples.

Regarding data scale and temporal coverage, the publicly curated AAER dataset aggregates enforcement releases issued between 1982 and 2021. The documents are organized at the level of corporate misstatement events and provide structured records that link AAER documents to specific events, together with their time spans. The current version contains approximately 4,278 AAER documents and corresponds to about 1,816 corporate misstatement events. It covers multiple industries and long time horizons. This makes it suitable for analyzing fraud risk under different economic cycles and regulatory environments.

To align with the theme of complex corporate transaction networks, this study treats firms or entities as nodes and constructs relational edges based on information that can be structurally extracted or mapped from AAER. Examples include entities jointly involved in the same misstatement event, repeated co-occurrences across different releases, and relational transaction or control clues aligned with public disclosures. This process yields graph-structured inputs for accounting fraud risk learning. Event-level violation labels provided by AAER are used to define node-level risk supervision. The graph structure captures potential collusive fraud, risk propagation, and chained operations among entities. This enables the model to learn more robust fraud representations within a structured relational context.

4.2 Experimental Results

This article first presents the results of the comparative experiments, as shown in Table 1.

Table 1. Comparative experimental results

Method	Acc	Precision	Recall	F1-Score
Qfnn-ffd[8]	0.842	0.835	0.819	0.827
FiFrauD[9]	0.867	0.861	0.846	0.853
FraudGNN-RL[10]	0.892	0.887	0.874	0.880
FraudGT[11]	0.904	0.899	0.886	0.892
Finstack-net[12]	0.911	0.907	0.895	0.901
Ours	0.927	0.923	0.912	0.917

As shown by the comparison in Table 1, the overall discriminative ability of the methods exhibits a progression from shallow to more advanced modeling. Earlier approaches rely mainly on local features or a single relational cue. They can capture some explicit anomaly patterns. Yet, in structured fraud scenarios within complex corporate transaction networks, they often fail to fully recover risk signals caused by multi-entity coordination, chain-like transmission, and concealed associations. As stronger relational modeling and representation learning mechanisms are introduced, models make more effective use of network context. Performance improves more stably. This trend highlights the critical value of transaction network structure in accounting fraud detection.

From the perspective of precision, the observed improvement indicates that stronger graph structural characterization helps reduce false positives. In accounting fraud detection, a false positive often means that compliant but structurally complex corporate relationships are misclassified as abnormal. This leads to wasted audit resources and biased risk handling. Our method maintains higher precision in the comparison. This suggests that the model better separates normal high-frequency transactions or intra-group fund movements from truly anomalous structural patterns. It remains effective even in corporate transaction networks with multi-level relations and noisy edges. It can focus on risk cues with clearer interpretability.

The joint performance on recall and F1 shows that our method not only reduces false positives but also decreases false negatives more effectively. This reflects more comprehensive coverage of concealed fraud patterns. Complex fraud is often implemented through dispersed transactions, cyclic paths, and coordinated actions across multiple entities. These strategies dilute single-point anomalies and make models based on isolated samples or limited neighborhoods more likely to miss cases. Higher recall and better F1 imply that the method aggregates information over a broader structural range. It can identify cross-node coordination and anomalous substructures. Risks that are deliberately disguised become more salient in the learned structural representations, which improves overall robustness.

The advantage in overall accuracy further suggests that the method achieves a better balance across metrics. It preserves detection coverage while maintaining reliable decisions. This balance is especially important for accounting fraud detection in complex corporate transaction networks. Real-world settings often involve class imbalance, heterogeneous relations, structural noise, and strategic adversarial behaviors. The results in the table indicate that our framework effectively integrates node attributes with relational structure. It forms more discriminative representations at both semantic and topological levels. It therefore delivers more stable and consistent overall performance under a unified risk scoring objective.

The number of layers in a graph neural network determines the depth of information propagation in the transaction network, thus affecting the model's ability to represent higher-order relationships and potential risk transmission chains. To verify the model's dependence on the depth of structural propagation, we examined the changing trends of our

algorithm on various evaluation metrics under different layer settings, and the experimental results are shown in Figure 2.

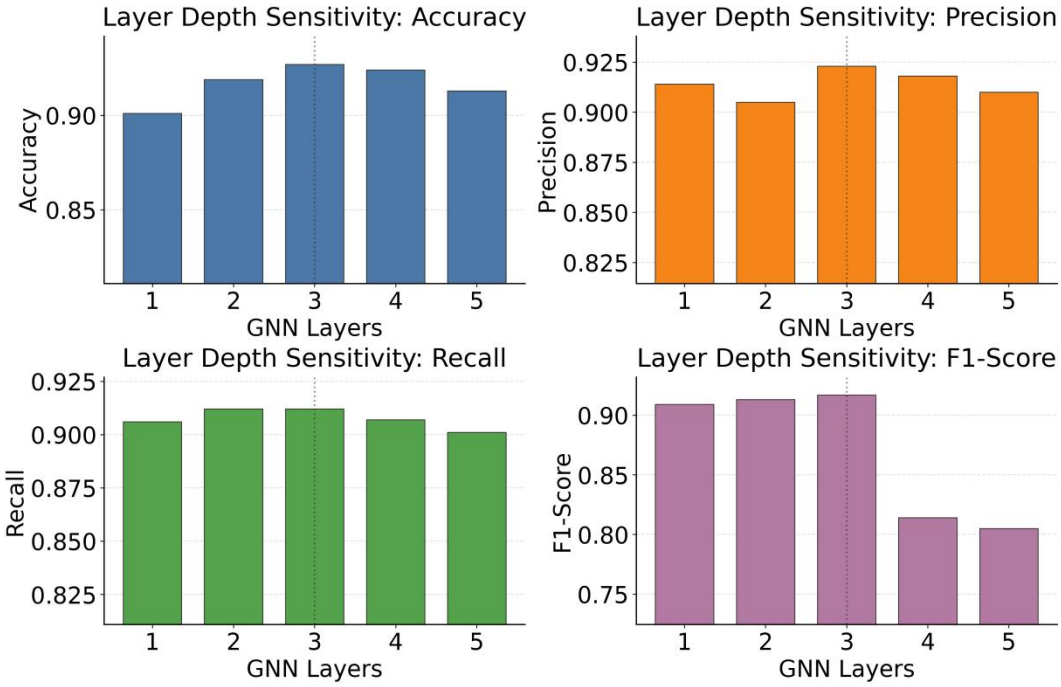

Figure 2. The impact of the number of layers in a graph neural network on experimental results

The overall patterns across the four subplots indicate that the effect of network depth is not linear or monotonic. It reflects a trade-off between information propagation depth and noise accumulation. In dense corporate transaction networks with many weakly related edges, shallow models rely more on local neighborhood signals. They preserve sensitivity to direct transaction relations but provide limited modeling of cross-entity and cross-chain risk propagation. As the number of layers increases, the model accesses higher-order structural context. It becomes more capable of capturing chain transactions, cyclic paths, and multi-entity coordination patterns. As a result, performance tends to improve first and then gradually saturate.

For accuracy and precision, the curves show a clear advantage at intermediate depths. This suggests that moderate propagation depth helps reduce misclassification. Accounting fraud detection places strong emphasis on false positive control, since excessive false alarms consume audit resources on non-critical entities. With a moderate number of layers, the model absorbs transaction structure semantics within a reasonable range. At the same time, it avoids dilution of decision boundaries caused by noisy distant neighbors. This enables better separation between normal complex fund flows and truly anomalous structural patterns. It reflects a stronger focus on structured risk cues.

The recall subplot exhibits smaller fluctuations and shows relative robustness to depth variation. This usually indicates that coverage of potential anomalous entities depends less on propagation depth than precise discrimination does. In other words, deeper structural propagation mainly contributes to reducing false positives and improving decision reliability, rather than simply expanding the range of detectable anomalies. This aligns with fraud characteristics in complex transaction networks. Fraudulent entities are often embedded in locally dense transaction communities. Limited neighborhood expansion is sufficient to expose structural abnormalities. Excessive expansion may instead introduce many irrelevant neighbors and weaken feature alignment and structural signals.

The F1 score subplot shows a clear decline at larger depths. This suggests that excessive propagation may cause over-smoothing or noise diffusion. The balance between precision and recall is then disrupted. In corporate transaction networks, very deep message passing tends to pull representations of different communities closer together. The distinction between normal and anomalous entities in embedding space becomes weaker. Distant transaction relations unrelated to fraud may also enter the decision process and amplify incidental co-occurrence noise. These sensitivity results therefore support the use of a moderate depth setting. It provides a more reasonable compromise between modeling higher-order risk chains and maintaining robust discrimination. This choice better matches the practical requirements of reliability and operability in accounting fraud detection.

Varying the attention head configuration changes how attention is distributed across the transaction network. This directly influences the aggregation strategy for multi-relational interactions and local structural cues. To examine how our algorithm depends on the granularity of attention decomposition, we track the model's output trend under

different head-count settings. The experimental results are shown in Figure 3.

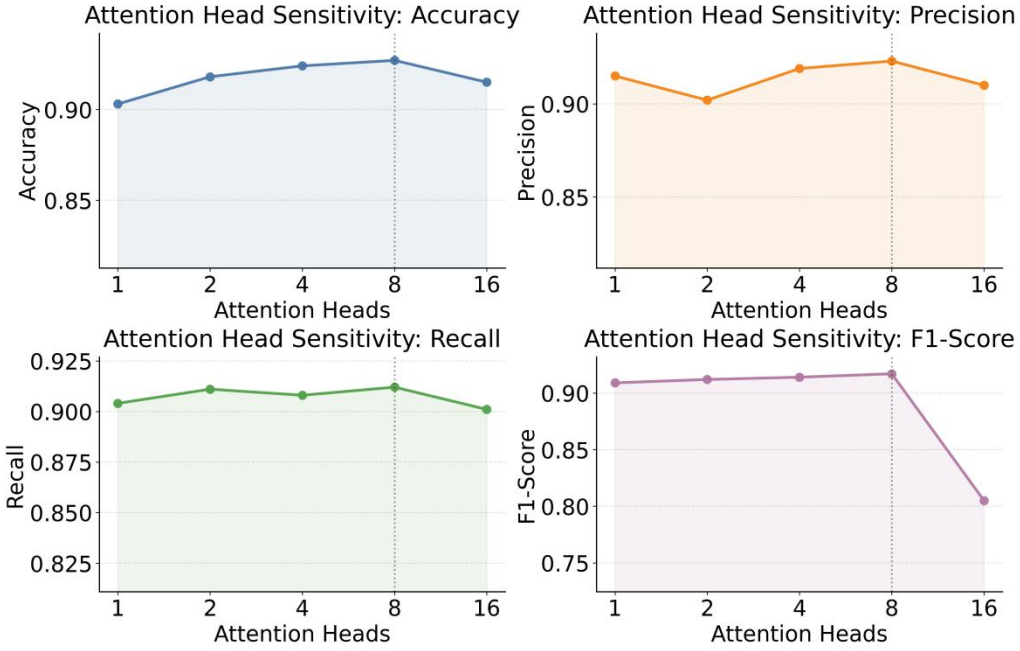

Figure 3. The effect of the number of attention heads on the experimental results

Variation in the number of attention heads has a pronounced impact on how the model represents multiple relations in corporate transaction networks. With fewer heads, the model is forced to compress diverse interaction signals into a limited subspace. It therefore focuses more on dominant relations or high-frequency transaction patterns. As the number of heads increases, the model can attend to different structural cues in parallel. This facilitates the separation of heterogeneous information such as fund flow directions, transaction intensity, and relational structure. Potential risk semantics in complex transaction networks are thus expressed more fully.

For several metrics, the curves show an initial improvement followed by a plateau. This indicates the existence of a reasonable range for attention decomposition granularity. Within this range, the multi-head mechanism effectively reduces bias introduced by a single attention view. It allows more balanced information aggregation across different neighborhoods and relational patterns. This behavior aligns well with accounting fraud scenarios. Fraud is rarely triggered by a single structural feature. It usually emerges from the coordination of multiple weak signals in transaction networks. Moderate multi-head attention helps capture these dispersed yet related anomaly cues.

As the number of attention heads continues to increase, a clear decline can be observed in some subplots. This suggests that overly fine attention partitioning may introduce redundant or noisy subspaces. In transaction networks, too many attention heads may excessively split structurally related signals. Each attention branch then receives only limited information. The ability to model complete risk patterns is weakened. Over decomposition may also amplify incidental co-occurrence or weak relations. This interferes with decision-making and reduces overall stability.

Overall, these results indicate that more attention heads do not necessarily lead to better modeling in complex corporate transaction networks. The number should match network scale, relation density, and risk structure complexity. A moderate configuration achieves a balance between multi-relational modeling and information aggregation efficiency. The model can distinguish different types of transaction interactions without dispersing focus due to structural noise. This finding further supports the necessity of the proposed structural design. It also highlights the critical role of appropriate attention granularity in improving the reliability and practical usefulness of accounting fraud detection.

Adjusting the neighborhood sampling size directly changes how much structural context is visible at each message-passing step. This, in turn, affects how the model aggregates local information and higher-order dependencies in the transaction network. The experimental results are shown in Figure 4.

The number of sampled neighbors directly affects the model's ability to perceive structure in corporate transaction networks. When the sampling size is small, each message passing step observes only a limited local neighborhood. The model relies mainly on direct transactions or one-hop neighbors for decision-making. This constrains its understanding of complex relational structures. As the sampling size increases, richer contextual information becomes available. Potential risk patterns hidden in multi-

entity interactions and indirect paths are more easily encoded into node representations.

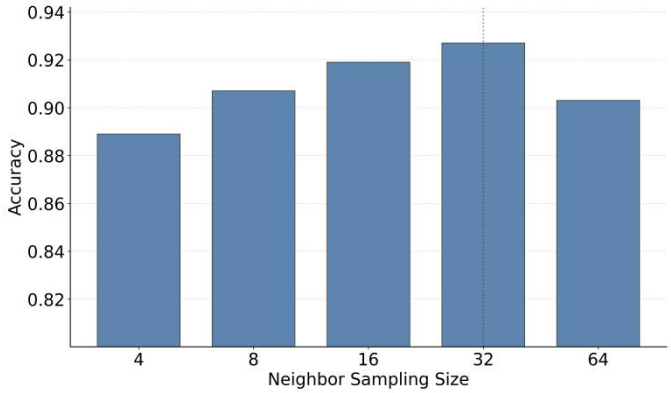

Figure 4. Sensitivity experiment of neighborhood sampling number to accuracy

At a moderate sampling scale, the model achieves a favorable balance between local granularity and global coverage. The sampled neighbors include both core transaction counterparts and a subset of key indirect relations. During propagation, the model can gradually integrate multi-level structural cues. This setting is particularly beneficial for accounting fraud detection. Chain transactions and related party dealings are better captured. Fraud often relies on a limited number of highly correlated entities rather than being evenly distributed across the entire network.

When the sampling size continues to increase, performance begins to decline. This reflects the introduction of many weakly related or irrelevant nodes through excessive neighborhood expansion. In corporate transaction networks, distant neighbors often correspond to normal transactions or noisy relations. Aggregating such information dilutes discriminative risk features. The model then struggles to maintain focus on key structural patterns. Large neighborhoods may also amplify the effects of degree imbalance. Highly connected nodes can dominate information propagation. Sensitivity to anomalous small structures is reduced.

Overall, this sensitivity analysis shows that a larger sampling size is not always better. The sampling strategy should match the structural characteristics of the transaction network. A moderate sampling scale preserves structural coverage while suppressing noise diffusion. It enables more effective capture of local anomalies and latent risk chains related to accounting fraud. This further indicates that careful control of information propagation range is a key factor in improving the stability and practical effectiveness of fraud detection in complex corporate transaction networks.

The proportion of training data determines the range of transaction patterns and risk structures that the model can cover, thus affecting its sufficiency in learning hidden fraud chains. As the number of training samples gradually increases, the model not only sees more diverse forms of normal transactions but also learns more easily the collaborative organization of abnormal behaviors in the network. To test the robustness of our method to changes in data scale, we examine the trend of recall as a function of different training proportions. The experimental results are shown in Figure 5.

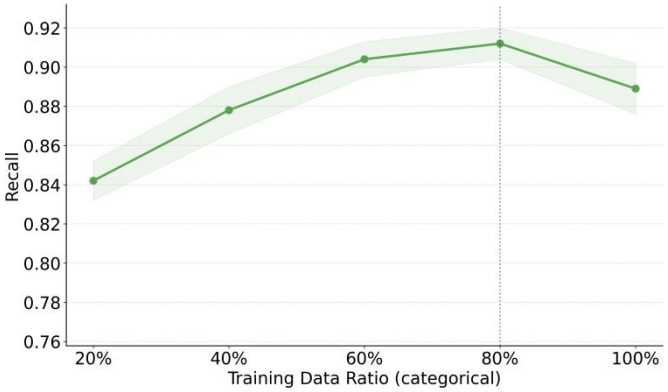

Figure 5. Sensitivity experiment of training data ratio to recall

The overall curve patterns show that increasing the proportion of training data significantly improves the model's coverage of fraudulent samples. This observation is consistent with the diversity of risk patterns in complex corporate transaction networks. Fraud often appears through multi-entity coordination, chained fund transfers, or combinations of related party transactions. Its structural forms are not fixed. Only when training data covers richer transaction topologies and behavioral variants can the model learn typical manifestations of these concealed patterns in the network. This leads to a stronger ability to discover anomalous entities.

When the training ratio increases from a low level to a moderate level, recall improves more noticeably. This indicates that newly added data mainly contributes to completing key structural patterns. For accounting fraud detection, a higher recall means lower missed detection risk. Missed cases often correspond to the most harmful concealed fraud chains. Improvements at the moderate stage usually imply that the model begins to capture cross-node combination features. Examples include abnormal fund cycles, asymmetric transaction fan out and fan in, and abnormal interactions related to network positions. The model can then identify potential fraudulent nodes under broader structural contexts.

As the training ratio continues to grow to a higher level, the curve gradually saturates. This suggests that learning core risk structures is nearly sufficient. Additional data mainly refines decision boundaries rather than introducing new patterns. The long tail nature of corporate transaction networks means that many new samples correspond to routine transaction behaviors. These samples help improve the modeling of normal patterns but provide diminishing marginal gains for recall. This phenomenon indicates that the proposed method absorbs key risk patterns efficiently. Once sufficient structural coverage is achieved, the model can stably extract effective anomaly cues from the graph context.

It is worth noting that a decline appears when the training ratio reaches its maximum. This suggests that a larger data scale does not always yield monotonic benefits. The effect may be influenced by changes in sample distribution and an

increase in difficult cases. Further data expansion may introduce more structurally similar but semantically normal entities. It may also add noisy edges and weak relations. These factors increase discrimination difficulty and weaken focus on anomalous signals. If added data intensifies the imbalance or introduces more complex confounding structures, the model may sacrifice sensitivity to minority fraud patterns to fit the overall data. Overall, this sensitivity result shows that training data coverage is critical for fraud detection in complex transaction networks. A balance between data expansion and structural noise control is still required to maintain stable recall of concealed fraud chains.

5. Conclusion

This study addresses accounting fraud detection in complex corporate transaction networks and proposes a unified analysis framework centered on graph-based modeling. From a relational perspective, it systematically characterizes transaction interactions among corporate entities and potential risk propagation mechanisms. Firms and their transactions are explicitly mapped into a structured network. Representations that integrate node attributes with network context are then learned. This approach goes beyond traditional paradigms based on isolated samples or static financial indicators. Fraud detection is no longer limited to single-point anomalies. Instead, abnormal behavior is understood through its organizational structure and collaborative patterns within the network. This shift provides a more interpretable and systematic solution for detecting concealed and structurally complex fraud.

From the perspective of practical value, this work offers direct implications for intelligent auditing, risk supervision, and corporate governance. In real business environments, regulatory targets are large in scale and highly interconnected. Rule-based inspection or manual sampling alone cannot effectively cover potential high-risk areas. The structured modeling approach adopted in this study helps prioritize entities or subnetworks with abnormal structural characteristics within massive transaction data. It provides data-driven support for audit resource allocation, risk warning prioritization, and continuous supervision. A network-oriented risk identification strategy better fits modern corporate structures with group operations, frequent related party transactions, and intensive cross-entity fund flows.

At the methodological level, this study further validates the applicability and extension potential of graph neural networks in accounting and financial analysis. Through multi-level information propagation in transaction networks, the model captures both local transaction behaviors and broader structural dependencies. It offers a new modeling tool for understanding anomaly mechanisms in complex economic systems. This structure-aware capability is not limited to accounting fraud detection. It can be naturally extended to other relation-intensive scenarios, such as financial risk transmission analysis, supply chain anomaly monitoring, and multi-entity collaborative behavior identification. A unified technical path for cross-domain risk modeling is thus provided.

Looking ahead, continued digitalization and expanding data sources will make transaction networks more dynamic and heterogeneous. Incorporating temporal evolution, relational heterogeneity, and uncertainty into graph modeling remains a key direction for improving real-world applicability. At the same time, interpretability, stability, and deployability will become increasingly important for regulatory and audit practice. By combining structured risk identification with business rules and audit logic, it is possible to build more reliable and actionable intelligent accounting risk analysis systems. Such systems can provide long-term support for maintaining market order and enhancing financial transparency.

References

[1] Tian Y, Liu G. Transaction fraud detection via spatial-temporal-aware graph transformer[J]. arXiv preprint arXiv:2307.05121, 2023.

[2] West J, Bhattacharya M. Intelligent financial fraud detection: a comprehensive review[J]. Computers & security, 2016, 57: 47-66.

[3] Motie S, Raahemi B. Financial fraud detection using graph neural networks: A systematic review[J]. Expert Systems with Applications, 2024, 240: 122156.

[4] Tong G, Shen J. Financial transaction fraud detector based on imbalance learning and graph neural network[J]. Applied Soft Computing, 2023, 149: 110984.

[5] Kadam P. Financial fraud detection using jump-attentive graph neural networks[C]//2024 International Conference on Machine Learning and Applications (ICMLA). IEEE, 2024: 628-635.

[6] Cheng D, Zou Y, Xiang S, et al. Graph neural networks for financial fraud detection: a review[J]. Frontiers of Computer Science, 2025, 19(9): 199609.

[7] Cheng D, Wang X, Zhang Y, et al. Graph neural network for fraud detection via spatial-temporal attention[J]. IEEE Transactions on Knowledge and Data Engineering, 2020, 34(8): 3800-3813.

[8] Innan N, Marchisio A, Bennai M, et al. Qfnn-ffd: Quantum federated neural network for financial fraud detection[C]//2025 IEEE International Conference on Quantum Software (QSW). IEEE, 2025: 41-47.

[9] Khodabandehlou S, Golpayegani A H. FiFrauD: unsupervised financial fraud detection in dynamic graph streams[J]. ACM Transactions on Knowledge Discovery from Data, 2024, 18(5): 1-29.

[10] Cui Y, Han X, Chen J, et al. FraudGNN-RL: a graph neural network with reinforcement learning for adaptive financial fraud detection[J]. IEEE Open Journal of the Computer Society, 2025.

[11] Lin J, Guo X, Zhu Y, et al. FraudGT: a simple, effective, and efficient graph transformer for financial fraud detection[C]//Proceedings of the 5th ACM International Conference on AI in Finance. 2024: 292-300.

[12] Cheng Z, Gui G, Tong K, et al. Finstack-net: Hierarchical feature crossing and stacked ensemble learning for financial fraud detection[C]//Proceedings of the 2025 2nd International Conference on Digital Economy, Blockchain and Artificial Intelligence. 2025: 444-447.