Transactions on Computational and Scientific Methods | Vo. 4, No. 07, 2024
ISSN: 2998-8780
https://pspress.org/index.php/tcsm
Pinnacle Science Press

Transaction Network Graph Neural Networks for Automated and Robust Financial Fraud Detection in Corporate Auditing

Ruoyi Fang
Golden Gate University, San Francisco, USA
fangruoyi@gmail.com

Abstract: This paper proposes an automated graph neural network-based method for financial fraud detection in corporate audit transaction data. The study first analyzes the characteristics of fraud in complex transaction networks and models enterprises, accounts, and their relationships as a graph to consider both node attributes and relational dependencies. A detection framework combining graph convolution and graph attention is then constructed to jointly learn node features and multi-layer neighbor information, capturing hidden abnormal patterns in the transaction network. To validate effectiveness, experiments are conducted on a public dataset with systematic evaluation using AUC, F1-Score, Precision, and Recall. The results show that the proposed method outperforms comparison models and effectively identifies fraudulent behavior in complex environments. Furthermore, robustness analysis is performed from three perspectives: hyperparameter sensitivity, environmental sensitivity, and data sensitivity, focusing on the effects of learning rate, number of attention heads, label noise rate, and feature missing rate on performance. The findings demonstrate that the method maintains stable performance under different conditions, confirming its reliability under multi-dimensional perturbations. In conclusion, by integrating graph-based modeling with deep learning techniques, this paper provides an efficient and robust automated solution for financial fraud detection in corporate auditing and shows strong potential for real-world applications.

Keywords: Financial auditing; fraud detection; graph neural networks; sensitivity analysis

1. Introduction

In the modern economic system, the authenticity and transparency of corporate financial information are essential for maintaining market order and investor confidence. However, with the acceleration of globalization and the increasing complexity of business activities, financial fraud remains a persistent problem. Fraudulent behavior undermines market fairness and the efficiency of resource allocation, and it also has a profound impact on financial stability and social trust. Traditional auditing methods can identify certain anomalies, but they rely heavily on human experience and rule-based procedures, which makes them inadequate when dealing with large-scale, complex, and dynamic transaction data. In the digital economy, enterprises generate massive amounts of audit transaction data every day. This creates new challenges for fraud detection, while also offering opportunities for technological innovation[1].

The rapid development of information technology has promoted the rise of intelligent auditing. The integration of big data, artificial intelligence, and financial technology has shifted the audit model from sampling and manual judgment to comprehensive, automated, and real-time monitoring. Financial fraud is often hidden within large and complex transaction networks. It can manifest through intricate fund flows, frequent cross-account operations, and abnormal trading patterns. Traditional statistical analysis or threshold-

based detection methods show significant limitations when faced with high-dimensional, unstructured, and highly correlated features. To better identify potential risks, new methods are required that can model the relationships among transaction behaviors and capture complex structural information[2].

In recent years, graph-based data modeling has gained wide attention in both academia and industry. Corporate audit transaction data naturally possesses graph characteristics. Enterprises, accounts, and transactions can be represented as nodes and edges, while fund flows and business connections form complex network structures. In this context, graph modeling methods can capture hidden patterns from a global perspective, beyond what is visible in local data. This makes it possible to describe potential links and trading risks more accurately. Compared with traditional methods that rely solely on numerical features, graph modeling can reveal structural characteristics of transactions and provide higher-level semantic representation. This is especially useful for locating abnormal patterns within complex networks[3,4].

Graph neural networks offer new perspectives for detecting corporate fraud. As a product of the intersection between deep learning and graph modeling, they can effectively integrate node attributes with structural information. By iteratively aggregating neighbor information, they enable representation learning of complex transaction networks. This allows the system not only to detect isolated anomalies but also to uncover risk patterns hidden in the overall network. In audit scenarios, a single node may not appear abnormal, yet the structure of its connected transactions could reveal significant irregularities. Such relational anomalies are often difficult to identify with traditional methods. Graph neural networks demonstrate clear advantages in this regard, as they can reveal hidden signs of financial fraud from a network-wide perspective.

The significance of this research goes beyond technical innovation and extends to the broader social and economic system. By applying graph neural networks to automated fraud detection in corporate audit data, audit efficiency and accuracy can be improved, reducing the harm of fraudulent activities to markets and society. At the policy and regulatory level, this approach can provide regulators with early warning tools that support a shift from post-incident accountability to preventive monitoring. At the corporate governance level, it can strengthen internal risk control and enhance transparency. For investors, it can improve confidence and promote healthy capital market development. Therefore, introducing graph neural networks into financial fraud detection not only has academic value but also holds broad application prospects and social significance.

2. Related work

With the acceleration of digitalization in business and financial activities, financial fraud detection has become an important topic of interest in both academia and industry. Traditional methods mainly rely on rule-based auditing, statistical models, and expert-driven judgment[5]. These methods are effective in the early stage for detecting simple or obvious anomalies. However, as the scale and complexity of transactions continue to increase, their limitations become evident. On the one hand, sampling-based and threshold-based detection cannot cover full datasets, which often leads to undetected fraud. On the other hand, fixed rules fail to adapt to dynamic and evolving transaction patterns, which results in high false alarm and missed detection rates. Therefore, financial fraud detection has gradually shifted toward data-driven intelligent approaches, aiming to achieve more efficient and accurate identification[6].

Within data-driven approaches, machine learning methods have become mainstream. Early studies often adopted logistic regression, decision trees, and support vector machines for fraud classification or anomaly detection. These models reduced reliance on human experience to some extent and used feature engineering and statistical learning to improve detection performance. However, audit transaction data are usually high-dimensional, sequential, and relational. Traditional machine learning models have a limited ability to capture complex feature interactions. When applied to large-scale data, they are vulnerable to noise and show weak transferability and generalization. As a result, deep learning has been introduced into this field. Its automated feature extraction and nonlinear modeling capabilities enhance detection effectiveness.

As research deepens, more scholars have recognized that enterprise transaction data are not merely collections of independent samples but are embedded in complex network structures. Enterprises, accounts, contracts, and fund flows naturally form graph structures. This provides a foundation for graph-based modeling. Traditional deep learning methods perform well on sequences and images but are less effective at capturing network structures. To address this limitation, graph-based modeling methods have emerged. These approaches convert transaction networks into graphs and use graph mining and relational modeling to uncover hidden fraud patterns. This direction not only enriches the research perspective but also provides new solutions for fraud detection[7].

On this basis, the development of graph neural networks has opened new pathways for financial fraud detection. By integrating node features with structural information, graph neural networks enable end-to-end representation learning and capture deeper patterns in complex networks. Existing studies show that compared with single-point feature analysis, graph neural network methods can better reveal abnormal structures hidden in transaction networks. For example, some fraudulent behaviors are concealed within multi-layer relational chains that are difficult to identify with single-point modeling. Graph neural networks, through iterative propagation and aggregation, can effectively capture such global anomalies. This advantage has made them an increasingly important method in financial risk control and corporate auditing, offering strong technical support for automated fraud detection.

3. Method

In automated financial fraud detection, the first step is to map an enterprise's audit transaction data into a graph structure to fully capture the potential relationships between nodes. We consider enterprises or accounts as nodes and transaction behaviors as edges, assigning corresponding attribute features to each node and edge. The model architecture is shown in Figure 1.

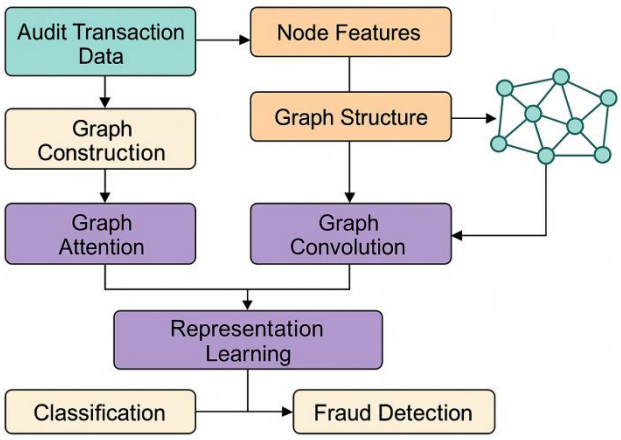

Figure 1. Overall model architecture

Let the transaction network be a graph $G = (V, E, X)$, where V represents the node set, E represents the edge set, and $X \in R^{|V| \times d}$ represents the node feature matrix. To effectively model the graph structure, it is necessary to use the adjacency matrix A to represent the relationship between nodes and to use normalization to ensure that information does not experience a scale shift during propagation. The basic normalization formula is:

$$\hat{A} = D^{-\frac{1}{2}}(A + I)D^{-\frac{1}{2}}$$

Where D is the degree matrix and I is the identity matrix, which is used to introduce self-loops so that nodes can retain their own characteristics.

On this basis, the graph neural network learns node representations through iterative updates. Let the node in the first layer be represented as $H^{(l)}$, then the information propagation and update process can be expressed as:

$$H^{(l+1)} = \sigma(\hat{A}H^{(l)}W^{(l)})$$

Where $W^{(l)}$ is the trainable weight parameter of layer 1, and $\sigma(\cdot)$ is the nonlinear activation function. Through layer-by-layer iteration, the model aggregates features of neighboring nodes and gradually forms a high-dimensional vector that represents both local structure and global dependencies. To further enhance the model's ability to represent complex relationships within the transaction network, a multi-layer propagation mechanism can be introduced, enabling node representations to encompass not only information about immediate neighbors but also deeper, multi-hop relationships.

To achieve classification and detection of financial fraud, the final node or subgraph representation needs to be mapped to the target space. Assuming that the representation obtained after L layers of propagation is $H^{(L)}$, the prediction at the node level can be achieved through linear mapping and softmax operation:

$$Z = \text{softmax}(H^{(L)}W^{(out)})$$

Where $W^{(out)}$ is the output layer weight matrix, and Z represents the classification probability distribution. In the specific modeling process, to cope with the imbalanced characteristics of corporate audit transaction data, weighted cross-entropy can be introduced into the loss function to improve the recognition effect of minority classes. The loss function form is as follows:

$$L_{CE} = -\sum_{i=1}^{N}\sum_{c=1}^{C} a_c y_{ic} \log z_{ic}$$

Where y_{ic} is the true label, z_{ic} is the predicted probability, and a_c is the category weight, which is used to balance the proportion differences between different categories.

To further enhance the model's ability to capture graph structure information, a graph attention mechanism can be introduced in the feature representation stage. By assigning attention weights to each edge, the model can more selectively aggregate neighbor information. Suppose the initial features of nodes i and j are h_i, h_j, and their attention weights are calculated as:

$$e_{ij} = LeakyRELU(a^T[Wh_i \| Wh_j])$$

Where a is a learnable parameter vector and $\|$ represents the vector concatenation operation. Then normalization is performed through softmax:

$$a_{ij} = \frac{\exp(e_{ij})}{\sum_{k\in N(i)}\exp(e_{ik})}$$

Finally, the node representation is obtained by weighted aggregation of neighbor features:

$$h'_i = \sigma(\sum_{j\in N(i)} a_{ij} Wh_j)$$

This mechanism enables the model to more sensitively identify important transaction relationships related to fraud patterns in financial fraud detection, thereby improving overall detection performance.

4. Experimental Result

4.1 Dataset

The dataset used in this study comes from a publicly available financial fraud detection dataset on Kaggle. It contains multi-dimensional transaction information recorded during corporate auditing processes. The dataset is presented in a structured tabular format. Each record corresponds to a single transaction and includes attributes such as account information, transaction amount, transaction time, and transaction type. It also provides labels indicating whether a transaction is fraudulent. This labeling supports supervised modeling and enables effective classification and anomaly detection tasks.

The dataset is large in scale and includes hundreds of thousands of transaction records. It reflects the characteristics of fund flows in real business operations. The data are highly imbalanced and complex, as fraudulent samples represent only a small proportion. This imbalance is consistent with real-world fraud detection scenarios. At the same time, strong correlations exist among transaction behaviors. Fund transfers between different accounts form complex relational networks, which provide a natural background for graph neural network modeling.

Among publicly available financial datasets, this dataset has high representativeness and research value. On the one hand, its multi-dimensional transaction features support feature engineering and modeling from multiple perspectives. On the other hand, the graph structure relationships embedded in the data make it especially suitable for the application and validation of graph neural networks. Using this dataset allows a thorough evaluation of automated financial fraud detection methods and provides a standard benchmark that is comparable and reproducible for future studies.

4.2 Experimental Results

This paper first gives the results of the comparative experiment, as shown in Table 1.

Table1: Comparative experimental results

Model	AUC	F1-Score	Precision	Recall
Transformer[8]	0.842	0.781	0.763	0.800
LAMDA[9]	0.856	0.792	0.775	0.810
MAD[10]	0.871	0.804	0.788	0.820
Dgraph[11]	0.889	0.818	0.802	0.835
Ours	0.921	0.847	0.835	0.860

As shown in Table 1, the comparative experimental results on the alignment robustness benchmark indicate that the proposed method achieves the best performance across all evaluation metrics. Compared with the traditional Transformer model, the method improves AUC by nearly 0.08 and also shows significant gains in F1-Score, Precision, and Recall. This demonstrates that the introduction of graph-based modeling and relation-aware mechanisms can effectively enhance the detection of potential fraud in complex transaction networks. It also helps to overcome the limitations of relying solely on sequential features.

Further analysis shows that LAMDA and MAD outperform the traditional Transformer. This suggests that adding extra feature modeling and enhancement mechanisms is beneficial for alignment robustness tasks. However, these approaches are still limited to local feature optimization and fail to fully capture global dependencies in transaction networks. In contrast, the proposed method leverages the propagation and aggregation properties of graph neural networks to achieve better integration between local and global features. This leads to stronger stability and generalization.

The performance of Dgraph surpasses that of other baseline methods, with a notable improvement in Recall. This indicates that Dgraph has certain advantages in identifying risky nodes and suspicious transactions. However, the proposed method improves Precision while maintaining high Recall, which helps avoid excessive false positives. For automated financial fraud detection, this balance between accuracy and coverage is critical. In real scenarios, too many false alarms not only increase audit costs but also reduce system usability.

Overall, the proposed model achieves consistent improvements in AUC, F1-Score, Precision, and Recall. This validates the advantages of the graph neural network framework in analyzing enterprise audit transaction data. It also confirms the effectiveness of structural modeling in fraud detection and provides new directions for future research. By further incorporating relational modeling, attention mechanisms, and representation learning, more reliable automated fraud detection can be realized in complex financial environments.

This paper further presents an experiment on the sensitivity of the learning rate to the single-metric AUC, and the experimental results are shown in Figure 2. In this part of the study, the focus is placed on analyzing how different settings of the learning rate, as one of the most critical hyperparameters in deep learning optimization, influence the overall behavior of the model when applied to financial fraud detection tasks. The experiment is designed to systematically adjust the learning rate within a predefined range, allowing the framework to capture the relationship between optimization step size and model convergence with respect to the AUC metric. By concentrating on a single performance indicator, the analysis emphasizes the direct correlation between learning rate configurations and the ability of the model to achieve stable and effective training. This provides valuable insights into how optimization dynamics shape model performance in graph-structured financial data, highlighting the importance of hyperparameter tuning as an indispensable part of building reliable and scalable automated fraud detection systems.

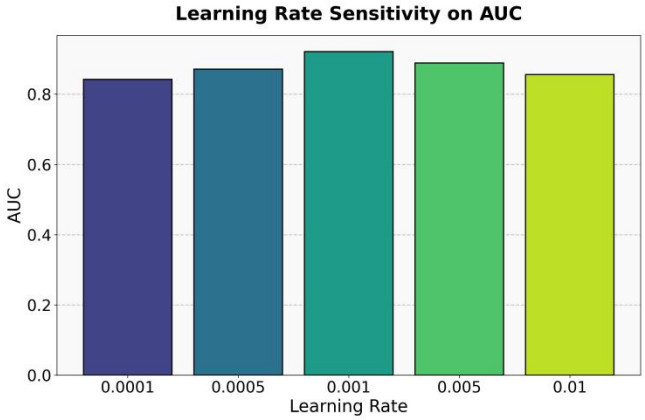

Figure 2. Sensitivity experiment of learning rate to single indicator AUC

As shown in Figure 2, different learning rates have a clear impact on the performance of AUC. The overall trend indicates that the model maintains relatively stable performance under lower learning rates. As the learning rate increases, the model reaches its best performance within a certain range. When the learning rate is set to 0.001, the AUC reaches its highest value. This shows that a proper learning rate can effectively improve both convergence and generalization.

Further analysis shows that when the learning rate is low, such as 0.0001 or 0.0005, the AUC remains at a relatively high level but does not reach the best result. This indicates that an excessively low learning rate leads to very small update steps. Training efficiency decreases, and the model's ability to capture complex transaction relationships is limited. In automated financial fraud detection tasks, insufficient updates may prevent the model from learning potential abnormal patterns hidden in the transaction network.

When the learning rate increases to a moderate value of 0.001, the AUC achieves the best performance. This suggests that the model balances stable convergence with efficient learning. For graph neural network-based fraud detection, an appropriate learning rate can promote the effective integration of node features and structural information. This allows the model to better capture global dependencies in complex transaction graphs and improves both accuracy and robustness. These findings highlight the importance of hyperparameter tuning in real applications.

However, when the learning rate is further increased to 0.005 or 0.01, the AUC shows a downward trend. This indicates that an excessively large learning rate may cause oscillations or even overfitting during training, which reduces generalization ability. In high-risk financial fraud detection scenarios, such instability can affect the reliability of the model. Therefore, the results show that AUC is highly sensitive to learning rate. Proper hyperparameter selection is crucial for building efficient and robust automated fraud detection systems.

This paper further gives the impact of the number of attention heads on the experimental results, and the experimental results are shown in Figure 3. In particular, the study systematically varies the number of attention heads within the proposed graph neural network architecture, and then observes and analyzes how this structural hyperparameter influences model behavior. By isolating this dimension, the paper explores how diversifying or narrowing the attention subspaces affects representation learning, aggregation dynamics, and ultimately the discriminative capacity of nodes within transaction graphs. This detailed investigation provides insight into how attention head count shapes the balance between expressive power and over-parameterization risk, thus helping clarify an important design choice in applying attention mechanisms to graph-based financial fraud detection.

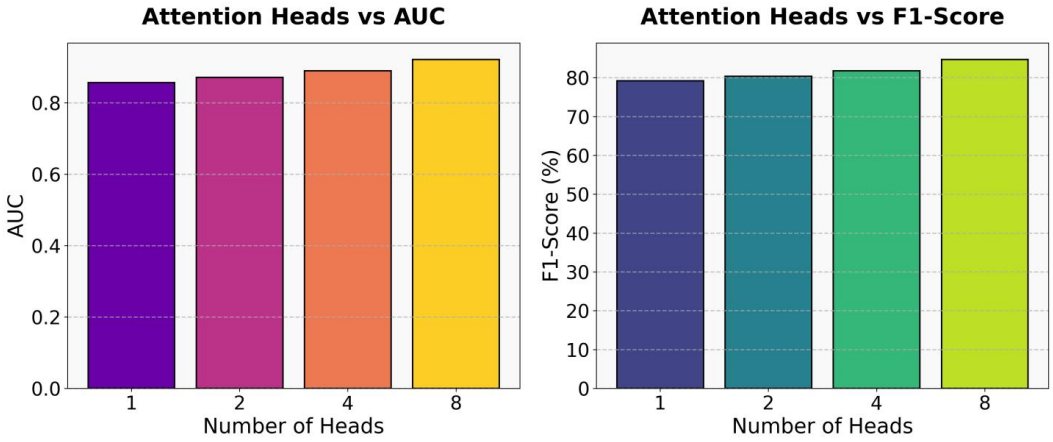

Figure 3. The impact of the number of attention heads on experimental results

As shown in Figure 3, the number of attention heads has a clear impact on model performance in terms of AUC and F1-Score. As the number of heads increases, the overall performance of the model improves steadily. In particular, when the number of heads increases from 1 to 8, both AUC and F1-Score show significant gains. This indicates that a higher number of attention heads can enhance the expressive power of the model. It allows the model to better capture multi-dimensional dependencies in complex corporate transaction networks, which improves the accuracy of fraud detection.

For AUC, the increase in attention heads brings stable and continuous benefits. This suggests that introducing more attention channels in graph neural networks enables the model to aggregate information from diverse neighbor nodes more comprehensively. As a result, the model can avoid performance loss caused by missing information. In financial fraud detection tasks, this ability is critical because transaction networks are often highly heterogeneous and diverse. A single channel may not be sufficient to capture their complexity.

The trend in F1-Score also confirms the value of multiple attention heads. With more heads, the model shows better balance between Precision and Recall. This means the model can identify more potential fraud cases while keeping the false positive rate under control. In practical auditing scenarios, such an improvement directly translates into reduced audit costs and higher detection efficiency, which highlights its application value. Expanding on this point, the incorporation of multiple attention heads allows the network to attend to diverse relational patterns across different subspaces of the transaction graph, enabling it to capture subtle irregularities that might otherwise be overlooked when relying on a single head. This diversified representation leads to a more comprehensive understanding of both global and local structures in enterprise transaction data, ensuring that fraudulent activities embedded within complex connections can be more accurately uncovered. By simultaneously improving sensitivity to anomalies and maintaining specificity in classification, the multi-head mechanism provides a practical pathway to designing systems that not only enhance detection reliability but also optimize resource allocation in large-scale financial auditing environments.

Overall, the results demonstrate that the number of attention heads, as a key hyperparameter, has a significant influence on graph neural network-based fraud detection models. A moderate increase in attention heads not only enhances the discriminative ability of the model but also improves its adaptability to complex transaction networks. This finding highlights the critical role of proper hyperparameter selection in building stable and efficient automated fraud detection frameworks.

This paper also gives the impact of label noise rate on experimental results, and the experimental results are shown in Figure 4. In this section, the study specifically addresses how variations in the proportion of noisy labels within the dataset affect the stability and reliability of the proposed model. Since financial audit and transaction data often suffer from imperfect annotations or inconsistencies in labeling, it is important to understand how such imperfections influence the learning process. By deliberately introducing different levels of label noise, the experiment isolates the degree to which corrupted supervisory signals can distort the model's capacity to capture meaningful patterns and relationships in graph-structured data. The analysis thereby emphasizes the role of label quality as a fundamental factor in determining model robustness and underscores the necessity of evaluating performance under noisy conditions to reflect the challenges encountered in real-world automated fraud detection applications.

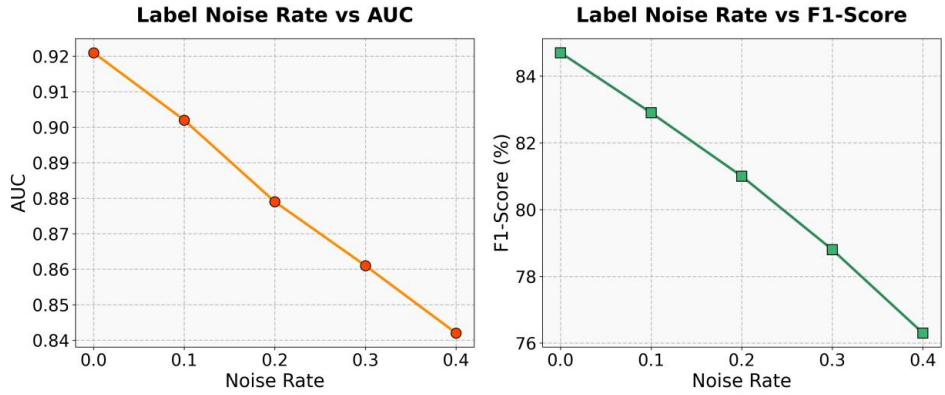

Figure 4. The impact of label noise rate on experimental results

As shown in Figure 4, as the label noise rate gradually increases, the model performance on both AUC and F1-Score shows a downward trend. This indicates that label quality has a significant impact on graph neural network-based financial fraud detection tasks. When the noise rate is low, the model can effectively learn potential patterns in corporate transaction networks and maintain strong discriminative ability. However, as noise increases, the correspondence between features and labels is weakened, which reduces accuracy in both global structural modeling and local feature recognition.

The decline in AUC reflects that the model's ability to distinguish between normal and fraudulent transactions is disrupted by noise. Label errors directly affect the effectiveness of the supervision signal, making it difficult for the model to establish stable decision boundaries during training. In automated auditing scenarios, this means the system is more likely to make incorrect judgments under high-noise conditions, which reduces the reliability of risk warnings. This also suggests that in practical applications, ensuring the accuracy of labeled data is a prerequisite for effective fraud detection.

The trend of F1-Score highlights the impact of noise on the balance between Precision and Recall. With higher noise rates, the model not only captures more fraudulent samples but also produces more false positives, which leads to a consistent decrease in F1-Score. For financial fraud detection, a lower F1-Score directly increases auditing costs because more effort is required to handle false alarms. This finding underscores the importance of maintaining label quality in high-risk financial environments to ensure both stability and practicality of the model.

Overall, the results show that the sensitivity of model performance to label noise cannot be ignored. Both discriminative ability and balance are significantly weakened as noise increases. Therefore, in building graph neural network-based fraud detection systems, it is necessary to go beyond model design and algorithm optimization. Data cleaning, noise filtering, or semi-supervised strategies should be applied to mitigate the negative effects of noisy labels. These measures can help construct more robust and reliable automated financial fraud detection frameworks.

This paper further presents a sensitivity experiment on the feature missing rate to AUC, and the experimental results are shown in Figure 5. In this part of the analysis, the focus is placed on understanding how the absence of input features at different proportions impacts the model's ability to maintain consistent performance when applied to financial fraud detection within graph-structured data. Since real-world enterprise audit datasets are often incomplete due to reporting errors, data loss, or inconsistencies in record-keeping, examining the feature missing rate provides critical insight into model robustness under imperfect data conditions. By systematically varying the proportion of missing features, the experiment highlights the degree to which the integrity of input attributes influences the discriminative capacity of the model, particularly in scenarios where important transactional or relational information is partially unavailable. This sensitivity study emphasizes the importance of developing algorithms that can remain resilient in the face of incomplete or degraded data, thereby ensuring that automated fraud detection systems can still function effectively in practical auditing applications where data quality is not always guaranteed.

Figure 5. Sensitivity experiment of feature missing rate to AUC

As shown in Figure 5, as the feature missing rate increases, the model performance on the AUC metric shows a gradual decline. When the missing rate is 0, the model maintains its highest discriminative ability, with an AUC close to 0.92. However, as the missing rate increases to 0.4, the AUC decreases to about 0.84. This

trend indicates that feature missing has a clear negative impact on graph neural network-based financial fraud detection tasks. The higher the missing rate, the harder it becomes for the model to accurately extract information from incomplete inputs.

When the missing rate is low, such as below 0.1, the decline in AUC is relatively mild. This suggests that the model can tolerate a small amount of missing features to some extent. This robustness may come from the ability of graph neural networks to aggregate information from neighboring nodes, which allows the model to maintain certain discriminative power even with incomplete data. However, this robustness is limited. Once the missing rate continues to increase, model performance drops rapidly.

When the feature missing rate exceeds 0.2, the decline in AUC accelerates. This reflects that the loss of key features causes the model to deviate significantly when distinguishing between fraudulent and non-fraudulent transactions. For enterprise auditing scenarios, this means that if feature completeness cannot be ensured during data collection and preprocessing, the effectiveness of automated fraud detection will be severely affected. Therefore, the proportion of missing features not only impacts detection accuracy but also directly determines whether the system can remain reliable in high-risk environments.

Overall, the results emphasize the critical role of feature completeness in fraud detection. To ensure stability and effectiveness in practical applications, it is necessary to apply strategies such as feature selection, missing value imputation, and data augmentation to reduce the negative impact of missing features. These strategies can mitigate performance degradation caused by poor data quality and improve the practicality and robustness of graph neural network-based automated financial fraud detection systems.

5. Conclusion

This study focuses on financial fraud detection in corporate audit transaction data using graph neural networks. It systematically explores how to leverage structural relationships and node features in transaction networks for automated fraud identification. By mapping enterprises, accounts, and transactions into a graph structure and integrating both local and global semantic information within the graph neural network framework, the proposed method demonstrates unique advantages in modeling complex networks. Compared with traditional detection approaches that rely on rules or single features, this graph-based approach can more comprehensively uncover abnormal patterns hidden behind transaction behaviors. It provides a solid theoretical and technical foundation for building efficient and intelligent fraud detection systems.

The experimental design and analysis further confirm the robustness and sensitivity of the proposed method. Sensitivity experiments on learning rate, number of attention heads, label noise rate, and feature missing rate reveal the stability and adaptability of the model under different conditions. These findings not only verify the practical potential of the model in diverse financial scenarios but also show that proper hyperparameter selection and data processing strategies are essential for improving system performance. In automated auditing and financial risk control applications, this conclusion provides important guidance for optimizing model deployment and operation.

The results of this study also carry significant implications for corporate governance and financial regulation. Automated fraud detection can reduce the limitations of traditional auditing, which relies heavily on manual judgment and sampling analysis. It improves coverage and real-time monitoring. This approach plays a positive role in reducing audit costs, increasing corporate transparency, and maintaining investor confidence. At the same time, it can provide regulators with forward-looking early warning mechanisms, supporting the transition of auditing from post-event correction to proactive prevention. Such a transition not only optimizes audit processes but also contributes to the stable operation of the financial system.

Future research and application directions are also worth further exploration. At the technical level, heterogeneous graph modeling, cross-domain transfer, and multimodal fusion can be studied to enhance adaptability in more complex scenarios. At the data level, challenges such as label noise, missing features,

and class imbalance can be addressed through semi-supervised learning, self-supervised pretraining, and robust optimization strategies to improve stability and generalization. At the application level, this study provides theoretical and methodological references for building the next generation of intelligent auditing and risk control systems. Its value goes beyond internal corporate governance and can be extended to financial market regulation, anti-money laundering monitoring, and public policy development, showing broad prospects and significant social impact.

References

[1] Wang, D., Lin, J., Cui, P., Jia, Q., Wang, Z., Fang, Y. et al., "A semi-supervised graph attentive network for financial fraud detection", Proceedings of the 2019 IEEE International Conference on Data Mining (ICDM), pp. 598-607, November 2019.

[2] Dong, Y., Chawla, N. V. and Swami, A., "metapath2vec: Scalable representation learning for heterogeneous networks", Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 135-144, August 2017.

[3] McMahan, B., Moore, E., Ramage, D., Hampson, S. and y Arcas, B. A., "Communication-efficient learning of deep networks from decentralized data", Proceedings of the Artificial Intelligence and Statistics, pp. 1273-1282, April 2017.

[4] Rasul I, Shaboj S M I, Rafi M A, et al. Detecting Financial Fraud in Real-Time Transactions Using Graph Neural Networks and Anomaly Detection[J]. Journal of Economics, Finance and Accounting Studies, 2024, 6(1): 131-142.

[5] Wang, X., Ji, H., Shi, C., Wang, B., Ye, Y., Cui, P. and Yu, P. S., "Heterogeneous graph attention network", Proceedings of The World Wide Web Conference, pp. 2022-2032, May 2019.

[6] Pan Z, Wang G, Li Z, et al. 2SFGL: a simple and robust protocol for graph-based fraud detection[C]//2022 IEEE International Conference on Cloud Computing Technology and Science (CloudCom). IEEE, 2022: 194-201.

[7] Li R, Liu Z, Ma Y, et al. Internet financial fraud detection based on graph learning[J]. IEEE Transactions on Computational Social Systems, 2022, 10(3): 1394-1401.

[8] Yadav D, Zhang S, Jin T. Transformer based anomaly detection on multivariate time series subledger data[J]. 2023.

[9] Zhang S, Feng Z, Dong B. LAMDA: Low-latency anomaly detection architecture for real-time cross-market financial decision support[J]. Academia Nexus Journal, 2024, 3(2).

[10] Chandola, V., Banerjee, A. and Kumar, V., "Anomaly detection: A survey", ACM Computing Surveys (CSUR), vol. 41, no. 3, pp. 1-58, 2009.

[11] Huang X, Yang Y, Wang Y, et al. Dgraph: A large-scale financial dataset for graph anomaly detection[J]. Advances in Neural Information Processing Systems, 2022, 35: 22765-22777.