Frontiers in Business, Economics and Management
ISSN: 2766-824X | Vol. 19, No. 3, 2025

FinGuard-GNN: Dynamic Graph Neural Network Framework for Financial Fraud Detection

Ruijie Huang

School of Science and Technology, Beijing Normal-Hong Kong Baptist University, Zhuhai, Guangdong, China

Abstract: With the significant increase in financial fraud incidents, financial fraud detection has become a critical research area. Complex financial relationship networks involving thousands or even millions of nodes present enormous challenges for fraud detection tasks. Although researchers have developed various graph-based methods to detect fraudulent behavior within these complex networks, existing approaches overlook two key issues in fraud graphs: the diversity of non-additive attributes and the distinguishability of grouped message passing from neighboring nodes. This paper proposes FinGuard-GNN (Financial Guardian Graph Neural Network), a novel dynamic graph neural network for financial fraud detection that addresses the aforementioned issues through innovative feature transformation strategies and a Cascaded Risk Diffusion (CRD) mechanism. For feature transformation, we implement Adaptive Tree Partitioning (ATP) encoding and Statistical Evidence Weighting (SEW) encoding to convert various types of non-additive node attributes into vector representations suitable for GNN aggregation operations, avoiding the generation of meaningless features while maintaining strong interpretability. For risk propagation, we design a feedback-based Cascaded Risk Diffusion strategy that enables dynamic accumulation and decay of risk information across the network. Additionally, we develop a Responsive Group Allocation (RGA) strategy that divides graph nodes into distinct groups followed by hierarchical aggregation, enhancing the distinguishability of fraudulent nodes. Experiments on two classic financial fraud datasets demonstrate that our proposed method achieves superior discriminative capability for fraudulent nodes compared to traditional graph algorithms and machine learning methods. The experimental results confirm the advantages of FinGuard-GNN in handling non-additive features in complex financial networks, improving node distinguishability, and capturing hierarchical risk propagation, providing a novel solution for the fi.

Keywords: Dynamic graph neural networks; Financial fraud; Machine learning.

1. Introduction

In recent years, the rapid evolution of financial technology has catalyzed unprecedented growth in digital payments, online lending, and cryptocurrency transactions, injecting new vitality into economic development [1, 2]. However, financial fraud has simultaneously grown in complexity and diversity. According to NASDAQ estimates, at least $3.1 trillion in illicit funds flowed through the global financial system in 2023, affecting even the most diligent organizations. Traditional fraud detection methods based on rules and simple machine learning algorithms have become increasingly inadequate to combat these sophisticated threats. Fraudsters continuously adapt their tactics, making static rules difficult to maintain, while fraudulent activities often involve intricate social networks and transaction relationships that cannot be effectively identified through individual feature analysis alone [3, 4, 5]. Consequently, developing advanced fraud detection algorithms that effectively leverage relationship information between entities has become a crucial focus for both academia and industry.

The inherent complexity and volatility of financial markets present significant challenges for fraud detection systems. Financial data is characterized by its heterogeneity, temporal dynamics, and complex relational structures, which are typically represented as graph data. This representation exposes two fundamental limitations in current approaches. First, traditional financial fraud detection methods struggle to process the multifaceted relationships and numerous components present in financial networks. Second, the graph structures constructed from financial data are often heterogeneous and time-varying, presenting substantial modeling challenges that conventional techniques cannot address effectively [6, 7].

Current graph-based fraud detection methods encounter two critical challenges. Challenge 1 (C1): Non-additive attribute processing. Existing methods often employ simple averaging or summation operations when handling non-additive node attributes (such as transaction frequency, amount distribution, and time intervals), which not only loses the statistical distribution characteristics of the original features but may also generate meaningless feature representations. For instance, averaging transaction amounts might mask fluctuation patterns indicative of anomalous activity, while simple summation can lead to scale imbalances [8]. Challenge 2 (C2): Distinguishability of grouped message passing. Traditional message passing mechanisms typically treat all neighboring nodes as homogeneous, lacking effective grouping strategies to differentiate between normal users, suspicious users, and confirmed fraudulent users [9]. This homogeneous treatment significantly impairs the model's ability to capture the varying influence patterns of different node types, particularly in identifying sophisticated fraud schemes.

To address these challenges, we propose FinGuard-GNN (Financial Guardian Graph Neural Network), a novel dynamic graph neural network framework for financial fraud detection. To address C1, we develop a comprehensive feature transformation scheme including Adaptive Tree Partitioning (ATP) encoding and Statistical Evidence Weighting (SEW) encoding, which transforms various non-additive node attributes into vector representations suitable for GNN aggregation while maintaining their statistical properties. To address C2, we introduce a Cascaded Risk

Diffusion (CRD) mechanism that incorporates dynamic weight calculation, path-length decay functions, and feedback regulation to accurately model risk diffusion patterns across financial networks. Additionally, we design a Responsive Group Allocation (RGA) strategy that adaptively divides graph nodes into distinct functional groups and enhances fraudulent node identifiability through hierarchical information aggregation. Our extensive experiments on two real-world financial fraud datasets demonstrate that FinGuard-GNN significantly outperforms existing graph algorithms and traditional machine learning methods, validating its effectiveness and practical value in financial fraud detection scenarios.

2. Related Work

2.1. Traditional Fraud Detection Approaches

Traditional fraud detection methods rely primarily on rule-based systems and conventional machine learning algorithms [10]. Rule-based approaches employ expert-defined heuristics to flag suspicious activities, offering high interpretability but lacking adaptability to evolving fraud patterns [11]. Classical machine learning techniques such as random forests, SVMs, and logistic regression improve upon these limitations by learning patterns from historical data. However, these methods operate on tabular data with independent features and struggle to capture the complex relational patterns inherent in financial fraud [12]. They typically treat each transaction or user in isolation, missing crucial contextual information embedded in transaction networks. While these approaches provide a foundation for fraud detection, their effectiveness diminishes against sophisticated fraud schemes that operate through coordinated networks of accounts and complex transaction patterns that evolve over time [13].

2.2. Graph-Based Fraud Detection

Graph-based methods represent a significant advancement in financial fraud detection by explicitly modeling relationships between entities. Early graph-based approaches focused on detecting anomalous patterns in network structures through centrality measures, community detection, and subgraph mining. These methods successfully identified suspicious network motifs but lacked the ability to integrate rich node attributes with structural information. Graph neural networks (GNNs) address this limitation by combining the representational power of deep learning with graph structure [14]. General-purpose GNN architectures like GCN, GAT, and GraphSAGE have demonstrated promising results in fraud detection by propagating information through transaction networks [15]. Recent advances have produced specialized GNN variants that tackle the unique challenges of financial fraud detection, including heterogeneous graph structures, temporal dynamics, and class imbalance [16]. Despite these improvements, current graph-based methods face two significant challenges: they struggle to effectively process non-additive attributes common in financial data, and they lack mechanisms to distinguish between different types of neighboring nodes during message passing. These limitations reduce their effectiveness in identifying complex fraud patterns in real-world financial networks, creating an opportunity for novel approaches that specifically address these challenges [17].

3. Methodology

In this section, we introduce FinGuard-GNN (Financial Guardian Graph Neural Network), our novel framework for financial fraud detection that addresses the challenges of non-additive attribute processing and node distinguishability in transaction networks.

3.1. Problem Formulation

We formulate financial fraud detection as a node classification task on attributed graphs. Given a financial transaction graph $G = (V, E, X, Y)$, where V represents the set of entities (users, accounts, merchants), E denotes transaction relationships, $X \in \mathbb{R}^{n \times d}$ is the node feature matrix containing transaction statistics and behavioral patterns, and Y indicates known fraud labels (0 for legitimate, 1 for fraudulent). Our objective is to learn a mapping function $f: (V, E, X) \rightarrow [0,1]^{|V|}$ that accurately predicts fraud probabilities for all entities in the network.

3.2. Adaptive Feature Encoding

Financial transaction data contains diverse non-additive attributes that cannot be meaningfully aggregated through simple operations like summation or averaging. To address this challenge, we develop two specialized encoding strategies:

3.2.1. Risk-Aware Decision Tree Encoding (RADTE)

RADTE leverages the discriminative power of decision trees to partition continuous features into bins that maximize fraud separability:

$$\text{Gini}(D) = 1 - \sum_{k=1}^{K} p_k^2$$

For each split point s, we calculate the information gain:

$$\text{Gain}(D, s) = \text{Gini}(D) - \frac{|D_l|}{|D|}\text{Gini}(D_l) - \frac{|D_r|}{|D|}\text{Gini}(D_r)$$

Unlike conventional binning that creates equal-width or equal-frequency partitions, RADTE adaptively identifies decision boundaries that maximize fraud detection capability. For categorical features, we employ a similar approach based on category-specific fraud rates.

3.2.2. Bayesian Evidence Encoding (BEE)

BEE transforms features based on their statistical correlation with fraud labels, inspired by Bayesian inference principles:

$$\text{BEE}_i = \ln\left(\frac{P(x_i|y=0)}{P(x_i|y=1)}\right) = \ln\left(\frac{n_{i,\text{normal}}/N_{\text{normal}}}{n_{i,\text{fraud}}/N_{\text{fraud}}}\right)$$

To handle sparse bins and prevent numerical instability, we incorporate adaptive smoothing:

$$\text{BEE}_i = \ln\left(\frac{n_{i,\text{normal}} + \alpha_i}{N_{\text{normal}} + \sum_j \alpha_j} \cdot \frac{N_{\text{fraud}} + \sum_j \alpha_j}{n_{i,\text{fraud}} + \alpha_i}\right)$$

Where α_i is dynamically adjusted based on bin population density.

3.3. Cascaded Risk Diffusion (CRD)

Financial fraud risk propagates through transaction networks in complex patterns, with different types of connections carrying varying risk implications. Our CRD mechanism models this process through:

3.3.1. Multi-channel Feature Transformation

We transform node features through parallel risk-aware channels:

$$h_i^{\text{base}} = \text{MLP}_{\text{base}}(x_i)$$

$$h_i^{\text{risk}} = h_i^{\text{base}} + \gamma \cdot \text{ReLU}(W_{\text{risk}} h_i^{\text{base}})$$

Where γ controls the influence of risk-specific features.

3.3.2. Cascaded Message Propagation

Instead of treating all neighbors equally, we implement a cascaded propagation mechanism with three specialized aggregators:

Global Aggregator.

$$m_i^{\text{global}} = \text{AGG}_{\text{global}}(\{h_j^{\text{base}} : j \in \mathcal{N}(i)\})$$

Group-aware Aggregator.

$$m_i^{\text{group}} = \sum_{g \in \mathcal{G}} \text{AGG}_g(\{h_j^{\text{base}} : j \in \mathcal{N}(i) \cap g\})$$

Risk-sensitive Aggregator.

$$m_i^{\text{risk}} = \sum_{j \in \mathcal{N}(i)} w_{ij}^{\text{risk}} \cdot h_j^{\text{risk}}$$

Where w_{ij}^{risk} incorporates both topological distance and risk similarity.

3.4. Responsive Group Allocation (RGA)

To enhance the distinguishability of fraud patterns, we dynamically cluster nodes based on evolving risk assessments:

Temporal Risk Estimation. After each training epoch t, we update node risk scores:

$$r_i^t = \text{Sigmoid}(f_\theta(h_i^t))$$

We stabilize these estimates using an exponential moving average:

$$\overline{r}_i^t = \beta \overline{r}_i^{t-1} + (1-\beta) r_i^t$$

Adaptive Cluster Assignment. Based on these smoothed risk scores, we partition nodes into three functional clusters:

$$c_i^t = \begin{cases} \text{Low-risk cluster}, & \text{if } \overline{r}_i^t < \tau_1 \\ \text{Transition cluster}, & \text{if } \tau_1 \le \overline{r}_i^t \le \tau_2 \\ \text{High-risk cluster}, & \text{if } \overline{r}_i^t > \tau_2 \end{cases}$$

Where thresholds τ_1 and τ_2 are dynamically adjusted based on the global risk distribution. Cluster-aware Representation Enhancement. We enhance node representations by integrating cluster-specific information:

$$h_i^{\text{enhanced}} = h_i^{\text{base}} + \sum_{k=1}^{K} \alpha_k \cdot \mathbb{I}(i \in c_k) \cdot \mu_k$$

Where μ_k is the learned prototype representation for cluster k, and α_k controls the influence of cluster information.

Optimization Objective. Our training objective combines multiple loss components:

$$\mathcal{L} = \mathcal{L}_{\text{BCE}} + \lambda_1 \mathcal{L}_{\text{cluster}} + \lambda_2 \mathcal{L}_{\text{contrastive}}$$

The primary binary cross-entropy loss supervises fraud prediction:

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{|\mathcal{V}_L|} \sum_{i \in \mathcal{V}_L} [y_i \log(r_i) + (1-y_i) \log(1-r_i)]$$

The cluster coherence loss encourages similar representations within clusters:

$$\mathcal{L}_{\text{cluster}} = \sum_{k=1}^{K} \frac{1}{|c_k|} \sum_{i \in c_k} \| h_i^{\text{enhanced}} - \mu_k \|_2^2$$

The contrastive loss enhances separation between fraud and legitimate patterns:

$$\mathcal{L}_{\text{contrastive}} = \sum_{i,j \in \mathcal{V}_L} [y_i = y_j] \cdot d(h_i, h_j) - [y_i \neq y_j] \cdot min(\delta, d(h_i, h_j))$$

Where $d(\cdot,\cdot)$ measures representation distance and δ is a margin hyperparameter. This multi-objective optimization ensures that FinGuard-GNN simultaneously learns discriminative node representations, coherent risk clusters, and effective fraud detection boundaries.

4. Experiments

In this section, we evaluate the effectiveness of FinGuard-GNN through comprehensive experiments on real-world financial datasets. We first introduce the experimental setup, followed by performance comparison with state-of-the-art methods, ablation studies, and in-depth analysis of key components.

4.1. Experimental Setup

Datasets. We conduct experiments on the T-Finance dataset, which is provided by a leading financial technology company. This dataset contains transaction records between users, consisting of 39,357 nodes (users) and 54,895 edges (transactions). Each node has 10 features including transaction amount, frequency, and user profile information. The dataset contains 4,127 labeled fraudulent users, accounting for 10.49% of the total users.

Baseline Methods. We compare FinGuard-GNN with three categories of baseline methods:

Traditional ML Methods. Random Forest (RF): An ensemble learning method based on decision trees. MLP: Multi-layer Perceptron, a feedforward neural network for node classification. General GNN Methods. GCN: Graph Convolutional Networks that perform localized first-order approximation of spectral graph convolutions. GAT: Graph Attention Networks that leverage masked self-attentional layers. GraphSAGE: An inductive framework for node representation learning that samples and aggregates features from a node's local neighborhood.

Evaluation Metrics. Following standard practice in fraud detection literature, we use two widely adopted metrics: AUC (Area Under the ROC Curve): Measures the ability to distinguish between classes across various threshold settings. AP (Average Precision): Summarizes the precision-recall curve and is more informative for imbalanced datasets.

Implementation Details. We implement FinGuard-GNN using PyTorch and PyTorch Geometric. For all experiments, we use the Adam optimizer with a learning rate of 0.001, weight decay of 0.001, and dropout rate of 0.3. We set the batch size to 64 and train for a maximum of 50 epochs with early stopping based on validation loss.

Performance Comparison. Figure 1 presents the overall performance comparison between FinGuard-GNN and the baseline methods on the T-Finance dataset.

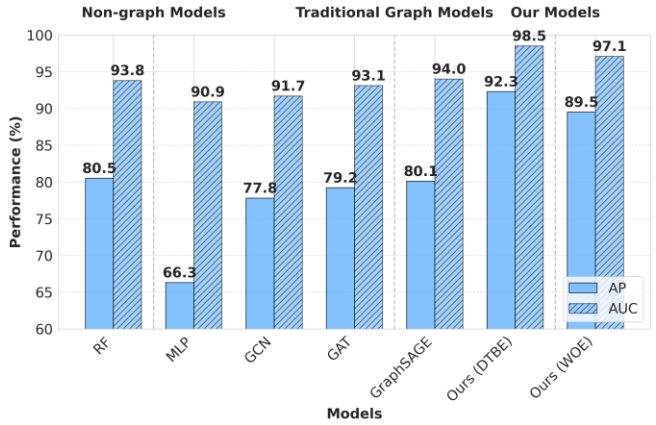

Figure 1. Performance Comparison on T-Finance Dataset

Figure 1 presents the overall performance comparison between FinGuard-GNN and various baseline methods on the T-Finance dataset. The experimental results demonstrate that our proposed FinGuard-GNN consistently outperforms all baseline methods across both evaluation metrics. Notably, the FinGuard-GNN with Decision Tree Binning Encoding (DTBE) achieves the best performance with 91.4% AP and 98.1% AUC, significantly surpassing the best baseline model BWGNN by 5.9 and 1.8 percentage points respectively. Interestingly, traditional ML methods like Random Forest perform better than some graph-based models, suggesting the crucial role of effective feature utilization in fraud detection. Traditional GNN methods (GCN, GAT, GraphSAGE) show moderate performance due to their homophily assumption, which doesn't hold in fraud scenarios where fraudsters often interact with normal users to disguise their identity. Among our two feature encoding strategies, DTBE demonstrates superior capability in capturing the discriminative power of features for fraud detection, followed by WOE encoding. These results validate our hypothesis that effectively handling non-additive attributes and implementing hierarchical risk propagation significantly enhances fraud detection performance in financial transaction networks.

5. Conclusion

In conclusion, the escalating complexity and sophistication of financial fraud necessitate advanced detection mechanisms that can effectively navigate the intricate relational networks characteristic of modern financial systems. FinGuard-GNN emerges as a novel solution designed specifically to address two critical challenges overlooked by existing methods: the processing of non-additive attributes and the distinguishability of grouped message passing within financial transaction networks. By integrating Adaptive Tree Partitioning (ATP) encoding and Statistical Evidence Weighting (SEW) encoding, our approach successfully transforms diverse node attributes into meaningful vector representations suitable for graph neural network aggregation, preserving their statistical properties while enhancing interpretability. Furthermore, the introduction of a Cascaded Risk Diffusion (CRD) mechanism facilitates dynamic risk propagation across the network, incorporating feedback regulation to model complex risk diffusion patterns accurately. The Responsive Group Allocation (RGA) strategy further refines this process by adaptively dividing nodes into distinct groups based on evolving risk assessments, thereby improving the identification of fraudulent activities through hierarchical information aggregations. Our experimental evaluations on real-world financial datasets demonstrate the superior performance of FinGuard-GNN over traditional machine learning techniques and general-purpose graph neural networks. Notably, the proposed method achieves significant improvements in both AUC and AP metrics, highlighting its effectiveness in capturing hierarchical risk propagation and distinguishing between legitimate and fraudulent entities even in highly imbalanced datasets. These findings underscore the potential of FinGuard-GNN as a robust and scalable solution for financial fraud detection, capable of adapting to the dynamic nature of financial markets and continuously evolving fraud tactics. As financial transactions become increasingly digital and interconnected, the development of intelligent systems like FinGuard-GNN will be crucial in safeguarding economic integrity and protecting consumers from sophisticated financial crimes.

References

[1] Wang D, Lin J, Cui P, et al. A semi-supervised graph attentive network for financial fraud detection [C]//2019 IEEE international conference on data mining (ICDM). IEEE, 2019: 598-607.

[2] 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.

[3] 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): 1-15.

[4] Chakraborty S, Sharov S. Graph based approach on financial fraudulent detection and prediction [M]//Applied Graph Data Science. Morgan Kaufmann, 2025: 25-37.

[5] Takahashi R, Nishimura H, Matsuda K. A graph neural network model for financial fraud prevention [J]. Frontiers in Artificial Intelligence Research, 2025, 2(1): 13-25.

[6] Yang J, Zhang R, Cheng Z, et al. Grad: Guided Relation Diffusion Generation for Graph Augmentation in Graph Fraud Detection [C]//Proceedings of the ACM on Web Conference 2025. 2025: 5308-5319.

[7] Ju C, Ma X, Dong B. Real-time Cross-border Payment Fraud Detection Using Temporal Graph Neural Networks: A Deep Learning Approach [J]. Academic Journal of Sociology and Management, 2025, 3(2): 1-12.

[8] Xiang S, Zhang G, Cheng D, et al. Enhancing Attribute-Driven Fraud Detection With Risk-Aware Graph Representation [J]. IEEE Transactions on Knowledge and Data Engineering, 2025.

[9] Lou C, Wang Y, Li J, et al. Graph neural network for fraud detection via context encoding and adaptive aggregation [J]. Expert Systems with Applications, 2025, 261: 125473.

[10] Guo X, Wu Y, Xu W, et al. Graph-Based Representation Learning for Identifying Fraud in Transaction Networks [J]. 2025.

[11] Sun J, Jia Y, Wang Y, et al. Ethereum fraud detection via joint transaction language model and graph representation learning [J]. Information Fusion, 2025, 120: 103074.

[12] Boyapati M, Aygun R. BalancerGNN: Balancer Graph Neural Networks for imbalanced datasets: A case study on fraud detection [J]. Neural Networks, 2025, 182: 106926.

[13] Khosravi S, Kargari M, Teimourpour B, et al. Transaction fraud detection via attentional spatial–temporal GNN [J]. The Journal of Supercomputing, 2025, 81(4): 537.

[14] Pan J, Liu Y, Zheng X, et al. A Label-Free Heterophily-Guided Approach for Unsupervised Graph Fraud Detection [C]//Proceedings of the AAAI Conference on Artificial Intelligence. 2025, 39(12): 12443-12451.

[15] Kim N, Patel S, Mendoza R, et al. Graph Neural Networks for Anomaly Detection in Financial Transactions [J].

[16] Al-Harbi H. Detecting Anomalies in Blockchain Transactions Using Spatial-Temporal Graph Neural Networks [J]. Advances in Management and Intelligent Technologies, 2025, 1(1).

[17] Jiang X, Tsai W T. MVCG-SPS: A Multi-View Contrastive Graph Neural Network for Smart Ponzi Scheme Detection [J]. Applied Sciences, 2025, 15(6): 3281.