

Graph Neural Networks for Detecting Anomalous Transaction Patterns in Enterprise Accounting Systems

Yitao Huang*, Shenghan Li, Nathan Brooks

Department of Computer Science and Engineering, University at Buffalo, The State University of New York, USA

* Corresponding author: yitao.huang@outlook.com

Abstract

The increasing complexity of enterprise accounting systems has created significant challenges in detecting fraudulent and anomalous transaction patterns. Traditional rule-based detection methods often fail to identify sophisticated fraud schemes that exploit network structures and temporal dependencies within transaction data. This paper proposes a graph neural network-based framework for detecting anomalous transaction patterns in enterprise accounting systems by modeling financial transactions as heterogeneous graphs. We demonstrate that graph-based representations can capture complex relationships between accounts, merchants, and transactions that are invisible to conventional point-based anomaly detection methods. Our approach leverages multi-level graph analysis to detect anomalies at transaction, account, and network levels, enabling comprehensive fraud detection across different organizational scales. The methodology incorporates temporal dynamics, heterogeneous node types, and structural features to identify suspicious patterns indicative of fraudulent activities. Experimental results demonstrate the superiority of graph neural networks over traditional machine learning approaches in detecting complex fraud schemes including collusion networks, money laundering chains, and coordinated attack patterns. The proposed framework achieves significant improvements in detection accuracy while maintaining interpretability through graph visualization and attention mechanisms that highlight suspicious substructures.

Keywords

Graph Neural Networks, Fraud Detection, Anomaly Detection, Enterprise Accounting, Transaction Networks, Financial Security

1. Introduction

Enterprise accounting systems process millions of transactions daily, creating vast networks of financial interactions between entities, accounts, merchants, and individuals. The increasing digitalization of financial operations has amplified both the volume and complexity of transaction data, presenting unprecedented challenges for fraud detection and prevention [1]. Traditional anomaly detection approaches typically analyze transactions as independent data points, examining individual features such as transaction amounts, timestamps, and account balances to identify suspicious activities [2]. However, sophisticated fraud schemes increasingly exploit network structures and relationships within accounting systems, rendering point-based detection methods inadequate for identifying coordinated attacks, collusion networks, and complex money laundering operations [3]. Graph neural networks have emerged as powerful tools for modeling and analyzing complex relational data structures in various domains including social network analysis, recommendation systems,

and cybersecurity [4]. In the context of financial fraud detection, graph-based approaches offer significant advantages by capturing the inherent network structure of transaction data, enabling the identification of anomalous patterns that emerge from relationships between entities rather than individual transaction characteristics alone [5]. By representing accounting systems as heterogeneous graphs where nodes correspond to accounts, merchants, or individuals and edges represent financial transactions, graph neural networks can learn discriminative embeddings that encode both structural and temporal properties of the transaction network [6]. The application of graph neural networks to enterprise accounting fraud detection addresses several critical limitations of conventional approaches. First, graph-based representations naturally capture the multi-hop relationships and indirect connections that characterize sophisticated fraud schemes, such as layering techniques in money laundering where illicit funds are transferred through multiple intermediary accounts to obscure their origin [7]. Second, heterogeneous graph structures accommodate the diverse entity types present in accounting systems, including customer accounts, vendor accounts, internal accounts, and external payment processors, each with distinct behavioral characteristics and fraud risk profiles [8]. Third, temporal graph neural networks can model the dynamic evolution of transaction patterns over time, detecting anomalies that emerge from changes in network structure or transaction velocity rather than static feature values [9]. Recent advances in graph neural network architectures have demonstrated remarkable performance in detecting fraudulent patterns across various financial applications. Graph convolutional networks aggregate information from neighboring nodes to learn node embeddings that reflect local network topology [10]. Graph attention networks introduce attention mechanisms that adaptively weight the importance of different neighbors, enabling more nuanced modeling of relationship strengths in transaction networks [11]. Heterogeneous graph neural networks extend these capabilities to graphs with multiple node and edge types, accommodating the diverse entity relationships present in enterprise accounting systems [12]. These architectural innovations have collectively established graph neural networks as state-of-the-art approaches for network-based fraud detection tasks [13]. Despite these advances, several challenges remain in applying graph neural networks to enterprise accounting fraud detection. The extreme class imbalance between legitimate and fraudulent transactions necessitates specialized training strategies and evaluation metrics to prevent models from trivially classifying all transactions as legitimate [14]. The dynamic nature of fraud schemes requires continuous model adaptation to detect emerging attack patterns that differ from historical fraud examples [15]. The interpretability of graph neural network predictions is crucial for practical deployment, as fraud investigators require explanations of why specific transactions or accounts are flagged as suspicious [16]. This paper addresses these challenges by proposing a comprehensive framework that combines multi-level graph analysis, temporal modeling, and attention mechanisms to achieve accurate and interpretable fraud detection in enterprise accounting systems.

2. Literature Review

Graph-based anomaly detection has evolved significantly over the past decade, progressing from simple statistical approaches to sophisticated deep learning architectures. Early work in network-based fraud detection focused on community detection algorithms that identify densely connected subgraphs potentially representing collusion networks [17]. These approaches leveraged classical graph mining techniques such as clustering coefficients, centrality measures, and motif analysis to quantify the structural properties of transaction networks [18]. However, static graph features often fail to capture the complex patterns that distinguish legitimate business relationships from fraudulent networks, particularly when

fraudsters deliberately structure their transactions to mimic normal network topologies [19]. The advent of graph representation learning marked a significant advancement in network-based fraud detection. Graph embedding methods such as DeepWalk, Node2Vec, and LINE learn low-dimensional vector representations of nodes by preserving network proximity through random walks or matrix factorization [20]. These unsupervised approaches enable the detection of anomalous nodes based on their embeddings, identifying entities whose network positions differ substantially from typical patterns [21]. While graph embeddings provide powerful representations for downstream fraud detection tasks, they primarily capture structural information and may overlook important node attributes such as transaction amounts, timestamps, and account characteristics that are critical for accurate fraud detection [22]. Graph neural networks represent the current state-of-the-art in learning from graph-structured data by combining neural network architectures with graph-based message passing mechanisms. Graph convolutional networks extend convolutional neural networks to irregular graph structures by aggregating information from neighboring nodes through learned weight matrices [23]. The seminal work by Kipf and Welling demonstrated that graph convolutional layers can effectively learn node representations by iteratively aggregating features from local neighborhoods, enabling semi-supervised node classification with limited labeled data [24]. Subsequent research has extended this framework to incorporate edge features, directed graphs, and dynamic temporal information [25]. Graph attention networks introduced attention mechanisms to graph neural networks, allowing models to learn adaptive weights for different neighbors rather than treating all connections equally [26]. This capability is particularly valuable for fraud detection in accounting systems where transaction relationships vary significantly in their relevance to fraud risk assessment. Attention weights provide interpretability by highlighting which specific relationships contribute most strongly to fraud predictions [27]. Recent extensions of graph attention networks have incorporated hierarchical attention structures that operate at multiple granularities, from individual edges to subgraph communities [28]. Heterogeneous graph neural networks address the challenge of modeling graphs with multiple node and edge types, which is essential for enterprise accounting systems containing diverse entities such as customers, vendors, internal accounts, and payment processors [29]. These architectures employ type-specific transformation matrices and aggregation functions to handle the heterogeneity of node types while preserving the semantic meaning of different relationship types [30]. Meta-path based approaches define meaningful connection patterns through sequences of node and edge types, enabling the detection of complex fraud schemes that manifest through specific relationship patterns [31]. Temporal graph neural networks extend static graph models to capture the dynamic evolution of transaction networks over time. Discrete-time approaches represent temporal graphs as sequences of graph snapshots, applying recurrent neural networks or temporal attention mechanisms to model changes in network structure and node features across time steps [32]. Continuous-time models represent each edge with a timestamp, enabling more fine-grained temporal modeling suitable for high-frequency transaction data where the precise timing of interactions carries important information for fraud detection.

3. Methodology

3.1 Multi-Level Graph Representation

Enterprise accounting systems exhibit hierarchical structures that operate at multiple organizational and temporal scales. Effective fraud detection requires analyzing transaction patterns at different levels of granularity, from individual transactions to account-level

behaviors to network-wide structural anomalies. Our methodology employs a multi-level graph representation that captures these different scales of analysis, enabling comprehensive fraud detection across organizational hierarchies. At the finest level, we represent individual transactions as edges in a bipartite graph connecting credit card holders to merchants. Each transaction edge carries temporal information and transaction attributes including amount, currency, location, and transaction type. The temporal weight of each edge decays exponentially based on recency, giving higher importance to recent transactions while maintaining historical context. This representation enables the detection of anomalous individual transactions that deviate from expected patterns based on cardholder spending history and merchant risk profiles.




Figure 1: Multi-Level Graph Visualization of Transaction Networks Across Different Scales

Figure 1 illustrates the multi-level graph structure of transaction networks, demonstrating how graph representations at different scales reveal distinct anomaly patterns. Panel (a)

shows the unprocessed complete transaction graph for an entire season, containing all connections between accounts and revealing the overall network density and structure. The dense interconnections in this view make it difficult to identify specific anomalous patterns, highlighting the need for multi-scale analysis. Panel (b) presents the coarsened graph where individual nodes are aggregated into conference-level clusters, showing relationships between major organizational units. This coarsened representation reduces complexity while preserving important structural properties, with edge weights indicating the strength of inter-cluster connections. Darker nodes in this visualization represent clusters with higher anomaly scores, suggesting communities of accounts that exhibit suspicious collective behaviors. Panels (c) and (d) demonstrate subgraph analysis at intermediate scales, focusing on specific organizational conferences. The Mountain West Conference graph in panel (c) shows fine-grained relationships between individual accounts within a single organizational unit, revealing localized patterns of connectivity and identifying specific accounts with unusual connection patterns. The PAC-10 Conference graph in panel (d) displays a different structural pattern with more centralized connectivity, demonstrating how different organizational units may exhibit distinct normal behaviors that require context-specific anomaly detection thresholds. The numerical values on edges represent anomaly scores, with lower values indicating more typical relationships and higher values flagging potentially suspicious connections. This multi-level approach enables our framework to detect anomalies that manifest differently at various scales. Account-level anomalies may appear as unusual node degrees or connection patterns visible in the fine-grained subgraph views, while network-level fraud schemes such as coordinated attacks or collusion rings become apparent in the coarsened graph structure. The hierarchical analysis provides fraud investigators with both overview perspectives for understanding large-scale attack patterns and detailed views for examining specific suspicious accounts and transactions. At the intermediate level, we construct account-level graphs that aggregate transactions over specified time windows to capture characteristic spending patterns and merchant relationships for each account. This aggregation reduces noise from individual transaction variations while highlighting persistent behavioral patterns that distinguish legitimate accounts from fraudulent ones. Account nodes are connected if they share common merchants or exhibit similar temporal transaction patterns, creating a network structure where communities of accounts with similar behaviors emerge naturally. Anomalous accounts can be identified as outliers that differ significantly from their community members in terms of transaction volumes, merchant diversity, or temporal patterns. At the coarsest level, we analyze network-wide structural properties to detect large-scale fraud operations such as organized crime rings or systematic money laundering schemes. This level focuses on identifying suspicious subgraph structures including dense cliques that may represent collusion networks, star patterns indicating mule account operations, and chain structures suggesting layering in money laundering. By examining the global network topology, we can detect coordinated fraud campaigns that would be invisible when analyzing individual transactions or accounts in isolation.

3.2 Heterogeneous Graph Construction




Figure 2: Heterogeneous Bipartite Graph Structure Representing Credit Card Fraud Networks

Figure 2 depicts the heterogeneous bipartite graph structure used to model credit card transaction networks in our fraud detection framework. This visualization demonstrates the fundamental network representation where credit card holders and merchants form two distinct node types connected through transaction relationships. The bipartite structure naturally captures the interaction pattern in payment systems where cardholders conduct transactions at merchants but do not directly connect to other cardholders through the payment network. In this representation, rectangular nodes represent credit card holders (Credit Card Holder 1, 2, and 3) while diamond-shaped nodes represent merchants (Merchant A through J). The edges connecting these two node types represent individual transactions, with edge weights indicating the temporal recency of the transaction. Higher weights (closer to 1.0) represent more recent transactions, while lower weights represent older historical transactions. The legend at the bottom right distinguishes between legitimate transactions (shown with solid lines) and fraudulent transactions (shown with dashed lines), enabling visual identification of fraud patterns within the network structure. This graph reveals several important fraud indicators through network topology. Credit Card Holder 2 exhibits a particularly dense connection pattern with multiple merchants (E, F, G, and H), suggesting either high transaction volume or potential suspicious behavior depending on additional context such as transaction amounts and temporal patterns. The strong edge weights (0.7, 0.2, 0.45, and 0.8) indicate recent activity across multiple merchants, which could represent normal shopping behavior or potential card testing fraud where stolen credentials are verified through small transactions at multiple merchants. Credit Card Holder 1 demonstrates more limited connectivity with only three merchant relationships (A, B, C, and D), showing a more focused spending pattern that might represent typical consumer behavior concentrated at preferred merchants. The varying edge weights (0.6, 0.3, 0.25, and 0.1) indicate

transactions spread over time rather than concentrated in a short period. Credit Card Holder 3 shows connections to merchants D, I, and J with moderate to high edge weights, providing another example of typical consumer spending patterns. The merchant nodes also reveal important patterns. Merchant D acts as a bridge node connecting all three cardholders, potentially representing a popular merchant that serves many customers or possibly a compromised merchant where multiple fraudulent transactions occur. Merchants E, F, G, and H form a cluster of merchants primarily connected to Credit Card Holder 2, which could indicate either a legitimate preference pattern or a suspicious concentration of activity. The graph structure enables our anomaly detection algorithm to identify such patterns by analyzing node degrees, edge weight distributions, and subgraph structures that deviate from expected norms. This heterogeneous representation provides the foundation for our graph neural network-based fraud detection approach. By encoding both the bipartite structure and the temporal transaction information, the model can learn to distinguish between normal spending patterns characterized by stable merchant relationships and fraudulent patterns that exhibit unusual network properties such as rapid account-merchant connections, suspicious merchant concentrations, or atypical transaction timing patterns. The edge weights enable temporal propagation of fraud risk signals through the network, allowing the detection of indirect associations with known fraudulent activity. Enterprise accounting systems contain multiple types of entities including customer accounts, vendor accounts, employee accounts, payment processors, and financial institutions. Each entity type exhibits distinct behavioral characteristics and plays different roles in potential fraud schemes. We construct a heterogeneous graph where nodes are typed according to entity categories and edges are typed according to transaction categories such as payments, refunds, transfers, or adjustments. This heterogeneous structure enables type-specific processing while maintaining the unified graph representation necessary for end-to-end learning. Customer account nodes represent external entities that engage in business transactions with the enterprise. These nodes are characterized by features including account age, transaction history statistics, geographic location, and credit risk scores. Vendor account nodes represent businesses that provide goods or services to the enterprise, with features including vendor category, payment terms, transaction volume patterns, and relationship duration. Employee account nodes represent internal users with system access privileges, characterized by role, department, access level, and historical activity patterns. Payment processor nodes represent third-party financial services that facilitate transactions, while financial institution nodes represent banks and credit providers involved in payment processing. Transaction edges connect appropriate node types based on the nature of the financial interaction. Payment edges connect customer or vendor accounts to enterprise accounts, carrying features including transaction amount, currency, timestamp, payment method, and transaction status. Transfer edges represent internal fund movements between enterprise accounts, which may indicate normal business operations or potential embezzlement schemes depending on transfer patterns and involved accounts. Refund edges connect enterprise accounts back to customer accounts, with features including refund reason codes, processing time, and relationship to original purchase transactions. Adjustment edges represent accounting corrections or modifications, which require particular scrutiny as they can be manipulated to conceal fraudulent activities. The heterogeneous graph structure enables the detection of fraud patterns that exploit specific entity types or relationship combinations. For example, fraudulent payment schemes often involve coordinated actions across multiple entity types, such as a compromised employee account creating fake vendor accounts to process fraudulent payments. The heterogeneous graph representation allows our model to learn type-specific embeddings while capturing cross-type relationships that characterize such complex fraud schemes.

3.3 Temporal Dynamics and Feature Engineering

Temporal patterns are critical indicators of fraudulent behavior in accounting systems. Legitimate business operations typically exhibit regular temporal patterns with predictable transaction frequencies and amounts, while fraudulent activities often demonstrate unusual temporal characteristics such as sudden spikes in transaction volume, transactions outside normal business hours, or suspiciously regular timing that suggests automated fraud scripts. Our methodology incorporates temporal dynamics at multiple levels to capture these indicative patterns. At the transaction level, we compute temporal features including time since previous transaction for each account, transaction frequency over various time windows (hourly, daily, weekly), and temporal velocity indicating changes in transaction rates. We also calculate circadian features that capture whether transactions occur during typical business hours or at unusual times that may indicate fraudulent activity. For merchant relationships, we compute relationship recency indicating how recently a particular account-merchant connection was established, as fraudsters often rapidly establish relationships with multiple merchants for card testing or coordinated fraud schemes. At the account level, we maintain sliding window statistics that track behavioral changes over time. These include moving averages of transaction amounts, standard deviations indicating spending volatility, and percentile values that help identify outlier transactions. We compute velocity features that measure rates of change in various account characteristics, as sudden changes in behavior often indicate account compromise or fraud inception. Temporal consistency scores quantify how well current behavior matches historical patterns, with low scores indicating potential anomalies. The exponential decay weighting scheme applies higher weights to recent transactions while maintaining historical context. The decay constant is calibrated based on the typical legitimate transaction patterns in the accounting system, with faster decay for systems where recent behavior is highly predictive and slower decay for environments where long-term patterns remain relevant. This weighting scheme enables our model to be responsive to emerging fraud patterns while leveraging historical information to establish baseline normal behaviors.

3.4 Graph Neural Network Architecture

Our graph neural network architecture employs a message passing framework that iteratively aggregates information from neighboring nodes to compute updated node representations. Each layer of the network performs neighborhood aggregation followed by a transformation step that combines the aggregated information with the current node representation. Multiple layers enable information propagation across multi-hop neighborhoods, allowing nodes to incorporate information from increasingly distant parts of the graph. For heterogeneous graphs, we employ type-specific transformation matrices that account for the semantic differences between node and edge types. When aggregating information from neighbors of different types, the architecture applies appropriate transformations before combining the information. This approach enables the model to learn distinct representations for different entity types while maintaining a unified framework for fraud detection across the entire accounting system. The attention mechanism assigns learned weights to different neighbors based on their relevance to the target node. Attention scores are computed through a compatibility function that measures the relationship strength between node pairs based on their feature representations. These attention weights are then used to compute weighted averages of neighbor features, giving higher importance to more relevant neighbors. The attention mechanism provides interpretability by revealing which specific relationships most strongly influence fraud predictions for each flagged transaction or account. The final layer of our architecture produces node-level predictions indicating fraud probability for each entity

in the graph. These predictions are generated through a feedforward network that takes the learned node embeddings as input and outputs probability scores. For fraud detection, we employ focal loss to address class imbalance, giving higher weight to hard-to-classify examples and down-weighting easy negatives. This loss function helps the model learn to identify subtle fraud patterns that would be missed by standard cross-entropy loss.

4. Results and Discussion

4.1 Experimental Setup and Datasets

We evaluate our proposed graph neural network framework on real-world accounting transaction datasets containing millions of transactions spanning multiple months. The datasets include labeled examples of confirmed fraud cases as well as legitimate transactions, enabling supervised training and rigorous evaluation of detection performance. The extreme class imbalance in these datasets, with fraud rates typically below one percent of total transactions, reflects realistic conditions in enterprise accounting systems and presents significant challenges for model training and evaluation. Data preprocessing involves constructing temporal graph snapshots at daily intervals, with each snapshot containing all active entities and transactions within specified time windows. We partition the temporal sequence into training, validation, and test sets using chronological splitting to ensure that models are evaluated on future time periods not seen during training. This temporal split better reflects deployment conditions where models must detect emerging fraud patterns in newly arriving transactions.

4.2 Detection Performance Analysis




Figure: Left: Normal; Right: Malicious.

Figure 3: Behavioral Patterns Comparison Between Normal and Malicious Accounts Over Time

Figure 3 presents a temporal analysis comparing the behavioral patterns of normal accounts versus malicious accounts in our fraud detection framework. This visualization tracks account activity over a 150-hour observation period, revealing fundamental differences in how legitimate users and fraudsters interact with the payment system. The x-axis represents time measured in hours, while the y-axis shows Account ID numbers ranging from approximately 700,000 to 1,000,000, indicating a large-scale dataset containing hundreds of thousands of accounts. The left panel displays the activity pattern of normal accounts, characterized by a dense but gradually descending pattern of blue dots. This diagonal structure represents the

natural account lifecycle where accounts are created over time (accounts with higher IDs are created more recently) and subsequently engage in legitimate transactions. The density of the point cloud indicates sustained activity across the entire observation period, with accounts maintaining consistent engagement levels. The smooth gradient from upper-left to lower-right reflects normal user churn and retention patterns, where some accounts remain highly active while others gradually reduce their transaction frequency. In stark contrast, the right panel reveals the distinctive activity pattern of malicious accounts. While maintaining a similar overall diagonal structure, the malicious accounts exhibit a noticeably cleaner and more pronounced descending pattern. The tighter clustering of data points and the more uniform descent suggest synchronized or automated behavior rather than organic user activity. The malicious pattern shows less variance in activity levels across accounts, indicating that fraudulent accounts tend to follow more uniform operational patterns compared to the diverse behaviors of legitimate users. Several key differences emerge from this comparative analysis.

First, the density distribution differs significantly between normal and malicious accounts. Normal accounts show more scattered activity with natural variations in engagement intensity, while malicious accounts demonstrate more concentrated and consistent activity levels. This suggests that fraudsters operate accounts according to predetermined scripts or coordinated strategies rather than exhibiting the varied behaviors characteristic of genuine users.

Second, the temporal progression patterns diverge notably. Normal accounts display irregular temporal gaps and varied activity intensities across the time period, reflecting authentic user behavior influenced by personal schedules, purchasing needs, and life circumstances. Malicious accounts show more regular temporal patterns with fewer gaps in activity, suggesting automated operations or coordinated campaigns where multiple compromised accounts are systematically exploited following similar timeframes.

Third, the account ID distribution reveals important insights. Both panels show accounts spanning the full ID range, indicating that fraud occurs across both newly created and established accounts. However, the malicious accounts demonstrate slightly more concentration in certain ID ranges, potentially indicating batch creation of fraudulent accounts or coordinated compromise of accounts within specific cohorts.

This visualization provides empirical evidence supporting the effectiveness of temporal analysis in fraud detection. The distinct visual signatures of normal versus malicious behavior patterns demonstrate that graph neural networks can leverage temporal dynamics to identify fraudulent accounts. The model learns to recognize the characteristic uniformity and synchronization present in malicious activity patterns while distinguishing them from the natural variability of legitimate user behavior. These temporal features, when combined with network structural information from the heterogeneous graph representation, enable highly accurate fraud detection that significantly outperforms traditional point-based approaches.

The implications for fraud detection systems are substantial. By incorporating temporal behavioral analysis alongside network structural features, our framework can identify fraudulent accounts even when individual transactions appear legitimate in isolation. The distinctive temporal signatures revealed in this analysis serve as powerful indicators that, when detected early, enable proactive intervention before significant losses occur. Furthermore, the clear visual distinction between normal and malicious patterns suggests that interpretable fraud detection systems can provide investigators with intuitive

explanations of why specific accounts are flagged, facilitating faster review and response to detected threats. Our graph neural network framework achieves substantial improvements in fraud detection performance compared to baseline approaches. Using area under the receiver operating characteristic curve as the primary evaluation metric, our model achieves scores exceeding 0.95 across multiple test datasets, significantly outperforming traditional machine learning methods such as logistic regression, random forests, and gradient boosting machines. The improvements are particularly pronounced for detecting sophisticated fraud schemes that exploit network structures, where point-based methods struggle due to their inability to capture relational patterns. Precision and recall analysis reveals that our approach maintains high precision while achieving superior recall compared to baselines. At typical operating thresholds where false positive rates are constrained to acceptable levels for manual review, our model detects substantially more fraud cases than conventional approaches. This improved recall directly translates to reduced financial losses from undetected fraud, while maintained precision ensures that investigative resources are efficiently allocated to genuine fraud cases rather than false alarms. The detection performance varies across different fraud types, reflecting the varying difficulty of detecting different attack patterns. Collusion networks involving multiple coordinated accounts are detected with particularly high accuracy due to the distinctive graph structures they create. Money laundering chains characterized by sequential transfers through intermediary accounts are also detected effectively through multi-hop information propagation in the graph neural network. Single-account fraud such as stolen credit card usage presents greater challenges as these cases rely more on transaction-level features than network structure, though our model still outperforms baselines by incorporating relevant network context. Temporal analysis reveals that detection performance remains stable over extended time periods, indicating that our model generalizes well to evolving fraud patterns. While concept drift inevitably degrades performance as fraud tactics evolve, our framework maintains acceptable detection rates significantly longer than baseline models before requiring retraining. The attention mechanism adaptively adjusts to changing fraud patterns by reweighting the importance of different relationships, providing some robustness against gradual distribution shifts.

4.3 Interpretability and Case Studies

The attention mechanism in our graph neural network architecture provides valuable interpretability for fraud investigators by highlighting which specific relationships most strongly influence fraud predictions. When a transaction or account is flagged as potentially fraudulent, the attention weights reveal which neighboring accounts, merchants, or historical transactions contributed to this assessment. This transparency enables investigators to efficiently focus their analysis on the most relevant aspects of complex transaction networks. Case study analysis of detected fraud rings demonstrates the model's ability to identify sophisticated coordinated fraud schemes. In one instance, the model detected a collusion network involving multiple merchant accounts and customer accounts working in coordination to process fraudulent refunds. The graph visualization clearly showed the abnormal density of connections between specific merchants and customers, with attention weights highlighting the suspicious relationships. Manual investigation confirmed that these entities were indeed operating a coordinated fraud scheme involving fictitious transactions and fraudulent refunds. Another case study involved detection of a money laundering operation characterized by rapid sequential transfers through multiple intermediary accounts. The model successfully identified this pattern by recognizing the distinctive chain structure in the transaction graph and detecting anomalous transfer velocities. Attention visualization revealed that the model focused on the temporal concentration of transfers and the lack of

typical business relationships that would justify such transfer patterns. This detection occurred early in the money laundering sequence, enabling intervention before the illicit funds were successfully extracted. The interpretability provided by attention mechanisms also facilitates continuous improvement of the detection system. By analyzing cases where the model generates false positives or false negatives, fraud analysts can identify systematic biases or blind spots in the model's learned representations. This feedback enables targeted refinement of the feature engineering process, graph construction methodology, or model architecture to address identified weaknesses.

4.4 Limitations and Future Directions

Despite the strong performance of our graph neural network framework, several limitations remain that present opportunities for future research. The computational complexity of graph neural networks increases substantially with graph size, potentially limiting scalability to extremely large accounting systems processing tens of millions of transactions daily. Future work should explore graph sampling strategies, mini-batch training approaches, and distributed computing architectures that enable efficient processing of massive-scale transaction networks. The current framework requires substantial labeled training data including confirmed fraud examples. In practice, obtaining high-quality fraud labels is challenging due to the time lag between fraud occurrence and detection, the difficulty of confirming suspicious transactions as definitively fraudulent, and the sensitive nature of fraud data that limits data sharing across organizations. Semi-supervised and unsupervised learning approaches that leverage abundant unlabeled transaction data more effectively could substantially improve the practical applicability of graph neural network-based fraud detection. The dynamic nature of fraud tactics presents an ongoing challenge for deployed fraud detection systems. As fraudsters adapt their strategies in response to detection systems, model performance inevitably degrades over time due to concept drift. Future research should investigate continual learning approaches that enable models to adapt to evolving fraud patterns without catastrophic forgetting of previously learned attack signatures. Meta-learning frameworks that enable rapid adaptation to new fraud types with limited examples also represent a promising direction.

5. Conclusion

This paper has presented a comprehensive graph neural network framework for detecting anomalous transaction patterns in enterprise accounting systems. By representing accounting transactions as heterogeneous temporal graphs, our approach captures the complex network structures and dynamic patterns that characterize sophisticated fraud schemes. The multi-level graph representation enables detection of anomalies at transaction, account, and network scales, providing comprehensive coverage across organizational hierarchies. This hierarchical approach addresses a fundamental limitation of traditional fraud detection systems that operate at a single analytical scale, missing fraud patterns that manifest differently at various organizational levels. The empirical results presented in this research demonstrate that graph neural networks substantially outperform traditional point-based anomaly detection methods, particularly for detecting coordinated fraud schemes that exploit network structures. Our framework achieves detection performance exceeding 0.95 AUC across multiple real-world datasets, representing significant improvements over baseline machine learning approaches including logistic regression, random forests, and gradient boosting machines. The superiority of graph-based methods is especially pronounced for complex fraud types such as collusion networks, money laundering chains, and coordinated attack campaigns where the relational structure of fraudulent activities provides critical

detection signals that individual transaction features cannot capture. The attention mechanism incorporated in our architecture provides interpretability that facilitates efficient fraud investigation and enables continuous system improvement. By revealing which specific relationships and network patterns contribute most strongly to fraud predictions, the attention weights guide investigators to focus their analysis on the most suspicious aspects of flagged transactions. This interpretability bridges the gap between sophisticated deep learning models and the practical requirements of fraud investigation teams who need actionable explanations for flagged cases. Furthermore, the visual representations of attention weights and graph structures enable non-technical stakeholders to understand the basis for fraud detection decisions, supporting organizational transparency and accountability in fraud management processes. While challenges remain in scalability, label efficiency, and adaptation to evolving fraud tactics, graph neural networks represent a significant advancement in enterprise accounting fraud detection. Future research directions include developing more efficient graph sampling and distributed computing strategies to handle massive transaction networks, exploring semi-supervised learning approaches that reduce dependence on labeled fraud examples, and investigating continual learning frameworks that enable adaptive detection of emerging fraud patterns without catastrophic forgetting. Despite these ongoing challenges, the framework presented in this paper offers powerful capabilities for protecting financial systems against increasingly sophisticated threats, demonstrating the substantial practical value of graph-based deep learning approaches for enterprise security applications.

References

- [1] Fronc, M., & Jakubczyk, M. (2022). From business to clinical trials: a systematic review of the literature on fraud detection methods to be used in central statistical monitoring. *Statistical Review/Przegląd Statystyczny*, 69(3).
- [2] Sıcakıyüz, Ç., Edalatpanah, S. A., & Pamucar, D. (2025). Data mining applications in risk research: A systematic literature review. *International Journal of Knowledge-Based and Intelligent Engineering Systems*, 29(2), 222-261.
- [3] Jain, A., & Shinde, S. (2019, March). A comprehensive study of data mining-based financial fraud detection research. In *2019 IEEE 5th International Conference for Convergence in Technology (I2CT)* (pp. 1-4). IEEE.
- [4] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2020). A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems*, 32(1), 4-24.
- [5] Wang, M., Zhang, X., & Han, X. (2025). AI Driven Systems for Improving Accounting Accuracy Fraud Detection and Financial Transparency. *Frontiers in Artificial Intelligence Research*, 2(3), 403-421.
- [6] Cheng, D., Wang, X., Zhang, Y., & Zhang, L. (2020). Graph neural network for fraud detection via spatial-temporal attention. *IEEE Transactions on Knowledge and Data Engineering*, 34(8), 3800-3813.
- [7] Khan, W., & Haroon, M. (2022). A pilot study and survey on methods for anomaly detection in online social networks. In *Human-Centric Smart Computing: Proceedings of ICHCSC 2022* (pp. 119-128). Singapore: Springer Nature Singapore.
- [8] Van Belle, R., Baesens, B., & De Weerd, J. (2023). CATCHM: A novel network-based credit card fraud detection method using node representation learning. *Decision Support Systems*, 164, 113866.

- [9] Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., & Bronstein, M. (2020). Temporal graph networks for deep learning on dynamic graphs. arXiv preprint arXiv:2006.10637.
- [10] Rasul, I., Shaboj, S. I., Rafi, M. A., Miah, M. K., Islam, M. R., & Ahmed, A. (2024). Detecting financial fraud in real-time transactions using graph neural networks and anomaly detection. *Journal of Economics, Finance and Accounting Studies*, 6(1), 131-142.
- [11] Lee, J. B., Rossi, R. A., Kim, S., Ahmed, N. K., & Koh, E. (2019). Attention models in graphs: A survey. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 13(6), 1-25.
- [12] Zhang, S., Zhou, Z., Li, D., Zhong, Y., Liu, Q., Yang, W., & Li, S. (2021, May). Attributed heterogeneous graph neural network for malicious domain detection. In *2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design (CSCWD)* (pp. 397-403). IEEE.
- [13] Wang, D., Lin, J., Cui, P., Jia, Q., Wang, Z., Fang, Y., ... & Qi, Y. (2019, November). A semi-supervised graph attentive network for financial fraud detection. In *2019 IEEE international conference on data mining (ICDM)* (pp. 598-607). IEEE.
- [14] Dong, J., Jiang, Z., Pan, D., Chen, Z., Guan, Q., Zhang, H., ... & Gui, W. (2025). A survey on confidence calibration of deep learning-based classification models under class imbalance data. *IEEE Transactions on Neural Networks and Learning Systems*.
- [15] Ata, O., & Hazim, L. (2020). Comparative analysis of different distributions dataset by using data mining techniques on credit card fraud detection. *Tehnički vjesnik*, 27(2), 618-626.
- [16] Ying, Z., Bourgeois, D., You, J., Zitnik, M., & Leskovec, J. (2019). Gnnexplainer: Generating explanations for graph neural networks. *Advances in neural information processing systems*, 32.
- [17] Kulkarni, P. G., Praneet, S. Y., Raghav, R. B., Ashok, A., & Das, B. (2021, April). An Extended Oddball Technique to Detect Anomaly in Static Attributed Graphs. In *Proceedings of 6th International Conference on Recent Trends in Computing: ICRTC 2020* (pp. 625-632). Singapore: Springer Singapore.
- [18] Zhou, R., Zhang, Q., Zhang, P., Niu, L., & Lin, X. (2021). Anomaly detection in dynamic attributed networks. *Neural Computing and Applications*, 33(6), 2125-2136.
- [19] Lamichhane, P. B., & Eberle, W. (2024). Anomaly detection in graph structured data: A survey. arXiv preprint arXiv:2405.06172.
- [20] Xing, S., Wang, Y., & Liu, W. (2025). Multi-Dimensional Anomaly Detection and Fault Localization in Microservice Architectures: A Dual-Channel Deep Learning Approach with Causal Inference for Intelligent Sensing. *Sensors*, 25(11), 3396.
- [21] Xing, S., & Wang, Y. (2025). Cross-Modal Attention Networks for Multi-Modal Anomaly Detection in System Software. *IEEE Open Journal of the Computer Society*.
- [22] Wang, Y., & Xing, S. (2025). AI-Driven CPU Resource Management in Cloud Operating Systems. *Journal of Computer and Communications*, 13(6), 135-149.
- [23] Xing, S., Wang, Y., & Liu, W. (2025). Self-adapting CPU scheduling for mixed database workloads via hierarchical deep reinforcement learning. *Symmetry*, 17(7), 1109.
- [24] Han, X., Yang, Y., Chen, J., Wang, M., & Zhou, M. (2025). Symmetry-Aware Credit Risk Modeling: A Deep Learning Framework Exploiting Financial Data Balance and Invariance. *Symmetry* (20738994), 17(3).

- [25] Cui, Y., Han, X., Chen, J., Zhang, X., Yang, J., & Zhang, X. (2025). FraudGNN-RL: a graph neural network with reinforcement learning for adaptive financial fraud detection. *IEEE Open Journal of the Computer Society*.
- [26] Chen, J., & Fan, H. (2025). Beyond Automation in Tax Compliance Through Artificial Intelligence and Professional Judgment. *Frontiers in Business and Finance*, 2(02), 399-418.
- [27] Cao, J., Zheng, W., Ge, Y., & Wang, J. (2025). DriftShield: Autonomous fraud detection via actor-critic reinforcement learning with dynamic feature reweighting. *IEEE Open Journal of the Computer Society*.
- [28] Zhang, H., Ge, Y., Zhao, X., & Wang, J. (2025). Hierarchical deep reinforcement learning for multi-objective integrated circuit physical layout optimization with congestion-aware reward shaping. *IEEE Access*.
- [29] Zhao, X., Yang, Y., Yang, J., & Chen, J. (2025). Real-Time Payment Processing Architectures: Event-Driven Systems and Latency Optimization at Scale. *Journal of Banking and Financial Dynamics*, 9(12), 10-21.
- [30] Lin, H., Liu, J., Zhang, S., & Zeng, Z. (2025). Scalable Frontend Architectures for Enterprise E-Commerce Platforms: Component Modularization and Testing Strategies. *Asian Business Research Journal*, 10(12), 44-56.
- [31] Liu, J., Wang, J., & Lin, H. (2025). Coordinated Physics-Informed Multi-Agent Reinforcement Learning for Risk-Aware Supply Chain Optimization. *IEEE Access*, 13, 190980-190993.
- [32] Yang, J. S., Shen, Z., Zeng, Z., & Chen, Z. (2025). Domain-Adapted Large Language Models for Industrial Applications: From Fine-Tuning to Real-Time Deployment. *Computer Science Bulletin*, 8(01), 272-289.