

X-MHGNN: An Explainable Multimodal Heterogeneous Graph Neural Network for Low-Latency Fraud Mitigation in Streaming Financial Networks

Shetul B Chandarana¹, Prof. (Dr.) Mayur Jani²

¹PG Scholar, Department of Computer Science and Engineering, School of Engineering and technology, Dr Subhash University, Junagadh, Gujarat

²Assistant Professor, Department of Information Technology, School of Engineering and technology, , Dr Subhash University, Junagadh, Gujarat.

Abstract- Detecting and preventing financial fraud are driven by high-volume transaction streams, interdependent financial entities, massive data imbalance, rapidly changing fraud patterns, and the need for low-latency decisions with robust explanations. This paper provides a comprehensive survey of the recent X-MHGNN literature (20 primary research + 5 review papers) publishing from 2021 to 2026, discussing the strengths and weaknesses of the previous works and identifying future research prospects. We first discuss the importance of the graph structure and GNN for fraud detection. Then we go through several representative graph learning and graph neural network (GNN) methods, multisource and multi-view learning, heterogeneous and temporal graph modelling, explainable artificial intelligence, temporal and streaming fraud detection, contrastive learning, dual-timescale learning, efficient graph transformers, and class-imbalance handling. We then conduct several systematic discussions and propose that the next-generation graph-based fraud detection methods should precisely model the multiple types of entities and temporal transaction relationships within the heterogeneous and temporal graphs, support multimodal evidence, leverage contrastive learning and efficient graph transformers, be able to provide low-latency, stable and explainable decisions in streaming-based environments, effectively handle the high data imbalance, and adapt to evolving fraud patterns. We finally point out the opportunities for the future X-MHGNN research, for example into the effectiveness of multimodality evidence in fraud detection, explainability evaluation metrics, and ablation studies of the different graph-based and explanation-based components. We conclude the paper with a discussion of future participants and directions.

Keywords: Financial fraud detection; Graph Neural Networks; Heterogeneous graphs; Multimodal learning; Explainable AI; Streaming analytics

1. INTRODUCTION

A. Background and Importance

Financial transactions and behaviour produce continuous streams of data. A financial event is therefore more than a row of numerical attributes; it is an interaction among entities. A customer may use several devices, an account may transfer funds to other accounts, multiple users may share a device or network location, and a merchant may interact with many customers. Identification of fraudulent behaviour patterns imbedded in these interactions is inherently challenging. Accordingly, recent developments in financial fraud detection have moved away from conventional machine learning and anomaly detection to graph-based learning [16]–[20].

Graph neural networks (GNNs) are a natural fit for learning from relational structure. Early works in finance

applied GNNs to role-aware relational modelling and to dealing with imbalanced labels [1], [2]. Subsequent research has applied GNNs to heterogeneous graphs, where multiple node and edge types preserve the semantics of financial relationships [4], [8], [12]. Other works applied temporal graph methods to account for the changing nature of financial networks [9], [11], while others explored graph transformer models for higher representation capacity and efficiency [6], [10].

B. Motivation

Fraud detection in finance is a crucial and growing field due to the increasing scale and sophistication of digital transactions. In this literature review, we aim to address the limitations of traditional fraud detection methods, which involve extensive manual feature engineering for each new fraud type and often fail to adapt to evolving fraud behaviours. We propose an approach that

incorporates dynamic heterogeneous graph neural networks (GNNs), which can effectively capture the evolving financial fraud behaviors and provide better fraud detection capabilities. In order to achieve this, we systematically review literature related to machine learning and deep learning-based and explainable graph learning models for financial fraud detection. This research is motivated by the fact that existing studies have only addressed single aspects of the application of GNNs to financial fraud detection, while having very few works that, to the best of our knowledge, report on a comprehensive approach that incorporates dynamic heterogeneous GNNs for fraud detection.

C. Problem Statement

The focus of this review essay is a question: How can dynamic heterogeneous graph neural networks (GNNs) provide a better approach for detecting financial fraud, identify specific challenges in detecting fraud, and examine the limitations of current fraud detection methods? There are different ways for researchers to work towards better fraud detection, for example, in a scenario where they have a single data source, when they are focusing on temporal dynamics, while in another scenario, they may be trying to address explainability or class imbalance. The literature matrix summarises recent research, indicating the areas of focus for each author. The research question is relevant in terms of industry and practical applications, because for a fraud detection system to be effective, it needs to integrate several functionalities, such as identifying the most recent opportunities for fraud, reasoning over many different types of data, describing anomalous behavior, and providing a quick risk assessment.

D. Objectives and Scope

- Review recent financial-fraud research using machine learning, deep learning, GNNs, heterogeneous GNNs, temporal GNNs, graph transformers, and explainable graph learning.
- Our research exercises and contributions are as follows: To classify the literature against the dimensions of graph structure, time modelling, multimodal fusion, imbalance, explainability and deployment efficiency;
- To compare the representative approaches against the dimensions of method, dataset, results, explainability, temporal, and operational aspects; To map the research gaps for the proposed X-MHGNN research trajectory; To formulate a research agenda and a

comparative evaluation framework for low-latency streaming fraud mitigation.




Fig. 1. provides an overview of the conceptual end-to-end framework for the proposed X-MHGNN research trajectory.

Figure 1 illustrates the research strategy. The strategy is conceptualized and is not yet validated experimentally. The paper reviews the literature that provides the evidence for each component, and the research contribution will focus on the integration of the components and their effectiveness for low-latency streaming fraud mitigation.

2.LITERATURE REVIEW

Our review draws primarily from the supplied X-MHGNN literature matrix that covers 2021-2026. The matrix reports 25 records: 15 published research papers, five review papers, and five survey papers. Our review synthesis is thematic, rather than chronological, and is disciplined by the following five review questions: (RQ1) in what manner are GNNs used for financial fraud detection; (RQ2) in what manner have heterogeneous, time-series and multi-modal graph learning developed; (RQ3) in what manner is explainability embedded; (RQ4) in what manner are streaming, scalability, and low-latency supported; and (RQ5) what research gaps are yet to be addressed in the X-MHGNN.

A. Evolution from Conventional Methods to Graph Learning

Conventional FinTech fraud detection includes expert rules, statistical models, supervised ML, anomaly detection, and deep learning. These techniques can still be meaningful in fraud detection because they are

effective for tabular transaction features, and in some instances are flexible in deployment [16]–[18]. Except for the fact that they lack a graph aspect, the transition to graph-based methods happens at the modeling unit level - working on transaction-level feature vectors instead of graph-level representations.

Moving to GNNs requires reinterpretation of the applied level - modeling the fraud alert level (i.e., the account) instead of the transaction. Xu et al. [1] employ role constraints and relational information in consumer-loan fraud, and Liu et al. [2] consider graph fraud in a severe label imbalance setting by proposing neighborhood and subgraph balanced sampling approaches. Motie and Raahemi [19] provide a systematic literature review of graph-based financial fraud detection, and Cheng et al. [20] provide an updated one.

B. Heterogeneous and Multi-View Graph Learning

In heterogeneous graph learning, learning is non-trivial since the network consists of multiple node and relation types. Kanezashi et al. [4] examined homogeneous versus heterogeneous GNNs on Ethereum fraud and found better performance with homogeneous vs heterogeneous variants, particularly R-GCN outperformed other heterogeneous methods. Multiple multi-view or attribute based methods [5], [7], [8], [12] show that different views of financial behavior can carry complementary information.

C. Temporal and Streaming Fraud Detection

Financial graphs are dynamic: transactions arrive over time, new entities enter the system, and fraud patterns evolve. Temporal Graph Networks and event-based temporal graphs therefore provide a more realistic representation than static graphs for streaming settings [9], [11]. Tian and Liu [6] incorporate spatial-temporal information through graph transformers, while recent work on dual-timescale graph learning models both short-term bursts and longer-term behavioral drift [15].

D. Explainable Graph Learning

Explainability is necessary when a fraud score triggers an operational intervention. In the reviewed literature we find heterogeneous fraud models with explanation components [8], interpretable and imbalance-aware graph learning [14], and counterfactual explanation mechanisms [15]. Explainability should be evaluated for

fidelity, stability, sparsity, and practical usefulness rather than treated as only a visualization technique.

E. Low-Latency and Efficient Fraud Detection

FraudGT [10] explicitly targets efficient graph-transformer processing for financial fraud and reports higher F1 together with improved throughput and lower latency than comparison methods. Reinforcement-learning and GNN fusion has also been studied for real-time financial fraud detection [13]. These studies demonstrate that model accuracy alone is insufficient for a streaming system: graph construction, feature extraction, neighborhood retrieval, model inference, and explanation generation all contribute to end-to-end delay.

Author / Year	Method	Dataset / Setting	Key result reported	Main limitation / relevance to X-MHGNN
Xu et al. (2021) [1]	Role-constrained GNN + CRF	Consumer-loan fraud; dataset details should be verified from full paper	Role-aware relational learning; demonstrated usefulness of multiple roles/relations	Does not jointly address multimodal fusion, XAI, streaming and low latency
Liu et al. (2021) [2]	PC-GNN; balanced subgraph + neighborhood sampling	YelpChi, Amazon; financial graph datasets M7/M9	Reported improved fraud-detection performance under class imbalance	Focuses primarily on imbalance; not an integrated temporal multimodal XAI system
Pereira & Murai (2021) [3]	GNN effectiveness analysis	Network fraud benchmarks; see original paper	Examines effectiveness of GNN approaches for network data	Not a unified deployment-oriented framework

Author / Year	Method	Dataset Setting	Key result reported	Main limitation / relevance to X-MHGNN	Author / Year	Method	Dataset Setting	Key result reported	Main limitation / relevance to X-MHGNN
Kan et al. (2022) [4]	R-GCN, HAN, HGT and homogeneous GNN comparison	Ethereum transaction network + phishing labels	Heterogeneous models performed better overall; R-GCN led overall metrics in the study	Cryptocurrency setting; multimodal XAI and streaming not central	Kim et al. (2024) [9]	Temporal Graph Network	Financial graph anomaly-detection benchmarks; verify exact dataset	Models time-dependent graph behavior	Explainability and multimodal fusion remain limited
Xiang et al. (2023) [5]	Attribute-driven graph representation; semi-supervised learning	Credit-card fraud dataset; verify exact split from paper	Uses attributes and graph structure for semi-supervised fraud detection	Limited integration of explanation and low-latency streaming	FraudGT (2024) [10]	Efficient Graph Transformer with edge-aware mechanisms	Multiple large-scale financial datasets; Anti-Money-Laundering data supported by released code	Reported 8–17% higher F1, about 2.4× throughput, and lower latency than comparison methods	Explainability and multimodal streaming adaptation are not the main contribution
Tian & Liu (2023) [6]	Spatial-temporal-aware Graph Transformer	Financial transaction graph; verify exact dataset from paper	Combines spatial and temporal information	Multimodal fusion and XAI remain limited	Saldaña-Ulloa et al. (2024) [11]	Event-based Temporal Graph Network	Online payment-platform transaction data; verify exact dataset details	Models transaction events in temporal graph form	Explanation and multimodal fusion require further integration
Shi & Zhao (2023) [7]	Hierarchical Graph Attention Network	Financial fraud data; verify exact dataset from paper	Integrates local and broader structural information	Does not provide a complete streaming/XAI/multimodal pipeline	Metapath-guided heterogeneous graph learning GN (2025) [12]	Metapath-guided heterogeneous graph learning	Financial fraud dataset; verify exact dataset from article	Uses semantic paths to exploit heterogeneous relations	Deployment latency and integrated XAI need further study
Supply-chain finance study (2024) [8]	Heterogeneous GNN + explanation	Supply-chain finance graph; verify exact dataset from article	Combines heterogeneous graph learning with fraud explanation	Low-latency streaming is not the central evaluation target	RL-GNN (2025) [13]	Reinforcement learning + GNN + commun	Real-time financial fraud setting; verify exact dataset	Targets context-aware real-time fraud detection	Unified multimodal heterogeneous explainability is not

Author / Year	Method	Dataset / Setting	Key result reported	Main limitation / relevance to X-MHGNN
	ity mining			the central focus
Interpretable graph learning (2026) [14]	Interpretable and imbalance-aware graph learning	Financial fraud benchmark; verify exact dataset from article	Integrates interpretability with imbalance handling	Streaming multimodal fusion and end-to-end latency remain open
Zheng (2026) [15]	TransFraudGN: dual-timescale GNN + contrastive learning + counterfactual explanation	IEEE-CIS, Amazon, Yelp	AUC-ROC 0.9487, 0.9361 and 0.9218 respectively; combines temporal modeling and explanation	Mobile-payment application is broader than its benchmark datasets; unified low-latency multimodal deployment remains open

Table 1. Summary of Representative Reviewed Research Papers

From Table 1, we conclude that progress is cumulative: imbalance aware GNNs for minority detection [2]; heterogeneous models that preserve relationship semantics [4]; temporal models to capture evolving interactions [9], [11]; efficient graph transformers to handle operational costs [10]; and lately explainable models that provide interpretable or counterfactual modelling [14], [15]. Existing research problem therefore has mainly to do with integrated systems and deployment validation rather than simply more graph layers.

3. RESEARCH GAP

I have seen GNNs used for financial fraud detection. For detecting fraud in networks especially with role-aware and imbalance-aware methods [1]–[3]. Heterogeneous GNNs help represent different entity types and their relations [4] [8] [12]. Temporal GNNs and event-based graph methods are used to capture how financial interactions change over time [9] [11]. Graph transformers are being explored for fraud detection with a focus on how fast they run and how much latency they add [6] [10]. Explainable graph learning and counterfactual approaches are also research areas [8] [14] [15]. Many reviews and surveys have gathered the progress of graph-based fraud detection and the problems it still faces [19]–[25]. The reviewed studies reveal eight recurring gaps.




Figure 2. Research Gap Analysis

First the study shows that multimodal evidence is not always fused with temporal graph learning. Second heterogeneous relationships and temporal dynamics are often split into parts or only partly combined. Third generating explanations is rarely tested with strict limits on latency. Fourth the costs of building a graph and searching its neighbourhood are not always added to the latency. Fifth, concept drift and the need for adaptation are still hard to solve. Sixth class imbalance and detecting negatives stay core problems. Seventh it is not clear if methods generalise across datasets or platforms. Eighth important operational numbers like p95 or p99 latency, throughput, calibration, explanation fidelity and memory

use are reported often than accuracy or F1 [19] [20] [25]. The main gap can be summed up: the literature has many of the pieces needed for X-MHGNN. There is little proof of one complete system that mixes heterogeneous multimodal data temporal updates, explainable results, heavy imbalance, concept drift and low latency. Our work should be seen as putting these pieces and testing them rather than claiming each part is brand-new.

4. METHODOLOGY

A. Conventional Machine Learning

The weakness of machine learning for fraud that spreads across networks is that relationships between people or accounts must be engineered by hand rather than learned directly from graph structure [16]–[18].

B. Deep Learning and Sequential Models

Deep learning models such as convolutional neural networks (CNNs), recurrent neural networks (RNNs), long short-term memory networks (LSTMs), gated recurrent units (GRUs), and transformers can learn complex or sequential patterns. Sequential models can capture transaction frequency, inter-event intervals, and behavioral changes over time [17], [18]. However, sequence-only representations do not explicitly encode the full set of relationships among people, accounts, devices, and merchants, motivating graph-based representations.

C. Graph Convolutional Networks (GCN)

Graph Convolutional Networks (GCNs) learn node representations by propagating and aggregating information from neighboring nodes. A standard GCN layer can be written as $H^{(l+1)} = \sigma(\tilde{D}^{(-1/2)} \hat{A} \tilde{D}^{(-1/2)} H^{(l)} W^{(l)})$, where $\hat{A}$ is the adjacency matrix with self-loops, $\tilde{D}$ is its degree matrix, $H^{(l)}$ is the node representation at layer l , $W^{(l)}$ is a learnable weight matrix, and σ is an activation function [26].

D. GraphSAGE

GraphSAGE learns node representations by aggregating information from sampled neighborhoods and supports inductive generation of embeddings for previously unseen nodes [27]. This property is useful in evolving

financial networks where new customers, devices, and merchants can continually enter the graph.

E. Graph Attention Networks

Graph Attention Networks (GAT) assign learned attention coefficients to neighboring nodes so that their relative contributions can be differentiated [28]. However, attention coefficients should not automatically be interpreted as causal explanations; dedicated explanation methods may therefore be required when explanations are used for fraud investigation.

F. Relational and Heterogeneous GNNs

Relational Graph Convolutional Networks (R-GCN) model different relation types using relation-specific transformations [29]. Heterogeneous Graph Attention Networks (HAN) use hierarchical attention over node-level and semantic-level information, including meta-path-based neighbors [30]. Heterogeneous Graph Transformers (HGT) use node- and edge-type-dependent parameters and attention to represent heterogeneous graphs, with relative temporal encoding for dynamic settings [31]. These models are appropriate when entities such as customers, accounts, merchants, devices, transactions, and locations have different semantic roles.

G. Temporal GNNs

Temporal GNNs incorporate event time and changing neighborhoods or node states, allowing models to represent evolving interactions more directly than static graphs [9], [11]. Financial fraud studies using temporal graph networks show how event-based transaction graphs can model changing interactions among users, cards, devices, and accounts [11]. The additional state updates can increase memory use and inference cost.

H. Multimodal Learning

Multimodal learning combines complementary evidence such as transaction attributes, behavioral sequences, device information, merchant information, location, text, and graph structure. In financial fraud, heterogeneous and multi-view approaches have combined multiple sources and graph views for detection and explanation [8], while meta-path-based approaches exploit different relational views in heterogeneous transaction networks [12]. An X-MHGNN design should encode each modality appropriately before fusion so that one information source does not dominate the representation.

I. Explainable AI for Graphs

Graph-model explanations can identify influential subgraphs and node features using GNNExplainer [32], while PGExplainer learns a parameterized explanation model for generating explanation masks across multiple instances [33]. In financial fraud detection, explanations should identify the features, entities, links, and temporal events that contribute to a decision and should be sufficiently stable and concise for analyst use.

J. Streaming and Low-Latency Processing

Streaming fraud detection requires continuous event ingestion, incremental graph updates, rapid feature computation, neighborhood retrieval, model inference, and explanation generation. Work on efficient graph transformers and temporal transaction graphs emphasizes that operational performance depends on the complete processing pipeline rather than neural-network inference alone [10], [11], [13].

5. CHALLENGES AND ENHANCEMENTS

A. Current Challenges

Financial-fraud graph research faces several practical and methodological challenges. First, access to transaction-level graph data is restricted by privacy, confidentiality, and the limited availability of representative datasets. Differences in how banks, payment platforms, e-commerce services, and cryptocurrency systems organize entities and transactions also make cross-dataset comparison difficult [8], [19], [20].

Continuous graph construction and low-latency inference are another challenge. Streaming systems must ingest events, update graph structures, retrieve relevant neighborhoods, compute features, run inference, and generate explanations within a short response time. The end-to-end pipeline can therefore become a bottleneck even when the neural model itself is computationally efficient [10], [11], [13].

Severe class imbalance remains important because fraudulent transactions are usually much fewer than legitimate transactions. A model can achieve high overall accuracy while still producing poor fraud recall or an excessive number of false positives. Imbalance-aware graph learning and evaluation with precision, recall, F1-score, and PR-AUC are therefore necessary [2], [5], [19]. Fraud patterns also change over time. New behaviors, entities, and transaction relationships can create concept

drift, reducing the reliability of models trained on older graph structures. Temporal graph methods and time-aware evaluation can help measure performance under evolving conditions [9], [11], [15].

Explainability is another challenge. Graph explanations should identify the relevant features, entities, relationships, and temporal events behind a prediction, while remaining sufficiently stable and understandable for analysts. Explanation methods such as GNNExplainer and PGExplainer provide useful mechanisms, but explanation quality and computational cost must also be evaluated [25], [32], [33].

Graph models may also be vulnerable to adversarial manipulation of nodes, edges, or attributes, and their performance may vary when transferred across datasets or financial platforms. Robustness testing and cross-domain evaluation are therefore required before practical deployment [20], [22], [24].

B. Suggested Enhancements and Experimental Framework

The proposed X-MHGNN framework should be evaluated against representative non-graph, sequential, homogeneous-graph, heterogeneous-graph, and temporal-graph baselines. Suitable baselines include Logistic Regression, Random Forest or XGBoost, LSTM or Transformer models, GCN, GraphSAGE, GAT, a heterogeneous GNN, and a temporal GNN [1], [6], [10], [11].

Ablation experiments should remove or replace one component at a time to measure the contribution of multimodal fusion, heterogeneous relation modelling, temporal modelling, imbalance handling, explainability, and low-latency optimization. Training, validation, and testing should be separated chronologically rather than by random splitting where future information could leak into model development [9], [11].

Model quality should be reported using precision, recall, F1-score, PR-AUC, ROC-AUC, and calibration measures rather than accuracy alone. Operational performance should include throughput, memory consumption, detection delay, and p50/p95/p99 end-to-end latency. Explanation quality, robustness to graph manipulation, and performance under concept drift should also be evaluated [10], [13], [25].

6. CONCLUSION

This review looked at how financial fraud detection has changed. It started with machine learning and anomaly detection. Now it uses heterogeneous temporal, explainable and efficiency-aware graph learning. The literature shows useful abilities. Role-aware. Imbalance-

aware sampling make graph fraud learning better. Heterogeneous graphs keep entity and relation types. Temporal GNNs show how interactions change over time. Graph transformers give representation and faster operation. Recent studies add counterfactual reasoning to financial fraud detection.

The main observation is that these abilities are still split in the literature. This creates a chance to build an integrated X-MHGNN framework. That framework would mix evidence, heterogeneous temporal graph learning, explainability, imbalance handling, concept-drift adaptation and low-latency streaming inference. I think the proposed framework is a research idea until we test it in experiments. To prove its value we need baselines, component ablations, time-aware tests, explanation checks, robustness trials and full operation metrics.

REFERENCES

- [1] B. Xu, H. Shen, B. Sun, R. An, Q. Cao, and X. Cheng, "Towards Consumer Loan Fraud Detection: Graph Neural Networks with Role-Constrained Conditional Random Field," AAAI Conference on Artificial Intelligence, 2021. doi:10.1609/aaai.v35i5.16582.
- [2] Y. Liu, X. Ao, Z. Qin, J. Chi, J. Feng, H. Yang, and Q. He, "Pick and Choose: A GNN-based Imbalanced Learning Approach for Fraud Detection," Proceedings of the Web Conference, pp. 3168–3177, 2021. doi:10.1145/3442381.3449989.
- [3] R. D. R. Pereira and F. Murai, "How effective are Graph Neural Networks in Fraud Detection for Network Data?," arXiv:2105.14568, 2021.
- [4] H. Kanezashi, T. Suzumura, X. Liu, and T. Hirofuchi, "Ethereum Fraud Detection with Heterogeneous Graph Neural Networks," arXiv:2203.12363, 2022.
- [5] S. Xiang, M. Zhu, D. Cheng, E. Li, R. Zhao, Y. Ouyang, L. Chen, and Y. Zheng, "Semi-supervised Credit Card Fraud Detection via Attribute-Driven Graph Representation," AAAI Conference on Artificial Intelligence, 2023.
- [6] Y. Tian and G. Liu, "Transaction Fraud Detection via Spatial-Temporal-Aware Graph Transformer," arXiv:2307.05121, 2023.
- [7] F. Shi and C. Zhao, "Enhancing financial fraud detection with hierarchical graph attention networks: A study on integrating local and extensive structural information," Finance Research Letters, 2023. doi:10.1016/j.frl.2023.104458.
- [8] "Heterogeneous graph neural networks for fraud detection and explanation in supply chain finance," Information Systems, vol. 121, 2024.
- [9] Y. Kim, Y. Lee, M. Choe, S. Oh, and Y. Lee, "Temporal Graph Networks for Graph Anomaly Detection in Financial Networks," arXiv:2404.00060, 2024.
- [10] J. Lin, X. Guo, Y. Zhu, S. Mitchell, E. Altman, and J. Shun, "FraudGT: A Simple, Effective, and Efficient Graph Transformer for Financial Fraud Detection," ICAIF, 2024.
- [11] D. Saldaña-Ulloa, G. De Ita Luna, and J. R. Marcial-Romero, "A Temporal Graph Network Algorithm for Detecting Fraudulent Transactions on Online Payment Platforms," Algorithms, vol. 17, no. 12, 552, 2024.
- [12] "Metapath-guided graph neural networks for financial fraud detection," 2025.
- [13] R. Renuga Devi, J. Emerson Raja, and Y. Boon Chin, "Reinforcement learning with graph neural network (RL-GNN) fusion for real-time financial fraud detection: a context-aware community mining approach," Scientific Reports, 2025.
- [14] "A novel graph learning framework for interpretable and imbalance financial fraud detection," 2026.
- [15] M. Zheng, "Causal-guided dual-timescale graph neural networks for explainable fraud detection in mobile payment systems," Journal of King Saud University Computer and Information Sciences, vol. 38, art. 634, 2026. doi:10.1007/s44443-026-00950-w.
- [16] K. Gubran Al-Hashedi and P. Magalingam, "Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019," Computer Science Review, 2021. doi:10.1016/j.cosrev.2021.100402.
- [17] W. Hilal, S. A. Gadsden, and J. Yawney, "Financial Fraud: A Review of Anomaly Detection Techniques and Recent Advances," Expert Systems with Applications, 2022. doi:10.1016/j.eswa.2021.116429.
- [18] A. Ali et al., "Financial Fraud Detection Based on Machine Learning: A Systematic Literature Review," Applied Sciences, vol. 12, no. 19, 9637, 2022.
- [19] S. Motie and B. Raahemi, "Financial fraud detection using graph neural networks: a systematic review," Expert Systems with Applications, 2024.
- [20] D. Cheng, Y. Zou, S. Xiang, and C. Jiang, "Graph neural networks for financial fraud detection: a review," Frontiers of Computer Science, 2025.

- [21] X. Ma, J. Wu, S. Xue, J. Yang, C. Zhou, Q. Z. Sheng, H. Xiong, and L. Akoglu, "A Comprehensive Survey on Graph Anomaly Detection with Deep Learning," arXiv:2106.07178, 2021.
- [22] H. Kim, B. S. Lee, W.-Y. Shin, and S. Lim, "Graph Anomaly Detection with Graph Neural Networks: Current Status and Challenges," IEEE Access, 2022. doi:10.1109/ACCESS.2022.3211306.
- [23] "Credit card fraud detection in the era of disruptive technologies: A systematic review," 2023.
- [24] "Graph Learning-Empowered Financial Fraud Detection: Progress and Future Directions," 2025.
- [25] Z. S. Khoozani, M. Seera, and C. P. Lim, "Towards Transparent Financial AI: A Systematic Review of Graph Learning and Explainable Methods for Credit Risk and Fraud Detection," 2026.
- [26] T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," in International Conference on Learning Representations (ICLR), 2017.
- [27] W. L. Hamilton, R. Ying, and J. Leskovec, "Inductive Representation Learning on Large Graphs," in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.
- [28] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph Attention Networks," in International Conference on Learning Representations (ICLR), 2018.
- [29] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. van den Berg, I. Titov, and M. Welling, "Modeling Relational Data with Graph Convolutional Networks," in The Semantic Web – 15th International Conference (ESWC 2018), LNCS 10843, pp. 593–607, Springer, 2018. doi:10.1007/978-3-319-93417-4_38.
- [30] X. Wang, H. Ji, C. Shi, B. Wang, P. Cui, P. S. Yu, and Y. Ye, "Heterogeneous Graph Attention Network," in Proceedings of the Web Conference (WWW '19), pp. 2022–2032, 2019. doi:10.1145/3308558.3313562.
- [31] Z. Hu, Y. Dong, K. Wang, and Y. Sun, "Heterogeneous Graph Transformer," in Proceedings of The Web Conference (WWW '20), pp. 2704–2710, 2020. doi:10.1145/3366423.3380027.
- [32] R. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec, "GNNExplainer: Generating Explanations for Graph Neural Networks," in Advances in Neural Information Processing Systems (NeurIPS), vol. 32, pp. 9240–9251, 2019.
- [33] D. Luo, W. Cheng, D. Xu, W. Yu, B. Zong, H. Chen, and X. Zhang, "Parameterized Explainer for Graph Neural Network," in Advances in Neural Information Processing Systems (NeurIPS), vol. 33, pp. 19620–19631, 2020.