International IT Journal of Research (IITJR), ISSN: 3007-6706
Volume 4, Issue 1, January-March 2026
Available online at: https://itjournal.org/index.php

Explainable Graph Convolutional Networks for Real-Time Financial Fraud Detection in High-Throughput Payment Streams

Dr. Tomasz Zieliński

Associate Professor, Institute of Computer Science, Warsaw University of Technology, Poland

ABSTRACT

The rapid growth of digital banking, online transactions, mobile payment applications, and electronic financial services has significantly increased the volume and complexity of financial transactions. Although these developments have improved transaction efficiency and accessibility, they have also created new opportunities for sophisticated financial fraud, including identity theft, account takeover, money laundering, payment manipulation, and coordinated fraudulent activities. Traditional fraud detection systems primarily rely on rule-based mechanisms and conventional machine learning algorithms, which often struggle to identify complex transaction relationships, evolving fraud patterns, and previously unseen fraudulent behaviors in high-throughput payment environments. To address these limitations, this study proposes an Explainable Graph Convolutional Network (X-GCN) framework for real-time financial fraud detection in large-scale payment streams. The proposed framework represents financial transactions as dynamic heterogeneous graphs, where nodes represent customers, accounts, merchants, devices, and transactions, while edges capture their relationships and interaction patterns. Graph convolutional networks are employed to learn complex structural dependencies and identify suspicious transaction behaviors through multi-layer neighborhood aggregation.

To improve transparency and interpretability, explainable artificial intelligence techniques, including attention-based explanations, feature attribution, and subgraph analysis, are integrated into the detection architecture. The framework incorporates continuous graph updates, temporal transaction embeddings, adaptive thresholding, and stream-processing mechanisms to support low-latency fraud identification under high transaction volumes. Its performance can be evaluated using precision, recall, F1-score, area under the receiver operating characteristic curve (AUC-ROC), false-positive rate, and transaction processing latency. The proposed approach aims to achieve a balance between detection accuracy, computational efficiency, scalability, and model interpretability. Furthermore, explainable fraud predictions can assist financial institutions, compliance teams, and risk analysts in understanding suspicious transaction patterns and making informed decisions. The framework provides a promising foundation for developing transparent, adaptive, and scalable fraud detection systems capable of addressing emerging financial threats in modern digital payment ecosystems.

Keywords: Explainable Artificial Intelligence, Graph Convolutional Networks, Financial Fraud Detection, Real-Time Payment Processing, Graph Neural Networks.

INTRODUCTION

The rapid advancement of digital technologies has transformed the global financial ecosystem by enabling faster, more accessible, and increasingly interconnected payment systems. Online banking, mobile payment applications, digital wallets, contactless transactions, real-time gross settlement systems, and electronic commerce platforms have significantly increased the volume of financial transactions conducted worldwide. These developments have improved financial inclusion, operational efficiency, and customer convenience. However, the expansion of digital payment infrastructures has also introduced significant security challenges, particularly the increasing sophistication, frequency, and complexity of financial fraud. Fraudulent activities such as credit card fraud, account takeover, identity theft, transaction laundering, phishing-based payments, synthetic identity fraud, and coordinated merchant fraud pose substantial risks to financial institutions, businesses, and consumers. Consequently, developing intelligent, accurate, and scalable fraud detection mechanisms has become an essential requirement for maintaining trust and security in modern financial systems.

Traditional financial fraud detection systems primarily depend on predefined rules, statistical techniques, and conventional machine learning algorithms. Rule-based approaches identify suspicious transactions through predetermined conditions, such as unusual transaction amounts, abnormal spending frequencies, geographical inconsistencies, and repeated payment attempts. Although these techniques are relatively simple to implement, they frequently struggle to recognize previously unseen fraud patterns and rapidly evolving fraudulent strategies. Similarly,

conventional machine learning approaches, including decision trees, random forests, support vector machines, and gradient boosting algorithms, can identify suspicious transaction characteristics by learning from historical data. However, these models generally emphasize individual transaction attributes and may overlook complex relationships among customers, financial accounts, merchants, devices, and transaction networks. Furthermore, increasing transaction volumes create additional computational challenges, particularly when financial institutions must evaluate thousands or millions of payment events within extremely short processing intervals.

Graph-based machine learning has emerged as a promising approach for addressing these limitations because financial transactions naturally form interconnected networks. In a financial transaction graph, customers, bank accounts, merchants, devices, and payment events can be represented as nodes, while relationships among these entities are represented as edges. Such representations enable fraud detection systems to investigate not only individual transactions but also the broader behavioral and relational structures associated with financial activities. For example, multiple fraudulent accounts may share common devices, payment addresses, merchants, or transaction destinations. Identifying these relationships can reveal coordinated fraud activities that may remain undetected through conventional transaction-level analysis.

Proposed Models and Methodologies

The proposed research introduces an **Explainable Graph Convolutional Network (X-GCN) framework for Real-Time Financial Fraud Detection in High-Throughput Payment Streams**. The primary objective is to develop an intelligent fraud detection architecture that combines graph-based representation learning, temporal transaction analysis, explainable artificial intelligence, and high-speed stream processing. The framework is designed to identify suspicious financial activities by examining individual transaction characteristics as well as complex relationships among customers, accounts, merchants, devices, and payment networks.

The proposed methodology consists of multiple interconnected stages, including financial transaction data collection, preprocessing, dynamic graph construction, graph convolutional learning, temporal behavior analysis, fraud classification, explainability generation, and real-time performance evaluation.

2.1 Overall Architecture of the Proposed Framework

The proposed X-GCN architecture consists of six major components:

1. Financial transaction data acquisition and preprocessing.
2. Dynamic heterogeneous financial graph construction.
3. Graph convolutional network-based feature extraction.
4. Temporal transaction behavior modeling.
5. Explainable fraud classification and prediction.
6. Real-time stream processing and performance optimization.

The architecture processes incoming financial transactions continuously and generates fraud probability scores. Transactions exceeding a predefined or adaptively learned risk threshold are identified for further investigation.
Graph learning for Financial Applications
aoxaustin.github.io

2.2 Financial Transaction Data Collection and Preprocessing

The first stage involves collecting financial transaction information from digital payment platforms, banking systems, credit card networks, and electronic commerce applications. The dataset contains transaction-level attributes and entity-related information necessary for identifying suspicious financial activities.

Important transaction features include:

- Transaction amount and currency.
- Transaction timestamp and frequency.
- Sender and receiver account identifiers.
- Merchant category and merchant identifier.
- Device and IP-related characteristics.
- Geographic transaction location.
- Payment channel and transaction type.
- Historical transaction behavior.
- Previous fraud labels and risk indicators.

The collected data undergoes preprocessing to improve data quality and model reliability. Missing values are handled through suitable statistical or categorical imputation techniques. Numerical features are normalized, categorical variables are encoded, and duplicate transaction records are removed. Additionally, timestamps are converted into

appropriate temporal representations to identify unusual transaction frequencies, sudden spending changes, and abnormal transaction sequences.

Feature normalization can be represented as:

$$X'=\frac{X-X_{\min}}{X_{\max}-X_{\min}}$$

Where:

- X' represents the normalized feature.
- X represents the original feature value.
- $X_{\min}$ represents the minimum observed value.
- $X_{\max}$ represents the maximum observed value.

This preprocessing stage establishes a consistent feature representation for subsequent graph-based learning.

2.3 Dynamic Heterogeneous Financial Graph Construction

Financial transactions naturally involve multiple interconnected entities. Therefore, the proposed framework represents financial activities using a dynamic heterogeneous graph.

The graph is defined as:

$$G_t=(V_t,E_t,X_t,R)$$

Where:

- V_t represents the set of graph nodes at time t.
- E_t represents the relationships between nodes.
- X_t represents node feature information.
- R represents different relationship types.

The graph contains several node categories:

Node Type	Description
Customer	Individuals initiating financial transactions
Account	Bank accounts and digital payment accounts
Merchant	Businesses receiving payments
Device	Mobile phones, computers, and payment devices
Transaction	Individual financial payment events
Location	Geographic locations associated with transactions

Edges represent financial interactions, including account ownership, payment transfers, device associations, merchant interactions, and transaction relationships.

Announcing Amazon Neptune ML: Easy, fast, and accurate predictions on graphs | AWS Database Blog
aws.amazon.com

A dynamic graph construction mechanism continuously incorporates newly arriving transaction events. This allows the framework to capture changing financial relationships and emerging fraud networks rather than relying exclusively on static historical transaction structures.

RESULTS & ANALYSIS

The proposed Explainable Graph Convolutional Network (X-GCN) framework is designed to improve financial fraud detection in high-throughput payment streams by combining graph-based feature learning, temporal transaction analysis, and explainable artificial intelligence. The evaluation focuses on the model's ability to identify fraudulent transactions, recognize complex financial relationships, minimize false-positive predictions, and maintain low processing latency under continuous transaction workloads.

The experimental evaluation considers important performance indicators, including accuracy, precision, recall, F1-score, AUC-ROC, false-positive rate, transaction processing latency, and throughput. Since no actual experimental dataset or execution results have been provided, the numerical values below are **illustrative hypothetical results** intended to demonstrate how the proposed framework could be evaluated. They should not be presented as experimentally verified findings.

3.1 Experimental Evaluation Environment

The proposed framework can be evaluated using historical financial transaction datasets containing both legitimate and fraudulent payment records. The dataset should include transaction amounts, timestamps, account identifiers, merchant information, device characteristics, and historical fraud labels.
The experimental environment can incorporate the following components:

Component	Proposed Configuration
Programming Language	Python
Deep Learning Framework	PyTorch
Graph Learning Framework	PyTorch Geometric
Streaming Platform	Apache Kafka
Stream Processing	Apache Flink
Graph Architecture	Dynamic Heterogeneous GCN
Temporal Model	Gated Recurrent Unit (GRU)
Explainability	SHAP, Integrated Gradients, Subgraph Explanations
Optimizer	Adam
Loss Function	Weighted Binary Cross-Entropy
Evaluation Strategy	Chronological Train-Test Split

The experimental design emphasizes realistic transaction processing conditions and evaluates both predictive performance and computational efficiency.

3.2 Fraud Detection Performance Analysis

The first experiment examines the ability of the proposed X-GCN architecture to distinguish between legitimate and fraudulent financial transactions.

The framework combines transaction-level features with graph-based structural information and temporal behavioral representations. This allows the model to recognize suspicious transaction patterns that may not be identified through conventional classification techniques.

Illustrative fraud detection performance

Performance Metric	Illustrative Result
Accuracy	98.42%
Precision	96.85%
Recall	97.20%
F1-Score	97.02%
AUC-ROC	99.10%
False Positive Rate	0.85%

The illustrative results indicate that the proposed architecture could achieve a favorable balance between fraud identification and legitimate transaction recognition.

The precision value suggests that a high proportion of transactions classified as fraudulent would be genuinely suspicious. Similarly, the recall value demonstrates the model's intended ability to identify a substantial proportion of actual fraudulent activities.

The F1-score provides an overall assessment of the balance between precision and recall, which is particularly important in financial fraud detection because both missed fraud and unnecessary transaction blocking can create significant operational consequences.

Comparative Analysis in Tabular Form

The comparative analysis evaluates the proposed **Explainable Graph Convolutional Network (X-GCN)** framework against conventional machine learning, deep learning, and graph-based fraud detection approaches. The comparison

focuses on fraud detection accuracy, precision, recall, F1-score, transaction processing latency, explainability, scalability, and adaptability to evolving financial fraud patterns.

4.1 Comparative Performance of Different Fraud Detection Models

Detection Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Logistic Regression	94.10	89.50	82.40	85.80
Random Forest	96.20	92.80	89.60	91.17
XGBoost	97.10	94.20	91.80	92.98
Conventional GCN	97.65	95.10	94.20	94.65
Temporal GCN	98.05	95.90	95.80	95.85
Proposed X-GCN	**98.42**	**96.85**	**97.20**	**97.02**

Note: *All numerical values are illustrative hypothetical results, not experimentally validated measurements.*
The comparison indicates that the proposed X-GCN architecture has the potential to achieve better classification performance by combining graph relationships, temporal transaction behavior, and explainable learning mechanisms.

4.2 Comparison of Architectural Features

Parameters	Traditional ML	Deep Learning	Conventional GCN	Proposed X-GCN
Transaction Analysis	Feature-based	Feature-based	Graph-based	Graph and temporal-based
Relationship Learning	Limited	Limited	Strong	Strong
Temporal Analysis	Limited	Model-dependent	Limited	Integrated
Fraud Pattern Recognition	Moderate	High	High	High
Explainability	Moderate	Low	Low	High
Dynamic Graph Support	No	No	Model-dependent	Integrated
Real-Time Processing	High potential	Model-dependent	Computationally demanding	Optimized design
Scalability	High potential	High potential	Graph-dependent	Distributed architecture
Complex Fraud Detection	Moderate	High	High	High

4.3 Comparative Analysis of Processing Efficiency

Model	Average Latency (ms)	Computational Complexity	Real-Time Suitability
Logistic Regression	4.5	Low	High
Random Forest	8.2	Moderate	High
XGBoost	6.8	Moderate	High
Conventional GCN	15.6	High	Moderate
Temporal GCN	18.4	High	Moderate
Proposed X-GCN	**12.1**	**Optimized graph computation required**	**High potential**

Note: *Latency values are illustrative and depend on hardware, graph size, batch configuration, and implementation.*

The proposed framework aims to balance graph-based detection capabilities with computational efficiency. Incremental graph updates, neighborhood sampling, and distributed processing are important for reducing processing overhead.

LIMITATIONS & DRAWBACKS

Although the proposed **Explainable Graph Convolutional Network (X-GCN) framework for Real-Time Financial Fraud Detection in High-Throughput Payment Streams** offers significant potential for improving financial security, several technical, computational, operational, and practical limitations must be considered. The integration of graph convolutional networks, temporal transaction modeling, explainable artificial intelligence, and real-time stream processing introduces challenges that may affect system performance, scalability, reliability, and practical deployment. The major limitations and drawbacks are discussed below.

5.1 High Computational Complexity

Graph convolutional networks require substantial computational resources to process large-scale financial transaction graphs. As the number of customers, accounts, merchants, devices, and transaction relationships increases, graph construction and neighborhood aggregation become increasingly demanding.

Processing millions of interconnected financial entities may require high-performance GPUs, distributed computing infrastructure, and optimized graph processing mechanisms. These requirements can increase implementation costs and operational complexity.

5.2 High Memory Requirements
The proposed framework maintains dynamic graph structures containing transaction relationships, node features, temporal embeddings, and historical information. Storing and updating these components requires considerable memory resources.

In high-throughput payment environments, continuous graph expansion may create memory management challenges, particularly when historical transaction relationships must be retained for long-term fraud analysis.

5.3 Real-Time Processing Latency
Although the proposed architecture incorporates stream-processing mechanisms, graph updates, temporal feature extraction, and explainability generation may introduce additional processing delays.

Financial payment systems frequently require extremely fast transaction authorization. Therefore, even small increases in prediction latency can affect customer experience and operational efficiency.

5.4 Data Imbalance Problems
Financial fraud datasets generally contain significantly more legitimate transactions than fraudulent transactions. This imbalance can create difficulties during model training because the classification algorithm may become biased toward the majority class.

Although weighted loss functions, resampling techniques, and adaptive thresholds can reduce this problem, detecting rare and previously unseen fraudulent activities remains challenging.

5.5 Data Privacy and Security Concerns
The proposed framework processes sensitive financial information, including customer identities, account relationships, transaction histories, device information, and payment behavior.

Integrating these data sources into a centralized or distributed graph structure creates potential privacy and cybersecurity risks. Unauthorized access, data leakage, and misuse of financial information may negatively affect customer trust and regulatory compliance.

CONCLUSION

The study, **"Explainable Graph Convolutional Networks for Real-Time Financial Fraud Detection in High-Throughput Payment Streams,"** presents an intelligent and scalable approach to addressing the growing challenges of financial fraud in modern digital payment ecosystems. The rapid expansion of online banking, digital wallets, electronic commerce, and real-time payment infrastructures has increased the complexity of fraudulent activities, creating a demand for advanced detection mechanisms capable of identifying suspicious transactions accurately and efficiently. Traditional rule-based and conventional machine learning approaches often face limitations in recognizing complex financial relationships, coordinated fraud networks, and continuously evolving fraudulent behaviors.

To overcome these challenges, the proposed Explainable Graph Convolutional Network (X-GCN) framework integrates dynamic heterogeneous graph construction, graph convolutional learning, temporal transaction analysis, and explainable artificial intelligence techniques. By representing customers, bank accounts, merchants, devices, and transactions as interconnected graph entities, the framework enables comprehensive analysis of financial relationships rather than relying exclusively on individual transaction characteristics. The integration of temporal modeling further supports the identification of unusual transaction sequences, abnormal payment frequencies, and emerging fraudulent patterns.

An important contribution of the proposed framework is its explainability mechanism, which incorporates feature attribution, attention-based relationship analysis, and subgraph explanations. These techniques can help financial investigators understand the factors contributing to suspicious transaction predictions and support transparent, accountable, and informed decision-making. Furthermore, the integration of real-time stream-processing technologies,

incremental graph updates, and optimized neighborhood sampling provides a foundation for processing large volumes of continuously arriving financial transactions.

The comparative evaluation methodology highlights the potential advantages of the proposed architecture over conventional machine learning and graph-based approaches in terms of fraud identification, temporal analysis, interpretability, and operational efficiency. The illustrative performance results indicate the potential to achieve high classification accuracy while maintaining relatively low prediction latency. However, these numerical results are hypothetical and require validation through reproducible experiments using real-world or appropriately anonymized financial transaction datasets.

REFERENCES

1. Akoglu, L., Tong, H., & Koutra, D. (2015). Graph-based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery, 29*, 626–688. https://doi.org/10.1007/s10618-014-0365-y
2. Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science, 17*(3), 235–255. https://doi.org/10.1214/ss/1042727940
3. Carcillo, F., Dal Pozzolo, A., Le Borgne, Y.-A., Caelen, O., Mazzer, Y., & Bontempi, G. (2018). SCARFF: A scalable framework for streaming credit card fraud detection with Spark. *Information Fusion, 41*, 182–194. https://doi.org/10.1016/j.inffus.2017.09.005
4. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys, 41*(3), Article 15. https://doi.org/10.1145/1541880.1541882
5. Cheng, D., Zou, Y., Xiang, S., & Jiang, C. (2025). Graph neural networks for financial fraud detection: A review. *Frontiers of Computer Science, 19*(9), Article 199609. https://doi.org/10.1007/s11704-024-40474-y
6. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems, 29*(8), 3784–3797. https://doi.org/10.1109/TNNLS.2017.2736643
7. Hamilton, W. L., Ying, R., & Leskovec, J. (2017). Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc.
8. Huang, Q., Yamada, M., Tian, Y., Singh, D., & Chang, Y. (2022). GraphLIME: Local interpretable model explanations for graph neural networks. *IEEE Transactions on Knowledge and Data Engineering, 35*(7), 6968–6972. https://doi.org/10.1109/TKDE.2022.3187455
9. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications, 100*, 234–245. https://doi.org/10.1016/j.eswa.2018.01.037
10. Kipf, T. N., & Welling, M. (2017). Semi-supervised classification with graph convolutional networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
11. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
12. Lin, J., Guo, X., Zhu, Y., Mitchell, S., Altman, E., & Shun, J. (2024). FraudGT: A simple, effective, and efficient graph transformer for financial fraud detection. In *Proceedings of the 5th ACM International Conference on AI in Finance*. https://doi.org/10.1145/3677052.3698648
13. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems* (Vol. 30). Curran Associates, Inc.
14. Motie, S., & Raahemi, B. (2024). Financial fraud detection using graph neural networks: A systematic review. *Expert Systems with Applications, 240*, Article 122156. https://doi.org/10.1016/j.eswa.2023.122156
15. Pang, G., Shen, C., Cao, L., & van den Hengel, A. (2021). Deep learning for anomaly detection: A review. *ACM Computing Surveys, 54*(2), Article 38. https://doi.org/10.1145/3439950
16. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). ACM. https://doi.org/10.1145/2939672.2939778
17. Rossi, E., Chamberlain, B., Frasca, F., Eynard, D., Monti, F., & Bronstein, M. (2020). Temporal graph networks for deep learning on dynamic graphs. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*.
18. Shams Khoozani, Z., Seera, M., & Lim, C. P. (2026). Towards transparent financial AI: A systematic review of graph learning and explainable methods for credit risk and fraud detection. *Artificial Intelligence Review, 59*, Article 246. https://doi.org/10.1007/s10462-026-11645-z
19. Sundararajan, M., Taly, A., & Yan, Q. (2017). Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning* (pp. 3319–3328). PMLR.
20. Van Vlasselaer, V., Bravo, C., Caelen, O., Eliassi-Rad, T., Akoglu, L., Snoeck, M., & Baesens, B. (2015). APATE: A novel approach for automated credit card transaction fraud detection using network-based extensions. *Decision Support Systems, 75*, 38–48.