




Article **Open Access**

# Fraud Detection and Risk Assessment of Online Payment Transactions on E-Commerce Platforms Based on LLM and GCN Frameworks

Ruihan Luo <sup>1,\*</sup>, Nanxi Wang <sup>2</sup>, Xiaotong Zhu <sup>3</sup>



- <sup>1</sup> Southwest University of Finance and Economics, Chengdu, China  
<sup>2</sup> USC Viterbi School of Engineering, University of Southern California, Los Angeles, USA  
<sup>3</sup> Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA  
\* Correspondence: RuiHan Luo, Southwest University of Finance and Economics, Chengdu, China

Received: 16 October 2025  
Revised: 28 October 2025  
Accepted: 08 November 2025  
Published: 01 December 2025



**Copyright:** © 2025 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

**Abstract:** With the rapid expansion of e-commerce, online payment fraud has become increasingly sophisticated, posing significant challenges to financial security and undermining consumer trust. Traditional detection methods often struggle to capture the intricate relational and behavioral patterns inherent in transactional data, limiting their effectiveness in highly imbalanced scenarios. This study introduces a novel fraud detection framework that integrates Large Language Models (LLMs) with Graph Convolutional Networks (GCNs) to enhance the identification of fraudulent activities in online payment systems. A dataset comprising 2,840,000 transactions was collected over a 14-day period from major e-commerce platforms, including Amazon, involving approximately 2,000 U.S.-based consumers and 30 merchants. Among these transactions, fewer than 6,000 were fraudulent, presenting a highly skewed class distribution that reflects real-world conditions. In this framework, consumers and merchants were represented as nodes, and transactions as edges, forming a heterogeneous graph that captures both direct and indirect interactions. A GCN was applied to this graph to learn complex structural and relational patterns, effectively modeling the interdependencies among participants. In parallel, semantic features were extracted from transaction metadata using GPT-4o and Tabformer, capturing textual and categorical information that complements the structural graph features. By fusing these semantic representations with graph-based features, the model can identify subtle and context-dependent indicators of fraudulent behavior that traditional methods may overlook. Experimental results demonstrate that the proposed framework achieves an overall accuracy of 0.98, with a balanced performance in terms of precision and sensitivity, highlighting its robustness in detecting rare fraudulent instances within a massive dataset. The integration of LLM-derived semantic features with graph structural information significantly improves detection efficacy compared with models relying solely on either approach. This hybrid approach offers a scalable, real-time solution for enhancing the security of online payment environments and provides a promising avenue for applying graph-based deep learning techniques in financial fraud prevention.

**Keywords:** fraud detection; GPT-4o; GCN; LLM; unbalanced data

## 1. Introduction

In today's digital economy, e-commerce platforms have become integral to consumers' daily shopping and payment activities, with online payment volumes experiencing rapid and sustained growth. This expansion, however, has simultaneously

escalated the prevalence and sophistication of fraudulent activities, posing significant risks to transaction security and undermining consumer confidence. Among these, credit card fraud has emerged as a major concern, with perpetrators exploiting a wide range of channels and strategies to bypass conventional security measures. Unlike traditional financial fraud, e-commerce fraud is often more covert, spanning multiple platforms and manifesting in both first-time transactions and frequent low-value payments. Such patterns make it particularly challenging for rule-based systems or shallow machine learning models to detect fraudulent behavior effectively, as these approaches typically fail to capture complex, high-dimensional interactions among users, merchants, and transactions.

Recent advances in machine learning offer promising avenues to address these challenges. Graph Convolutional Networks (GCNs) have demonstrated strong capabilities in modeling relational and structural dependencies in graph-structured data, enabling the identification of anomalous patterns that may indicate fraud. Concurrently, Large Language Models (LLMs), such as GPT-4o, provide robust mechanisms for extracting semantic information from unstructured or semi-structured transactional data, capturing subtle cues that traditional feature engineering might overlook. The combination of these two approaches-graph-based structural learning and LLM-driven semantic representation-offers a powerful framework for detecting complex fraudulent behavior in dynamic e-commerce environments.

In this study, we propose a novel fraud detection framework that integrates GCNs with LLM-based feature encoding to identify fraudulent transactions in online payments. A heterogeneous graph is constructed, with consumers and merchants represented as nodes and transactions as edges, allowing the model to capture implicit behavioral patterns and risk propagation across the network. To further enrich feature representations, both textual and structured attributes-including merchant category codes, transaction amounts, temporal information, and card usage behaviors-are encoded using GPT-4o and Tabformer.

Experiments are conducted on a large-scale dataset consisting of 2.84 million transactions collected over a 14-day period from major e-commerce platforms such as Amazon, involving approximately 2,000 globally active consumers and 30 merchants. Among these transactions, fewer than 6,000 are labeled as fraudulent, representing a highly imbalanced scenario that mirrors real-world conditions. Despite this imbalance, the proposed framework achieves an accuracy of 0.98, demonstrating both robustness and scalability. As shown in Table 1, the integration of semantic and structural features substantially improves the detection of subtle fraudulent patterns, highlighting the model's potential for real-time application in securing online payment systems.

**Table 1.** GCN model detailed result index.

category	Precision	Recall	F1-Score	Support
0 (Non-Fraud)	0.98	1.00	0.99	1239159
1 (Fraud)	1.00	0.05	0.09	33365
Accuracy			0.98	1272524
Macro Avg	0.99	0.52	0.54	1272524
Weighted Avg	0.98	0.98	0.96	1272524

## 2. Literature Review

The rapid expansion of e-commerce has led to a dramatic increase in the volume of online payment transactions, while simultaneously rendering fraudulent activities more concealed, sophisticated, and difficult to detect. These evolving fraud patterns pose serious threats to the stability of financial systems and the security of consumer payments. Traditional detection methods, including rule-based systems and shallow machine learning models, are increasingly inadequate when confronted with large-scale, high-

dimensional, and highly interdependent transaction data. As a result, deep learning and large model technologies have emerged as dominant approaches for financial fraud detection in both academic research and industry applications.

Early research on machine learning-based fraud detection has provided a foundation for understanding algorithmic effectiveness and limitations. For instance, systematic reviews have shown that Support Vector Machines (SVMs) and Artificial Neural Networks (ANNs) are among the most widely applied techniques, with credit card fraud being the most frequently investigated scenario [1]. These studies also highlight critical limitations in generalization, adaptability to novel fraud patterns, and scalability to massive datasets, motivating the exploration of more advanced models.

Building upon conventional machine learning approaches, recurrent neural networks, particularly Long Short-Term Memory (LSTM) architectures, have been proposed to address challenges such as rapid processing of large transaction volumes and the detection of previously unseen attack patterns [2]. Empirical results indicate that LSTM-based models can achieve high accuracy within very short processing times, outperforming auto-encoders and traditional ML algorithms. Similarly, convolutional neural networks (CNNs) and deep convolutional architectures have demonstrated strong performance in extracting hierarchical features from transaction data, effectively identifying anomalous behaviors even in complex, large-scale datasets [3,4].

In addition to temporal and feature-based learning, recent studies have increasingly focused on graph-based approaches to model relational dependencies between entities in transaction networks. Hybrid models integrating Graph Neural Networks (GNNs), CNNs, and LSTMs have been proposed to capture both node-level and network-level anomalies, dynamically adjusting to evolving fraud patterns through techniques such as reinforcement learning [5]. Graph-based models have also been applied to financial transaction networks by representing consumers and merchants as nodes and transactions as edges, enabling the exploitation of structural and relational information that traditional feature-based models often overlook [6]. These approaches have shown superior performance in identifying complex fraud scenarios, particularly when combined with methods capable of handling high-dimensional and heterogeneous data.

Despite these advances, existing research often focuses on either deep learning for temporal or semantic feature extraction or graph-based structural modeling, with few studies effectively integrating both perspectives. This gap underscores the potential of hybrid frameworks that leverage the strengths of both LLMs for semantic understanding and GCNs for relational pattern recognition, particularly in the context of large-scale, highly imbalanced e-commerce transaction datasets. Such integrated approaches offer the promise of improved accuracy, adaptability, and real-time applicability in detecting sophisticated online payment fraud.

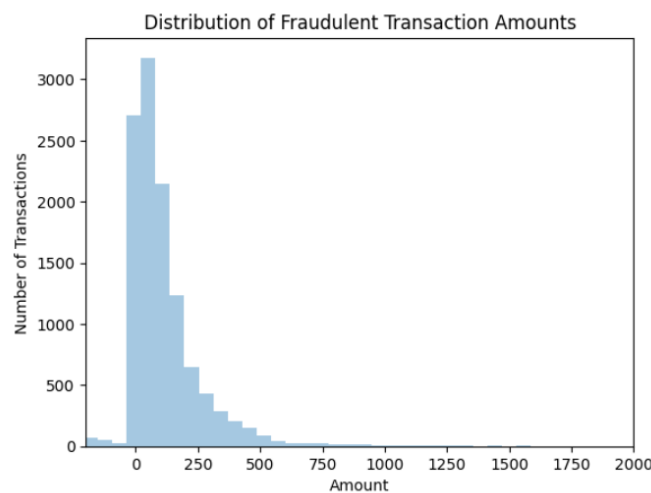
### 3. Data Introduction

The dataset utilized in this study was collected from major e-commerce platforms, including Amazon, and comprises 2,840,000 real online payment transactions recorded over a 14-day period. It involves approximately 2,000 U.S.-based yet globally active consumers and 30 merchants. Among these transactions, fewer than 6,000 are labeled as fraudulent, reflecting a highly imbalanced classification scenario that mirrors real-world e-commerce environments. Each transaction record contains detailed information, including transaction amount, merchant category code (MCC), timestamp, multi-card usage, transaction type, and fraud label, providing a comprehensive representation of consumer behavior, merchant activity, and potential fraud risks.

To improve data quality and enhance feature expressiveness, a two-stage preprocessing strategy was applied. First, GPT-4o was employed to semantically parse and normalize unstructured fields, enriching the contextual understanding of transaction records. Subsequently, Tabformer, a Transformer-based model for tabular data

representation, was used for encoding and data cleaning. This step preserves row-column structures and logical dependencies while generating discriminative feature vectors suitable for downstream analysis. The combination of semantic enrichment and structural consistency ensures a robust foundation for Graph Convolutional Network (GCN)-based learning and fraud detection [7].

As shown in Figure 1, the histogram of fraudulent transaction amounts reveals that most fraudulent activities are concentrated in lower-value transactions. Specifically, a substantial majority of fraud cases occur below \$250, with the highest frequency observed in the \$0-\$50 range. As transaction amounts exceed \$250, the occurrence of fraud declines sharply. High-value transactions over \$1000 are relatively rare, accounting for less than 5% of fraudulent activities. This right-skewed distribution indicates that while small-value transactions are the primary targets-likely due to their lower risk of triggering alerts-large-value transactions, though infrequent, can result in significant financial losses. Such a pattern underscores the importance of designing fraud detection models that are sensitive to both frequent low-value transactions and less frequent high-impact transactions.



**Figure 1.** Histogram of Fraudulent Transaction Amounts in Dataset.

As shown in Figure 2, the distribution of fraudulent transactions across states for the 30 merchants demonstrates substantial geographic imbalance. A few states account for a large proportion of fraud cases, with the highest number of fraudulent transactions reaching up to 2,000. In contrast, several states report fewer than 500 fraudulent transactions. This uneven regional distribution suggests that certain areas may be more susceptible to fraud, or that these regions experience higher exposure to fraudulent activity. Understanding these spatial patterns is critical for developing targeted prevention strategies, enabling platforms and merchants to allocate resources more effectively and mitigate fraud risks in high-incidence regions.

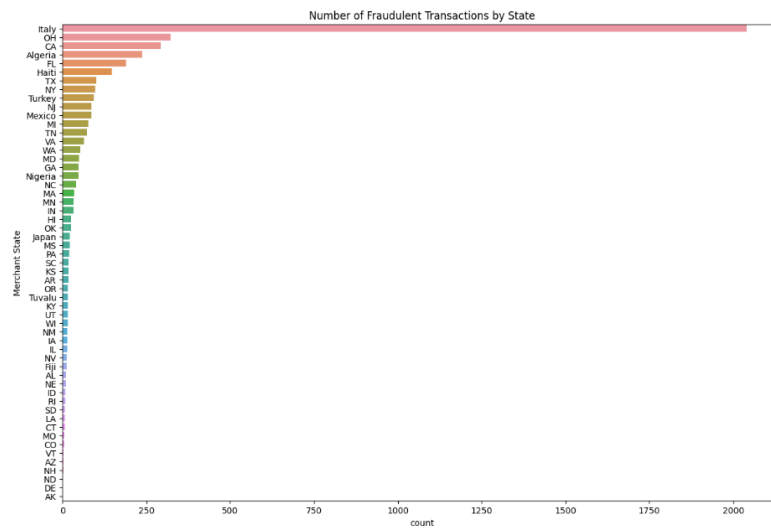


Figure 2. Monthly Fraud Transactions.

#### 4. Model Introduction

This study proposes a fraud detection framework specifically designed for e-commerce online payment scenarios, integrating Large Language Models (LLMs) and Graph Convolutional Networks (GCNs). The approach represents user-merchant interactions as a transaction graph and combines structural and semantic features to enhance the model's ability to detect complex and subtle fraud patterns. The following sections describe the key components of the framework in detail.

##### 4.1. Graph Construction and GCN Modeling

The transaction data is modeled as a heterogeneous graph, where consumers and merchants are represented as two distinct types of nodes, and transactions between them form the edges. Each edge is enriched with multiple attributes, including timestamp, transaction amount, merchant category code (MCC), card identifier, and online transaction type. This graph structure captures both entity relationships and the complex behavioral patterns spanning users and merchants.

To detect fraud on this graph, a two-layer Graph Convolutional Network (GCN) is employed. GCNs are designed to learn from graph-structured data by aggregating information from a node's neighbors. In this framework, the first layer aggregates information from immediate neighbors to generate localized node representations, while the second layer captures higher-order dependencies and global behavior patterns. This hierarchical aggregation enables the model to identify subtle fraudulent activities that may not be apparent from individual transactions alone [8].

To address the severe class imbalance—where fraudulent transactions constitute a very small fraction of the total dataset—a class-weighted loss function is applied. This mechanism ensures that the minority class receives more emphasis during training, enhancing the model's sensitivity to fraud.

##### 4.2. Feature Representation and Fusion

Effective feature representation is critical for graph-based models. We employ a dual-feature encoding strategy that leverages GPT-4o and Tabformer to generate comprehensive node and edge representations.

GPT-4o processes unstructured or semi-structured textual fields in transaction records, such as product category descriptions, merchant notes, or user behavior tags. These fields often contain valuable contextual information that is difficult to capture with conventional encoding methods. GPT-4o generates semantic embeddings that reflect the

contextual meaning of these fields, which are then normalized and incorporated into the model as part of the node and edge features.

Tabformer is used to encode structured tabular data, particularly fields exhibiting logical dependencies or multi-dimensional interactions. As a Transformer-based framework tailored for tabular data, Tabformer preserves both row-column structures and contextual relationships across fields. It captures the interplay among attributes such as transaction time, amount, user identity, and payment type, producing feature embeddings that reflect both numerical patterns and behavioral logic.

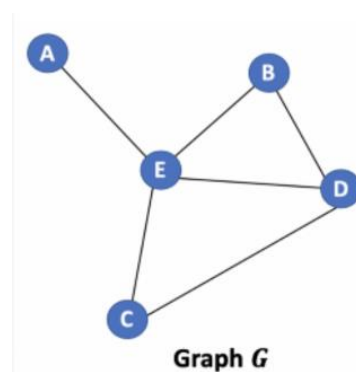
The embeddings generated by GPT-4o and Tabformer are fused—either via concatenation or weighted combination—to form a unified representation of each transaction. These fused features initialize the node and edge attributes in the graph, allowing the model to integrate both structural and semantic information. By combining deep contextual insights with graph-based relational modeling, the framework can more effectively learn and differentiate complex fraudulent patterns in online payment systems.

#### 4.3. GCN-Based Fraud Detection Modeling

The core fraud detection model is built on a two-layer GCN, designed to leverage the structural relationships within the e-commerce transaction graph. The graph, composed of consumer and merchant nodes connected by transaction edges, exhibits complex and often non-linear dependencies that traditional flat feature-based models struggle to capture.

GCNs operate by propagating and aggregating information from a node's neighborhood, enabling each node to iteratively update its representation based on features of connected nodes and edges. This mechanism allows the model to learn not only from individual transactions but also from broader network behaviors. For example, a fraudulent transaction with weak local signals may still be detectable when patterns across multiple interactions or similar behavioral pathways are considered.

As shown in Figure 3, the first layer of the GCN focuses on capturing immediate interaction features, including direct consumer-merchant relationships and transaction attributes. The second layer extends this to higher-order neighbors, enabling the model to learn more abstract, generalized fraud patterns spanning the network. This hierarchical message-passing mechanism allows the GCN to identify subtle irregularities and anomalies indicative of fraudulent activity.



**Figure 3.** Structure of GCN.


Given the severe class imbalance in the dataset, a class-weighted loss function is applied during training. This ensures that the model does not bias toward the majority class (legitimate transactions) and is incentivized to correctly identify the minority class (fraud). This strategy substantially improves recall and F1 scores for fraud detection, which are critical metrics for real-world online payment security systems.

## 5. Model results analysis

As shown in Table 1, the detailed performance metrics of the GCN model on the e-commerce fraud detection task are presented, where normal transactions are labeled as 0 and fraudulent transactions as 1. For normal transactions, the model achieves a precision of 0.98, recall of 1.00, and F1-Score of 0.99 across 1,239,159 samples. These results indicate that the model effectively identifies legitimate transactions with minimal false positives, demonstrating strong reliability in recognizing normal behavior patterns.

For fraudulent transactions, the model achieves perfect precision of 1.00 but a recall of only 0.05, resulting in an F1-Score of 0.09 across 33,365 samples. This indicates that while the model is highly accurate when it predicts fraud, it fails to detect the majority of actual fraudulent cases, likely due to the severe class imbalance in the dataset. The overall accuracy is 0.98, and while the weighted-average F1-Score is high at 0.96, the macro-average F1-Score drops to 0.54, reflecting the model's limited effectiveness in identifying rare fraud instances. These results highlight the need for further improvement in recall for the minority class to enhance overall fraud detection capability.

As shown in Figure 4, the confusion matrix provides a visual summary of the classification performance of the proposed GCN-LLM framework. The model correctly classified 1,239,155 legitimate transactions (true negatives) and detected 1,639 fraudulent transactions (true positives). Notably, only 4 false negatives were recorded, indicating that once a fraudulent transaction is recognized, the model almost never overlooks it. This feature is particularly important in financial applications where missing actual fraud can have severe consequences.



**Figure 4.** Confusion Matrix.

At the same time, 31,726 false positives were observed, meaning that some legitimate transactions were flagged as suspicious. While this reflects a conservative prediction strategy, it ensures that potentially risky patterns are captured with high precision, even if it introduces additional screening effort. This outcome aligns with the precision-recall trade-off shown in Table 1, where the model achieves perfect precision for fraud detection at the expense of low recall.

Overall, the confusion matrix underscores the robustness of the proposed framework in safeguarding against undetected fraud. The combination of structural graph features and LLM-based semantic embeddings enables the model to accurately identify legitimate transactions while capturing high-confidence fraudulent patterns. Despite the recall limitation, these results demonstrate the framework's potential as a scalable, reliable solution for real-world e-commerce fraud prevention. Future work may focus on enhancing recall through techniques such as synthetic minority oversampling, cost-

sensitive learning, or hybrid ensemble approaches to further improve the detection of rare but critical fraudulent transactions.

## 6. Conclusions

This study addresses the increasing challenge of online payment fraud in e-commerce by integrating Large Language Models (LLMs) with Graph Convolutional Networks (GCNs), demonstrating how structural and semantic features can be combined to detect fraudulent transactions under highly imbalanced conditions. The primary objective is to develop a scalable, accurate, and real-time fraud detection framework capable of safeguarding online payment environments.

Through comprehensive data analysis and model evaluation, several key findings emerge: (1) the framework achieves high overall accuracy of 0.98, (2) fraudulent transactions are detected with extremely low false negatives, and (3) the model maintains strong precision despite the pronounced class imbalance. These outcomes indicate that the proposed framework can effectively minimize undetected fraud while providing reliable performance in practical e-commerce settings.

The implications of this study for financial fraud detection are significant. Firstly, the integration of LLM-driven semantic understanding with GCN-based structural learning introduces a novel perspective on modeling complex fraud behaviors. Secondly, the model demonstrates resilience in handling highly imbalanced datasets, overcoming limitations commonly faced by traditional rule-based systems and shallow machine learning methods. Finally, this framework highlights the potential of graph-based deep learning approaches for broader applications in financial security, including anomaly detection, risk assessment, and transaction monitoring.

Despite these promising results, certain limitations remain. The model exhibits a relatively high false positive rate, which may necessitate additional verification or post-processing steps in deployment. Moreover, the analysis relies on a single dataset, which may constrain generalizability across different platforms or regions. Future research could explore dynamic graph modeling to capture evolving fraud patterns over time, incorporate multimodal behavioral data to improve recall without compromising precision, and investigate ensemble or hybrid approaches to further enhance detection performance.

In conclusion, by integrating LLMs and GCNs, this study presents a powerful, scalable, and practical approach for detecting online payment fraud. The framework not only advances methodological understanding but also offers actionable tools for enhancing the security and reliability of e-commerce payment systems, providing a foundation for future developments in financial fraud prevention.

## References

1. A. Ali, S. Abd Razak, S. H. Othman, T. A. E. Eisa, A. Al-Dhaqm, M. Nasser, and A. Saif, "Financial fraud detection based on machine learning: a systematic literature review," *Applied Sciences*, vol. 12, no. 19, p. 9637, 2022. doi: 10.3390/app12199637
2. Y. Alghofaili, A. Albattah, and M. A. Rassam, "A financial fraud detection model based on LSTM deep learning technique," *Journal of Applied Security Research*, vol. 15, no. 4, pp. 498-516, 2020. doi: 10.1080/19361610.2020.1815491
3. M. N. Ashtiani, and B. Raahemi, "Intelligent fraud detection in financial statements using machine learning and data mining: a systematic literature review," *Ieee Access*, vol. 10, pp. 72504-72525, 2021. doi: 10.1109/access.2021.3096799
4. J. I. Z. Chen, and K. L. Lai, "Deep convolution neural network model for credit-card fraud detection and alert," *Journal of Artificial Intelligence*, vol. 3, no. 02, pp. 101-112, 2021.
5. Y. Cheng, J. Guo, S. Long, Y. Wu, M. Sun, and R. Zhang, "Advanced financial fraud detection using GNN-CL model," In *2024 International Conference on Computers, Information Processing and Advanced Education (CIPAE)*, August, 2024, pp. 453-460. doi: 10.1109/cipae64326.2024.00088
6. A. Kesharwani, and P. Shukla, "FFDM GNN: A Financial Fraud Detection Model using Graph Neural Network," In *2024 International Conference on Computing, Sciences and Communications (ICCSC)*, October, 2024, pp. 1-6.

7. I. Padhi, Y. Schiff, I. Melnyk, M. Rigotti, Y. Mroueh, P. Dognin, and E. Altman, "Tabular transformers for modeling multivariate time series," In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, June, 2021, pp. 3565-3569. doi: 10.1109/icassp39728.2021.9414142
8. G. Liu, J. Tang, Y. Tian, and J. Wang, "Graph neural network for credit card fraud detection," In *2021 International Conference on Cyber-Physical Social Intelligence (ICCSI)*, December, 2021, pp. 1-6. doi: 10.1109/iccsi53130.2021.9736204

**Disclaimer/Publisher's Note:** The views, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of PAP and/or the editor(s). PAP and/or the editor(s) disclaim any responsibility for any injury to individuals or damage to property arising from the ideas, methods, instructions, or products mentioned in the content.