

Design and Evaluation of FraudNet: An Intelligent Graph-Based Architecture for Fraud Ring Detection

Boochi Adikeshavulu^{*1}, Dr. S. Usharani²

¹ Student, Department of Computer Applications, Viswam Engineering College, Andhra Pradesh, India

² Professor, Department of Computer Applications, Viswam Engineering College, Andhra Pradesh, India

²anjanasundar80@gmail.com

Abstract — Financial fraud continues to evolve with increasing sophistication, particularly through coordinated fraud rings that exploit distributed transaction patterns to evade detection. Traditional rule-based systems and tabular machine learning models are limited in capturing relational dependencies across entities such as accounts, devices, and identities. This paper presents FraudNet AI, a graph-based deep learning framework that models financial transactions as interconnected nodes within a transaction graph to identify complex fraud structures. The proposed system utilizes the IEEE-CIS Fraud Detection dataset and constructs a heterogeneous graph by linking transactions through shared attributes including device identifiers, IP addresses, billing information, and email domains. A hybrid Graph Neural Network (GNN) architecture combining Graph Convolutional Networks (GCN), Graph Attention Networks (GAT), and GraphSAGE is employed to learn rich node embeddings that encode both intrinsic transaction features and neighborhood relationships. The framework includes preprocessing, graph construction, model training with class imbalance handling, and deployment via a Flask-based API and interactive dashboard. Experimental results demonstrate improved performance over baseline models in terms of AUC-ROC, precision, and recall.

Keywords: Graph Neural Networks; Fraud Detection; Fraud Rings; Transaction Graph; Graph Attention Networks; GCN; GAT; GraphSAGE; Financial Security.

1. Introduction

The rapid growth of digital financial systems has significantly increased exposure to fraudulent activities. Among these, fraud rings represent a major challenge due to their coordinated and distributed nature. These networks exploit multiple accounts, devices, and identities, making detection through traditional methods highly ineffective. Existing fraud detection systems primarily rely on rule-based mechanisms or supervised machine learning models such as decision trees and gradient boosting. While effective for isolated anomalies, these approaches fail to capture relational patterns among entities.

As fraud rings inherently involve interconnected behaviors, ignoring these relationships leads to high false negatives. Recent advancements in deep learning have introduced graph-based methods capable of modeling relationships between entities. Graph Neural Networks (GNNs) enable learning from both node features and graph topology, making them suitable for fraud detection tasks. This paper introduces FraudNet AI, a graph-driven framework that transforms tabular transaction data into a structured graph and applies advanced GNN architectures. The system enhances detection accuracy while providing interpretable insights through graph visualization and attention mechanisms.

2. Literature Survey

Rule-based systems were among the earliest fraud detection solutions, relying heavily on expert knowledge and predefined thresholds. Although interpretable, these systems lacked adaptability to evolving fraud patterns and failed to scale with the volume and complexity of modern digital transactions. Machine learning models such as Random Forests and XGBoost improved detection accuracy by capturing nonlinear relationships in tabular data. However, they treat transactions independently, limiting their effectiveness in identifying collaborative fraud patterns that span multiple entities. Deep learning approaches, including recurrent neural networks, introduced temporal modeling but still failed to capture relational dependencies across entities. Recent studies have highlighted the importance of graph-based methods, where transactions are represented as nodes and relationships as edges. Graph Convolutional Networks (GCNs) aggregate neighborhood information, while Graph Attention Networks (GATs) assign importance weights to connections. GraphSAGE enables inductive learning for unseen nodes. Despite these advancements, most systems lack integration with real-world deployment pipelines and visualization tools. FraudNet AI addresses these limitations by combining multi-variant GNN architectures with a complete deployment framework.

3. Proposed Work / Methodology

The FraudNet AI framework follows a structured five-stage pipeline, from raw data ingestion through model deployment, ensuring modularity and reproducibility at each stage.

3.1 Data Preprocessing

- The IEEE-CIS Fraud Detection dataset is prepared through the following steps:
- Missing values handled using statistical imputation (mean/median/mode strategies)
- Categorical features encoded using label encoding and one-hot encoding
- Numerical features standardized using z-score normalization
- Temporal features transformed to capture cyclical patterns (hour, day-of-week)

3.2 Graph Construction

- Transactions are converted into a structured graph representation:
- Transactions represented as nodes with feature vectors derived from preprocessed attributes
- Edges created based on shared attributes: device ID, IP address, email domain, billing address
- Edge weights assigned based on attribute similarity and temporal proximity between transactions

3.3 GNN Model Architecture

- A hybrid GNN architecture is designed to leverage complementary learning strategies:
- GCN (Graph Convolutional Network): Aggregates neighborhood features through spectral convolution
- GAT (Graph Attention Network): Applies learned attention weights to connections, highlighting suspicious linkages
- GraphSAGE: Enables scalable inductive learning for unseen transaction nodes at inference time

The combined architecture integrates outputs from all three GNN variants, improving robustness and generalization across both seen and unseen fraud patterns.

3.4 Training Strategy

- Weighted cross-entropy loss function to address severe class imbalance (fraud is rare)

- Dropout regularization and batch normalization to prevent overfitting
- Residual connections to ensure gradient stability across deep network layers

3.5 Deployment

- Flask REST API for real-time transaction scoring and fraud probability output
- Interactive dashboard for graph visualization, monitoring, and alert management

4. Results and Discussion

The proposed model was evaluated using the IEEE-CIS Fraud Detection dataset with stratified train-test splitting to preserve the fraud-to-legitimate ratio across partitions. Table 1 presents the performance comparison against established baseline models.

Table 1: AUC-ROC performance comparison of FraudNet AI against baseline models on the IEEE-CIS dataset.

Model	AUC-ROC
Logistic Regression	0.83
Random Forest	0.87
XGBoost	0.89
FraudNet AI (Proposed)	0.93

The GNN-based model significantly outperforms all traditional approaches by leveraging relational information encoded in the transaction graph. With an AUC-ROC of 0.93, FraudNet AI demonstrates a 4-point improvement over XGBoost and a 10-point improvement over Logistic Regression.

4.1 Key Observations

- Fraud rings form dense clusters in the transaction graph, making them visually and computationally distinguishable from legitimate transaction subgraphs.
- The GAT attention mechanism identifies critical inter-transaction connections, providing interpretable signals for fraud analysts.
- Improved recall for the minority fraud class is achieved through the weighted loss function, reducing false negatives in high-risk detection scenarios.

4.2 Discussion

While training the GNN requires greater computational resources than tabular models, inference remains efficient and operates well within real-time latency requirements. The model demonstrates strong scalability and real-world applicability.

The use of GraphSAGE enables the system to generalize to unseen transaction nodes at inference time, a critical capability for production fraud detection systems where new accounts and devices continuously appear.

5. Conclusion

FraudNet AI presents a novel graph-based approach for detecting coordinated fraud rings in financial systems. By integrating advanced GNN architectures—GCN, GAT, and GraphSAGE—with a complete deployment pipeline including Flask API and an interactive dashboard, the system achieves high accuracy and interpretability simultaneously. The framework successfully bridges the gap between theoretical graph learning and practical fraud detection systems. The experimental results confirm that modeling relational transaction data as a graph unlocks detection capabilities that tabular approaches fundamentally cannot achieve. Future work includes real-time streaming graph processing, heterogeneous graph modeling, and enhanced explainability through GNNExplainer integration.

References

- [1] T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," International Conference on Learning Representations (ICLR), 2017.
- [2] P. Velickovic et al., "Graph Attention Networks," International Conference on Learning Representations (ICLR), 2018.
- [3] W. Hamilton et al., "Inductive Representation Learning on Large Graphs," Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [4] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System," ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), 2016.
- [5] L. Breiman, "Random Forests," Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
- [6] M. Fey and J. E. Lenssen, "Fast Graph Representation Learning with PyTorch Geometric," ICLR Workshop on Representation Learning on Graphs and Manifolds, 2019.
- [7] IEEE-CIS, "Fraud Detection Dataset," Kaggle Competition Dataset, 2019.
- [8] A. Paszke et al., "PyTorch: An Imperative Style, High-Performance Deep Learning Library," Advances in Neural Information Processing Systems (NeurIPS), vol. 32, 2019.

- [9] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," Journal of Machine Learning Research (JMLR), vol. 12, pp. 2825–2830, 2011.
- [10] Z. Ying et al., "GNNExplainer: Generating Explanations for Graph Neural Networks," Advances in Neural Information Processing Systems (NeurIPS), 2019.