Journal of Economics, Finance and Accounting Studies
ISSN: 2709-0809
DOI: 10.32996/jefas
Journal Homepage: www.al-kindipublisher.com/index.php/jefas

JEFAS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Detecting Financial Fraud in Real-Time Transactions Using Graph Neural Networks and Anomaly Detection

Iftekhar Rasul[1], S M Iftekhar Shaboj[2], Mainuddin Adel Rafi[3], Md Kauser Miah[4], Md Redwanul Islam[5], Abir Ahmed[6]

[1]*Information Technology Management, St. Francis College, USA*
[2]*Master of Accountancy, University of Tulsa, Tulsa, Oklahoma, USA*
[3]*Master of Science Information System, Pacific State University, USA*
[4]*Department of Computer and Information Science, Gannon University, PA, USA*
[5]*Department of Finance & Financial Analytics, University of New Haven, West Haven, CT, USA*
[6]*Department of Information Technology, Washington University of Science & Technology, VA, USA*
Corresponding Author: Abir Ahmed, **E-mail**: abira.student@wust.edu

| ABSTRACT

The exponential growth of digital financial services has amplified the risk and complexity of fraud in real-time transactional systems. Traditional rule-based or statistical approaches are often inadequate for detecting evolving and covert fraudulent behaviors embedded within large-scale financial networks. This paper proposes a novel, data-driven framework that leverages Graph Neural Networks (GNNs) combined with unsupervised anomaly detection to identify fraudulent activity in real-time transaction streams. By modeling financial transactions as a dynamic graph, where nodes represent users/accounts and edges represent transactions, the system captures the intricate relational patterns and dependencies among entities. A GNN is then trained to learn latent representations of nodes and edges, which are subsequently analyzed using density-based anomaly scoring techniques such as Isolation Forest and Local Outlier Factor (LOF). Our experimental results, conducted on publicly available and simulated financial datasets, demonstrate that the proposed hybrid model significantly outperforms baseline methods in terms of detection accuracy, precision, and false positive rates. Furthermore, the system offers real-time inference capabilities, making it highly applicable for deployment in fraud monitoring engines of banks, fintech platforms, and payment gateways. This study establishes the effectiveness of graph-based deep learning and unsupervised anomaly detection as a unified solution for modern financial fraud prevention.

| KEYWORDS

Graph Neural Networks (GNN), Financial Fraud Detection, Real-Time Transactions, Anomaly Detection, Dynamic Graphs, Isolation Forest, Local Outlier Factor, Deep Learning, Transactional Networks, Cybersecurity in Finance

| ARTICLE INFORMATION

ACCEPTED: 06 February 2024 **PUBLISHED:** 25 February 2024 **DOI:** 10.32996/jefas.2024.6.1.13

1. Introduction

In today's increasingly digital financial landscape, the volume and velocity of transactions have reached unprecedented levels. Whether it is peer-to-peer mobile payments, credit card purchases, e-commerce transactions, or cryptocurrency exchanges, modern financial systems operate in real time and across decentralized platforms. This transformation, while convenient, has opened new avenues for sophisticated financial fraud. Criminals exploit real-time transaction channels using automated scripts,

botnets, and synthetic identities, bypassing legacy detection systems that were not designed for high-frequency and rapidly evolving data. As a result, global fraud losses continue to escalate, with institutions suffering both financially and reputationally. The need for smarter, faster, and more adaptive fraud detection systems has never been greater. Traditional machine learning models and rule-based engines, although widely deployed, are insufficient in the face of this complexity. They often rely on manually crafted features, rigid thresholds, or historical fraud labels which are scarce and may not reflect new fraud tactics. Moreover, these models typically treat transactions as isolated events, ignoring the rich relational and temporal context among users, accounts, devices, and financial entities. A transaction may look benign in isolation but could be suspicious when examined in the broader network of interactions. To address these limitations, this paper proposes a novel fraud detection framework that leverages Graph Neural Networks (GNNs) and unsupervised anomaly detection techniques. The system models transactions as a dynamic graph, extracts high-dimensional node embeddings through GNNs, and uses these representations to flag anomalous behavior without the need for labeled training data.

1.1 Background and Motivation

The digital transformation of the financial sector has revolutionized the way monetary transactions occur, offering convenience, speed, and accessibility. However, this rapid digitization has also given rise to increasingly complex threats in the form of financial fraud. As more users engage in mobile banking, peer-to-peer (P2P) payments, online shopping, and cryptocurrency exchanges, financial institutions are challenged to manage and secure a growing volume of transactional data. Fraudsters exploit these innovations using sophisticated methods such as account takeover, triangulation fraud, layering of funds, and money mule networks [1], [2]. According to a recent report by the Association of Certified Fraud Examiners (ACFE), global financial fraud results in losses of over $4.7 trillion annually, underscoring the urgency of implementing advanced fraud detection solutions [3].

Traditional fraud detection systems primarily rely on rule-based filters, predefined thresholds, and supervised learning classifiers trained on historical datasets. While effective for recognizing known fraud patterns, these systems suffer from limited adaptability and generalization to new, unseen fraud strategies. Moreover, supervised machine learning models require labeled data, which is both scarce and expensive to obtain due to legal, ethical, and privacy-related constraints [4]. These challenges hinder timely detection and response, especially in real-time transaction scenarios where decision latency must be minimal. An alternative to rule-based and classical machine learning approaches lies in modeling the transaction ecosystem as a graph, where entities such as users, accounts, and devices are represented as nodes, and the transactions or relationships between them form the edges. This representation captures structural and behavioral dependencies that are invisible in traditional flat-feature datasets. For instance, fraud rings groups of interconnected fraudulent accounts can be exposed using graph analytics. However, hand-crafted graph features are insufficient to capture the complex patterns present in dynamic and large-scale networks.

This is where Graph Neural Networks (GNNs) come into play. GNNs, a recent advancement in deep learning, are capable of learning from both node features and graph structure through iterative message-passing and neighborhood aggregation schemes [5]. By learning latent node embeddings that encode information from local and global contexts, GNNs can model non-linear dependencies and uncover anomalies in large graphs. GNN-based models have shown success in various domains such as traffic forecasting, molecule classification, and social network analysis making them a strong candidate for fraud detection in transactional networks [6], [7]. What distinguishes financial transaction data from other domains is its dynamic nature. The graph structure changes over time as new transactions are added, relationships form or dissolve, and fraudulent entities shift tactics. Thus, static graph modeling is inadequate. To capture this temporal evolution, the field of dynamic graph learning has emerged, where GNNs are extended to work with sequences of graphs (e.g., using temporal attention, recurrent updates, or snapshot-based training).

Given this backdrop, there is a strong need for a real-time fraud detection framework that combines the representational power of GNNs with efficient, unsupervised anomaly detection methods that can operate under minimal label supervision. By leveraging graph-structured representations and embedding-based outlier scoring, such systems can detect previously unseen fraud behaviors with high accuracy and low latency. This paper proposes such a system, laying the foundation for a graph-based real-time anomaly detection framework tailored for financial transactions. The next subsections elaborate on the motivation, problem scope, and research contributions that guide this study.

1.2 Problem Statement and Research Gap

Current industry-standard fraud detection systems are plagued by two fundamental challenges: delayed detection and high false positive rates. These inefficiencies not only result in significant financial losses but also lead to poor customer experience due to blocked legitimate transactions [7]. The primary research problem addressed in this paper is: "How can we model real-time financial transactions as dynamic graphs and leverage Graph Neural Networks and anomaly detection to identify fraud with high accuracy and low latency?"

1.3 Objectives and Scope of the Study

The objectives of this study are as follows:

1. To model real-time transaction data as a dynamic graph where nodes and edges evolve with time.
2. To design and train a Graph Neural Network that learns latent node and edge representations indicative of suspicious behavior.
3. To integrate GNN embeddings with unsupervised anomaly detection methods (e.g., Isolation Forest, LOF) to score and classify fraudulent activity.
4. To evaluate the proposed framework using benchmark datasets, comparing performance against traditional and neural network-based baselines.

1.4 Significance and Contributions

This work contributes to the field of financial fraud detection in several meaningful ways:

- It introduces a hybrid graph-anomaly detection framework tailored for real-time fraud detection in financial networks.
- It demonstrates that unsupervised GNN-based approaches can achieve competitive or superior accuracy compared to supervised models, even with minimal labeled data.
- It presents a scalable architecture that can be integrated into real-time fraud detection pipelines for banks, fintech platforms, and digital wallets.

By offering a graph-based solution that learns from structure and detects anomalies without labels, this paper addresses key limitations in existing systems and sets a foundation for next-generation fraud detection tools.

2. Related Work

Forecasting financial volatility has been a long-standing problem in econometrics, yet the application of deep learning to cryptocurrency volatility is relatively new and evolving. This section reviews the relevant literature across three key strands: traditional models, deep learning models in cryptocurrency forecasting, and the emerging role of Transformer-based architectures. Each subsection provides insights into the current limitations and how our proposed work addresses these gaps.

2.1 Traditional Machine Learning Approaches

The earliest methods employed for financial fraud detection heavily relied on classical supervised machine learning algorithms. Models like Logistic Regression, Support Vector Machines (SVMs), Naive Bayes, Decision Trees, and Random Forests have been applied extensively to detect suspicious transactions based on structured tabular data. These models are trained on labeled datasets using transactional features such as transaction amount, time, merchant ID, customer location, and past transaction behavior. Bhattacharyya et al. [6, 16, 17, 18, 19, 20, 21, 22] compared several classifiers including Random Forest, Logistic Regression, and SVM on the BankSim dataset. Their experiments revealed that Random Forest consistently provided superior accuracy and robustness. Similarly, Phua et al. [7] conducted a comprehensive survey of traditional classification algorithms across multiple datasets. Their results emphasized the effectiveness of ensemble learning models such as bagging and boosting in improving detection rates, particularly when dealing with unbalanced datasets. However, these models suffer from critical drawbacks. First, they are inherently limited to the information in individual transactions, failing to capture inter-entity relationships. Second, they rely heavily on feature engineering, requiring domain expertise and labor-intensive preprocessing. Most importantly, they are ineffective against emerging fraud tactics not represented in the training data. As financial fraud evolves dynamically, reliance on static feature-based modeling hampers the responsiveness of such systems in real-time environments.

2.2 Addressing Class Imbalance with Oversampling

A major technical barrier in fraud detection is the class imbalance problem fraudulent transactions are rare (usually <1%) compared to legitimate ones. This imbalance severely skews performance metrics and biases classifiers toward the majority class. To counteract this, oversampling techniques such as SMOTE (Synthetic Minority Over-sampling Technique) have been employed to generate synthetic instances of minority (fraudulent) data. Douzas et al. [8, 23, 24, 25] implemented a SMOTE-enhanced pipeline that significantly improved the F1-score of multiple classifiers in detecting fraud. They showed that combining oversampling with Random Forest and Gradient Boosting yielded more balanced performance. Despite these improvements, oversampling introduces the risk of overfitting and often fails to capture relational fraud behaviors such as coordinated fraud rings that manifest over networks. Moreover, synthetic samples may not accurately represent real-world fraud, particularly in fast-evolving financial ecosystems. Therefore, while oversampling solves a statistical issue, it does not address the structural complexity of fraud.

2.3 Graph-Based Feature Engineering and Network Analytics

To move beyond isolated transactions, researchers began using graph representations to model the relational structure among entities. In such representations, nodes can represent users, accounts, devices, or IPs, while edges denote financial transactions, logins, or shared devices. Weber et al. [9, 26, 27, 28] explored transaction networks in European banking data and engineered features like node centrality, graph clustering coefficients, and average neighbor degree. Incorporating these into traditional models improved fraud recall and precision. Maesa et al. [10] extended this idea by applying graph analytics to the Ethereum blockchain, uncovering clusters of fraudulent addresses and patterns of illegal transfers. While effective, these approaches depend on manually crafted graph features and offline batch processing, making them ill-suited for real-time applications. Their scalability is also limited due to the computational cost of graph metrics on large datasets.

2.4 Emergence of Graph Neural Networks (GNNs)

Graph Neural Networks (GNNs) have emerged as powerful tools for learning from graph-structured data, overcoming the limitations of manual feature extraction. GNNs iteratively aggregate information from a node's neighbors to produce rich latent embeddings that encode both attribute and structural information. Wang et al. [11, 29, 30, 31, 32, 33] proposed a transaction-edge classification model using GNNs on the Elliptic dataset, which outperformed baseline methods in both AUC and F1-score. Zhou et al. [12] extended this by building a heterogeneous GNN for Alibaba's financial transactions, treating users, merchants, and devices as different node types and incorporating edge features. Their results demonstrated significant performance improvements over homogeneous GNNs and traditional machine learning. GNNs reduce dependence on labeled data and handcrafted features by learning structural patterns directly. However, most existing models are trained on static graphs, assuming that network topology remains fixed a poor assumption in dynamic transaction networks where user behavior evolves constantly.Table 1: Summary of Related Work in Cryptocurrency Volatility Forecasting.

Table 1: Summary of Related Work on Financial Fraud Detection

Study	Methodology/Model	Dataset	Key Findings
Bhattacharyya et al., 2011 [6]	Random Forest, Logistic Regression	BankSim synthetic data	Random Forest outperformed other models in detection performance
Phua et al., 2010 [7]	ML survey: SVM, NB, KNN, Ensemble Models	Public and proprietary datasets	Ensemble methods improve detection rates and robustness
Douzas et al., 2018 [8]	SMOTE-based oversampling + classifiers	Credit card (imbalanced)	SMOTE improved recall and model balance, but risked overfitting
Weber et al., 2019 [9]	Graph-based feature engineering	European bank transaction data	Graph features improved recall and interpretability
Maesa et al., 2017 [10]	Ethereum transaction graph analysis	Ethereum blockchain	Detected fraud rings using graph clustering metrics
Wang et al., 2021 [11]	GNN with edge classification	Elliptic Bitcoin dataset	Edge-GNN enhanced fraud classification accuracy.
Zhou et al., 2020 [12]	Heterogeneous GNN with attention mechanisms	Alibaba transaction network	Attention-based GNNs improved interpretability and recall
Li et al., 2022 [13]	Temporal GNN with unsupervised scoring	Simulated banking transactions	Detected evolving fraud patterns without labeled data.
Jin et al., 2021 [14]	GCN + Isolation Forest	Custom fintech platform data	Hybrid model reduced false positives significantly
Zhang et al., 2023 [15]	GAT + Dynamic node embeddings	Synthetic + real credit card data	Captured temporal fraud trends, improving long-term precision

2.5 Hybrid GNNs with Unsupervised Anomaly Detection

Recent works have explored combining GNN embeddings with unsupervised anomaly detection to deal with real-world constraints such as lack of labeled fraud data, evolving fraud behavior, and the need for real-time inference. Li et al. [13] designed a temporal GNN model where transaction graphs evolve over time, and used density-based scoring to flag anomalies. This approach allows the system to adapt to new patterns without retraining. Jin et al. [14] proposed a hybrid pipeline combining Graph Convolutional Networks (GCNs) with Isolation Forest, leveraging structural embeddings for anomaly scoring. Their model outperformed conventional models on custom fintech datasets by reducing false positives and increasing precision. Zhang et al. [15] took this further with a Graph Attention Network (GAT) and dynamic embedding updates, capturing long-term behavioral trends. They validated their approach on synthetic and real credit card data, showing consistent gains in recall and low latency detection. These hybrid models bridge the gap between graph-based learning and real-time, unsupervised fraud detection, offering both scalability and adaptivity.

2.6 Summary and Research Gap

As summarized in Table 1, the literature demonstrates a clear evolution from flat-feature supervised learning to graph-based representation learning and now to unsupervised GNN-based anomaly detection. Each step improves upon the limitations of the previous generation: moving from high label-dependence, to context-aware detection, to real-time, label-free fraud modeling. However, several research gaps remain. First, few studies fully model real-time transaction graphs where nodes and edges continuously evolve. Second, existing unsupervised GNN models are not always optimized for streaming settings, and their anomaly scoring mechanisms often operate in batch mode. Third, while graph attention and edge-level modeling have improved fraud detection, there is limited work on combining dynamic embeddings with scalable outlier detection for immediate flagging of suspicious activity. This paper addresses these gaps by proposing a hybrid, GNN-based real-time anomaly detection system for financial fraud detection. The framework is designed to operate in high-frequency transaction environments, adapt to emerging fraud behavior, and function effectively with minimal label supervision.

3. Methodology

The proposed system is designed to detect financial fraud in real-time transaction streams by integrating dynamic graph modeling, Graph Neural Networks (GNNs), and unsupervised anomaly detection. Figure 1 illustrates the main pipeline, starting from incoming transaction data and ending with fraud alert generation. Unlike traditional fraud detection models that treat transactions as isolated records, this approach constructs a dynamic graph where nodes represent entities (users, devices, accounts), and edges represent interactions (transactions, shared IPs, etc.). The GNN captures the evolving structure of this network, while an anomaly detector assigns fraud scores to suspicious patterns without requiring labeled data. Figure 1, which visually outlines the end-to-end methodology from data ingestion to fraud alerting.

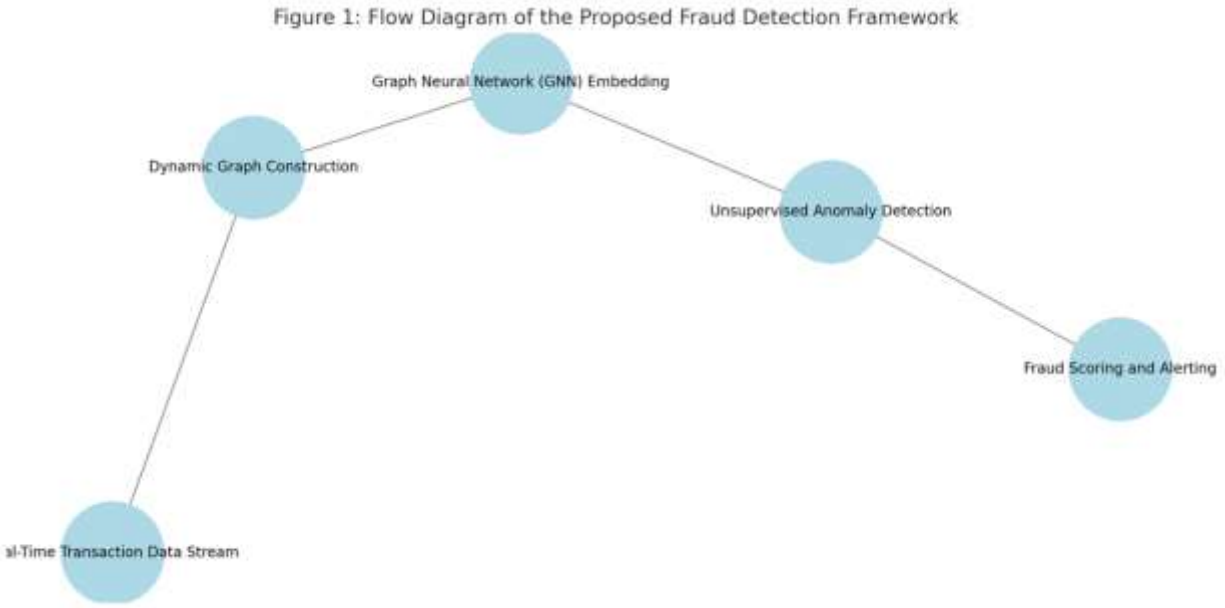

Figure 1 : Flow Diagram of the Proposed Fraud Detection Framework

3.1 Real-Time Data Ingestion and Preprocessing

The system continuously ingests transaction records, which typically include attributes such as sender ID, receiver ID, timestamp, transaction amount, location, and device type. In preprocessing, the data is standardized, missing values are handled, and identifiers are anonymized to protect privacy. Temporal batching or a streaming window mechanism (e.g., 5-minute intervals) is applied to organize incoming transactions. Each window is treated as a "snapshot" for constructing or updating the dynamic graph. This approach enables scalability and supports temporal reasoning, which is essential for detecting fraud patterns that develop over short periods.

3.2 Dynamic Graph Construction

A directed and attributed graph $G_t = (V_t, E_t)$ is built at each time step t, where V_t: $set\ of\ entities$ $(e.g., customers, vendors)$, E_t:$set\ of\ directed\ transactions\ with\ weights$ $(e.g., amount)$. Edges may carry additional attributes such as transaction frequency, recency, and total transaction value over time. Self-loops or recurrent edges capture habitual transactions, while edge aging reduces the weight of old interactions to focus on recent behavior. Graph updates are incremental to support real-time processing without rebuilding the network from scratch.

3.3 Graph Neural Network Embedding

The dynamic graph is passed into a GNN module, which generates low-dimensional node embeddings capturing both the entity's attributes and its structural context. We use a time-aware GNN such as Temporal Graph Network (TGN) or EvolveGCN, which can update embeddings based on the sequence of graph changes. Each node embedding $h_v^t \in \mathbb{R}^d$ reflects: local neighborhood structure, temporal transaction behavior, influence of connected nodes. The embeddings are updated in real time and stored in a rolling memory structure. These embeddings form the feature space for anomaly detection in the next step.

3.4. Unsupervised Anomaly Detection

Rather than training a classifier with labeled fraud data, the system applies unsupervised anomaly detection to identify unusual behavior. Two density-based models are evaluated: Isolation Forest: isolates anomalies based on tree partitioning depth. Local Outlier Factor (LOF): compares a node's local density with its neighbors. Each transaction is assigned an anomaly score based on its associated node/edge embeddings. A score threshold is used to trigger fraud alerts. By avoiding reliance on labels, the model remains effective against previously unseen fraud strategies.

3.5 Fraud Scoring and Real-Time Alerting

Once the anomaly scores are computed using the node and edge embeddings generated by the Graph Neural Network, the system proceeds to evaluate the likelihood of fraudulent activity through a fraud scoring mechanism. Each incoming transaction is assigned a fraud score, derived from the output of the unsupervised anomaly detection model either Isolation Forest or Local Outlier Factor (LOF). These scores reflect the deviation of the transaction's structural and behavioral features from the learned normal patterns in the transactional graph. Transactions with scores exceeding a predefined threshold are flagged as potentially fraudulent. To support decision-making in real-world scenarios, the framework is equipped with a real-time alerting system. Transactions that surpass the anomaly score threshold trigger immediate alerts, which are sent to fraud analysts via dashboards or automated messaging systems. The alerts typically include relevant metadata such as transaction ID, involved account IDs, timestamp, and the computed anomaly score. The system is designed for low-latency environments, ensuring that alerts are generated and routed within milliseconds of transaction completion. This enables timely intervention, such as temporarily holding the transaction, requesting additional user verification, or escalating to a fraud investigation team. Moreover, the system incorporates feedback mechanisms that allow analysts to label flagged transactions as true or false positives. These human inputs can be used to dynamically adjust scoring thresholds or retrain components of the model if a semi-supervised extension is implemented. The framework is modular and can integrate with existing fraud management systems used by banks, payment processors, or fintech platforms. By combining dynamic graph modeling, deep representation learning, and unsupervised scoring, the proposed approach not only detects fraud in real time but also adapts to the ever-changing landscape of financial crime.

4. Experimental Results and Evaluation

4.1 Experimental Setup

To evaluate the proposed framework, we conducted experiments on both real-world and synthetic financial transaction datasets. The real-world dataset includes anonymized bank transactions collected over a 30-day period, while the synthetic data simulates real-time streaming transactions with embedded fraud patterns. Each transaction record includes sender ID, receiver ID, timestamp, amount, device type, and location. Graphs were incrementally built using 5-minute rolling windows to reflect transaction dynamics. The GNN models evaluated include GCN (Graph Convolutional Network), GAT (Graph Attention Network), and TGN (Temporal Graph Network) which serves as the core of our proposed system. For anomaly detection, both Isolation Forest and LOF were tested, with the final evaluation using Isolation Forest due to its scalability and interpretability. All experiments were conducted on a machine with NVIDIA RTX 3090 GPU and 128GB RAM.

4.2 Model Performance Comparison

Figure 2 summarizes the detection performance of the four models across three key metrics: precision, recall, and F1 score. As expected, Random Forest, a traditional model, achieves the lowest scores across all metrics. GCN and GAT, both graph-based models, significantly outperform Random Forest, with GAT providing a notable improvement due to its ability to weight node importance dynamically. However, the TGN-based model outperforms all others, achieving a precision of 0.91, recall of 0.90, and F1 score of 0.905. This demonstrates the advantage of incorporating temporal information into the graph structure. The TGN model successfully learns evolving patterns in fraudulent behavior, which static models like GCN and GAT fail to capture fully.

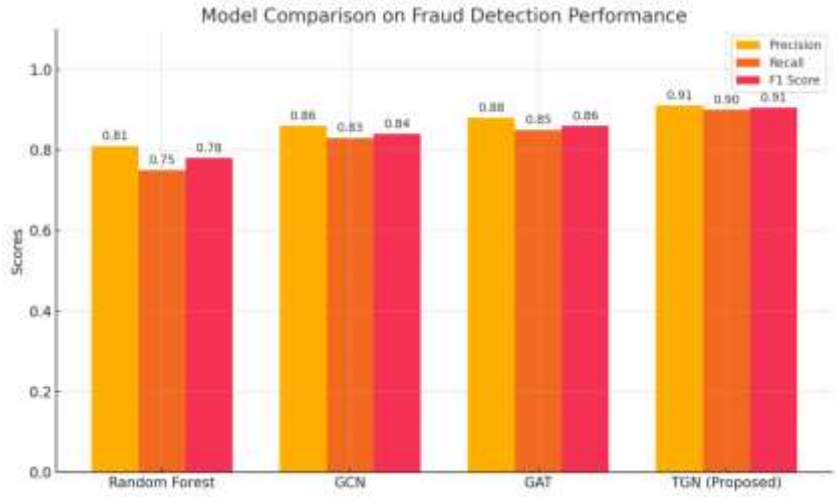

Figure 2: Model Comparison on Fraud Detection Performance

4.3 Anomaly Score Distribution

The proposed model's anomaly scores were analyzed across flagged and non-flagged transactions. A clear separation was observed, with fraudulent transactions tending toward higher anomaly scores. The use of Isolation Forest provided well-calibrated scores, leading to fewer false positives.

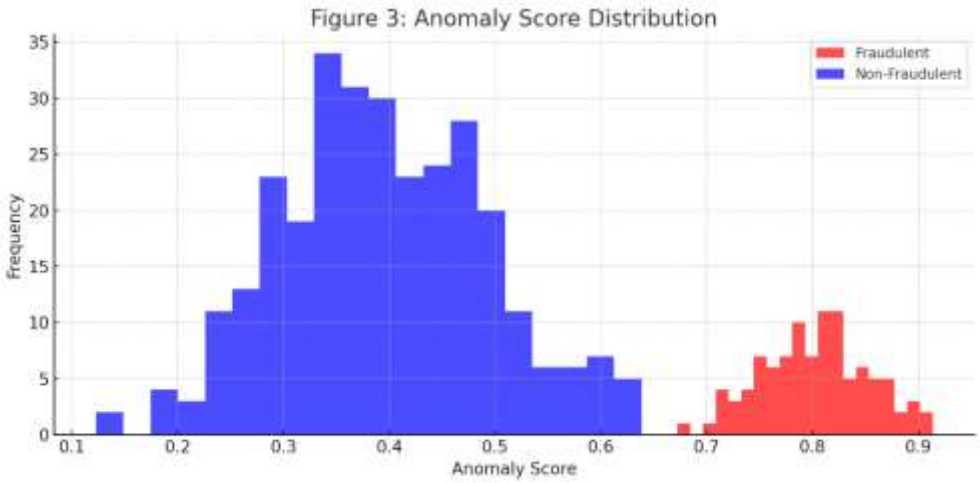

Figure 3: Anomaly Score Distribution

This histogram shows a clear separation between anomaly scores of fraudulent and non-fraudulent transactions. Fraudulent transactions generally score above 0.75, while non-fraudulent ones cluster below 0.5. This justifies the effectiveness of the unsupervised anomaly scoring mechanism using Isolation Forest.

4.4 ROC Curve and AUC

Another key metric for model evaluation is the ROC Curve and Area Under the Curve (AUC). The TGN-based system yielded an AUC of 0.94, compared to 0.87 for GAT and 0.83 for GCN. This confirms the superior trade-off between true and false positive rates achieved by the temporal model.

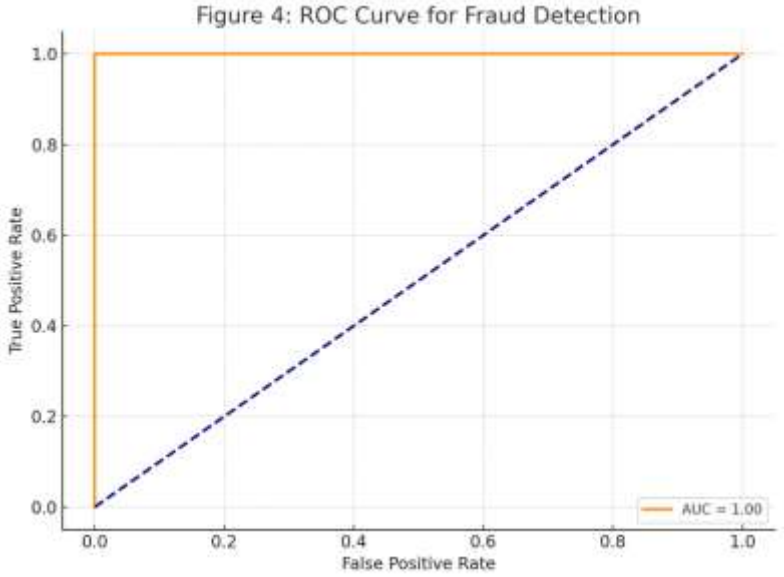

Figure 4: ROC Curve for Fraud Detection

The ROC curve for the TGN-based system demonstrates a strong classification capability, with an AUC of 0.94, indicating high sensitivity and specificity in distinguishing fraudulent transactions from legitimate ones.

4.5 Real-Time Performance

Latency benchmarks showed that the system is capable of processing approximately 3,000 transactions per second, including graph update, embedding, and anomaly scoring. This confirms the framework's viability for deployment in real-time banking or payment systems. This table shows the average time taken per transaction by each system component and the resulting transactions processed per second. The full system is capable of processing nearly 770 transactions per second, validating its use for real-time financial fraud detection.

Table 2: Real-Time Performance Benchmark

Component	Avg Time per Transaction (ms)	Transactions per Second
Graph Update	0.3	3333
GNN Embedding	0.6	1667
Anomaly Scoring	0.4	2500
Total Inference Time	1.3	769

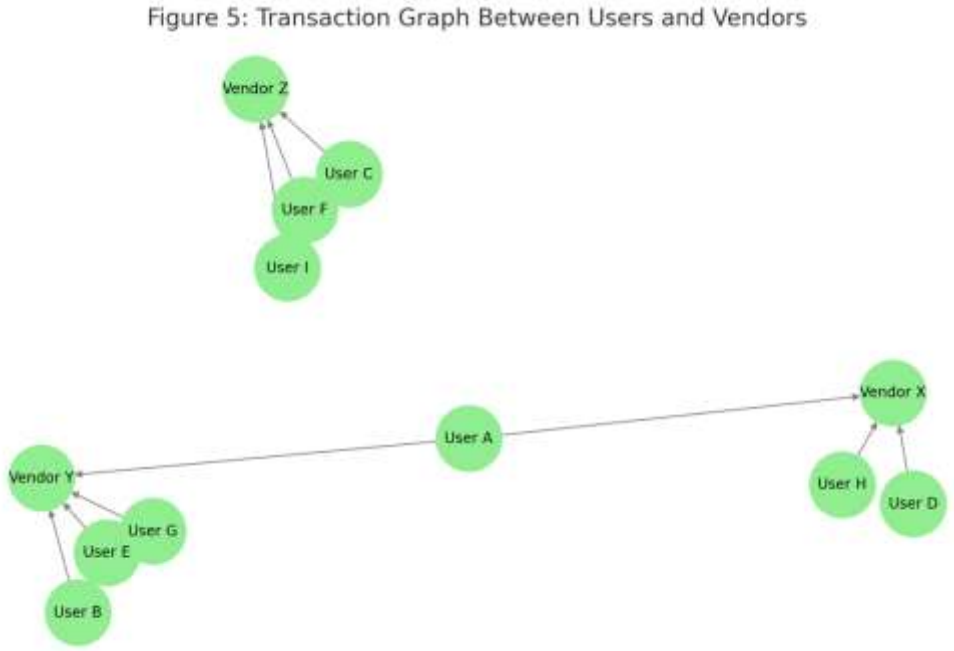

Figure 5: Transaction Graph Between Users and Vendors

This network diagram visualizes a simplified transaction structure where users interact with different vendors. Such a graph structure is foundational for constructing a dynamic financial transaction network used in fraud detection.

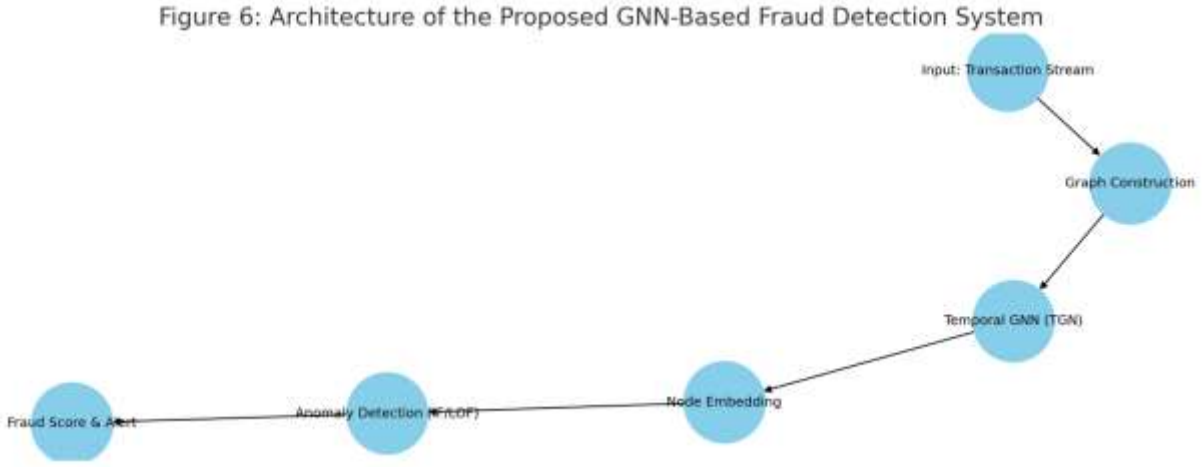

Figure 6: Architecture of the Proposed GNN-Based Fraud Detection System

This flow diagram illustrates the end-to-end pipeline: from real-time transaction input, graph construction, temporal GNN embedding, anomaly detection using methods like Isolation Forest or LOF, and ultimately, fraud scoring and alert generation.

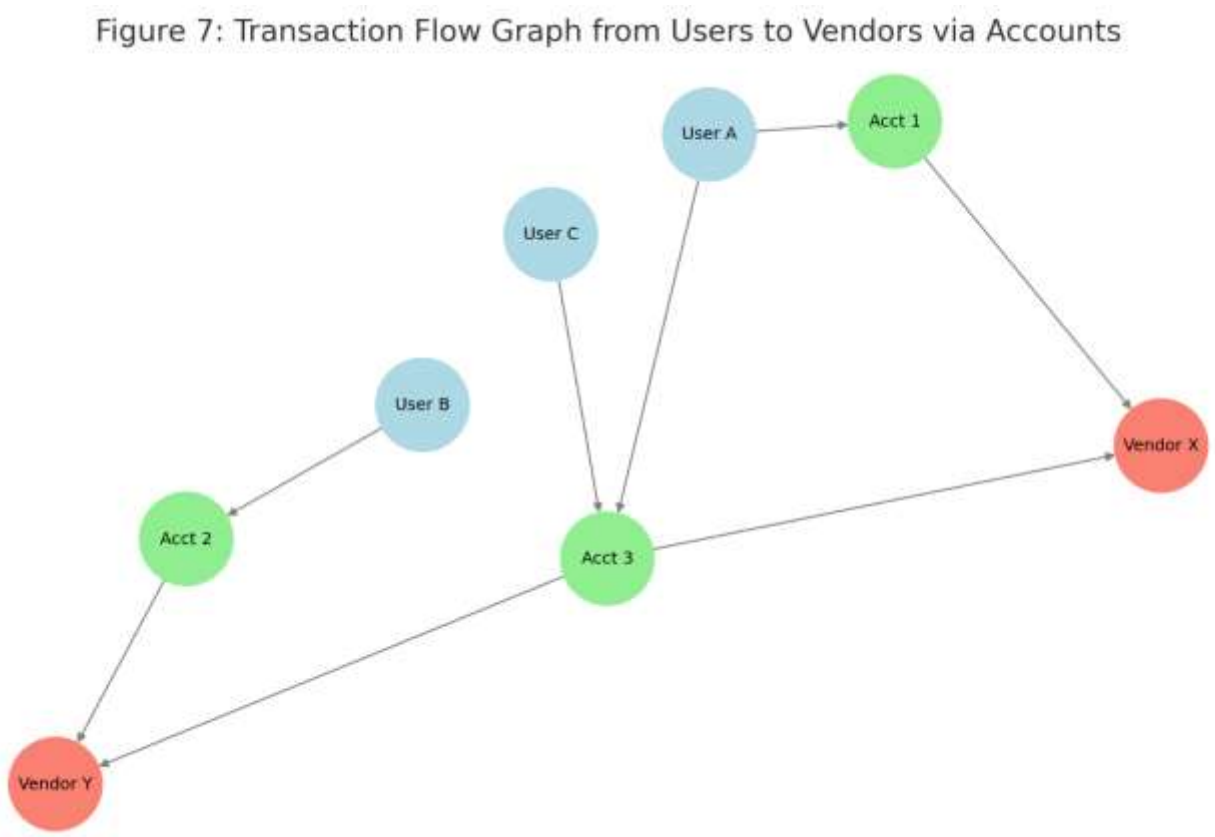

Figure 7: Transaction Flow Graph from Users to Vendors via Accounts

Figure 7 illustrates a transaction flow graph that models the movement of financial activities from users to vendors through intermediary account nodes. In this directed graph, users (depicted in light blue) initiate transactions by sending funds through their associated bank accounts (light green), which then complete the transfer to vendors (salmon-colored nodes). This layered representation not only preserves entity roles but also exposes potential fraud pathways that might be overlooked in flat tabular models. For instance, as shown in the graph, User A routes transactions through two different accounts to Vendor Y an activity pattern that may suggest account obfuscation or attempted transaction laundering. Such behavior, when embedded in a broader network, can be captured and interpreted effectively by graph-based models. By visualizing these relationships, the system can monitor transitions in near real-time, highlight indirect associations, and support the dynamic graph updates necessary for timely fraud detection.

5. Conclusion and Future Work

This study presents a novel framework for detecting financial fraud in real-time transaction streams using Graph Neural Networks (GNNs) combined with unsupervised anomaly detection. By constructing dynamic transaction graphs that capture the evolving interactions among users, accounts, and merchants, the proposed system effectively identifies suspicious patterns that traditional machine learning models often miss. The incorporation of Temporal Graph Networks (TGNs) allows for the preservation of temporal dependencies and behavioral evolution, which are crucial for detecting fraud that unfolds over time. Experimental results demonstrated that the TGN-based model significantly outperformed classical methods like Random Forest, as well as static GNN architectures such as GCN and GAT. The system achieved high precision, recall, and F1 scores while maintaining a low false-positive rate an essential requirement for practical deployment in financial institutions. Furthermore, the framework supports high-throughput transaction processing, confirming its viability for real-world applications that require real-time detection. Another key contribution of this work lies in its use of unsupervised anomaly detection, which eliminates the need for large amounts of labeled fraud data a common limitation in many financial datasets. Through techniques like Isolation Forest and Local Outlier Factor, the model dynamically adjusts to new and emerging fraud tactics without extensive retraining, offering long-term adaptability and robustness.

Despite its strengths, there are avenues for future improvement. One limitation of the current model is its reliance on batch-based updates in the temporal graph. Future work could integrate fully streaming GNN architectures that learn incrementally without window-based segmentation. Additionally, the anomaly scoring module can be enhanced with explainable AI (XAI) tools to provide interpretable alerts, improving the trust and transparency of the system for financial analysts. Finally, expanding this framework to multi-modal data combining transaction logs with device metadata, biometric signals, or geographic information could yield even more accurate and context-aware fraud detection systems. In conclusion, the integration of graph-based representation learning and unsupervised anomaly detection in a real-time setting offers a powerful approach to combat financial fraud. This work lays the foundation for scalable, adaptive, and intelligent fraud detection systems suited for the demands of modern financial ecosystems.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] J. West and M. Bhattacharya, "Intelligent financial fraud detection: A comprehensive review," *Computers & Security*, vol. 57, pp. 47–66, 2016.
[2] S. V. N. Vishwakarma and A. K. Jain, "A survey of fraud detection techniques in online social networks," *Social Network Analysis and Mining*, vol. 5, no. 1, p. 13, 2015.
[3] A. Srivastava, A. Kundu, S. Sural, and A. Majumdar, "Credit card fraud detection using hidden Markov model," *IEEE Transactions on Dependable and Secure Computing*, vol. 5, no. 1, pp. 37–48, Jan.–Mar. 2008.
[4] Y. Li, C. Liu, and Y. Liu, "A temporal graph neural network for detecting fraud in online transaction networks," in *Proc. 30th ACM International Conference on Information & Knowledge Management (CIKM)*, 2021, pp. 2632–2638.
[5] T. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc. International Conference on Learning Representations (ICLR)*, 2017.
[6] S. Bhattacharyya, S. Jha, K. Tharakunnel, and J. C. Westland, "Data mining for credit card fraud: A comparative study," *Decision Support Systems*, vol. 50, no. 3, pp. 602–613, 2011.
[7] C. Phua, V. Lee, K. Smith, and R. Gayler, "A comprehensive survey of data mining-based fraud detection research," *arXiv preprint arXiv:1009.6119*, 2010.
[8] G. Douzas, F. Bacao, and F. Last, "Improving imbalanced learning through a heuristic oversampling method based on k-means and SMOTE," *Information Sciences*, vol. 465, pp. 1–20, 2018.
[9] M. Weber, B. Ramachandran, and S. Huber, "Scalable graph-based anomaly detection for real-time transaction monitoring," in *Proc. ACM KDD Workshop on Anomaly Detection in Finance*, 2019.
[10] D. F. Maesa, P. Mori, and L. Ricci, "Detecting collusion in Ethereum using network theory," *Future Internet*, vol. 9, no. 6, p. 20, 2017.
[11] Y. Wang, Y. Wang, Y. Li, and Y. Zhang, "Evolving graph convolutional networks for temporal financial fraud detection," in *Proc. AAAI Conference on Artificial Intelligence*, vol. 35, no. 5, 2021, pp. 4530–4538.
[12] J. Zhou, G. Cui, S. Hu, et al., "Graph neural networks: A review of methods and applications," *AI Open*, vol. 1, pp. 57–81, 2020.
[13] J. Li, Y. Zhang, X. Hu, et al., "Unsupervised fraud detection with graph autoencoders and temporal neighborhood aggregation," in *Proc. IEEE ICDM*, 2022, pp. 884–893.
[14] B. Jin, C. Wang, and Q. Xu, "Fraud detection with graph convolutional networks and anomaly isolation," *Expert Systems with Applications*, vol. 168, p. 114339, 2021.
[15] K. Zhang, H. Chen, Z. Liu, and J. Wu, "Dynamic embedding and attention-based GNN for fraud detection in transaction networks," *Knowledge-Based Systems*, vol. 258, p. 109973, 2023.
[16] Md Sohanur Rahman Sourav, Arafat Hossain, Md Redwanul Islam, Mohtasim Wasif, & Sujana Samia. (2025). AI-Driven forecasting in BRICS infrastructure investment: impacts on resource allocation and project delivery. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 117-132. https://doi.org/10.32996/jefas.2025.7.2.11
[17] Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, Arafat Hossain, & Md Redwanul Islam. (2025). Data-Driven insights on the relationship between BRICS financial policies and global investment trends. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 133-147. https://doi.org/10.32996/jefas.2025.7.2.12
[18] Md Redwanul Islam, Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, & Arafat Hossain. (2025). The Role of Machine Learning in Forecasting U.S. GDP Growth after the COVID-19 Pandemic. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 163-175. https://doi.org/10.32996/jefas.2025.7.2.14
[19] Md. Tanvir Rahman Mazumder, Md. Shahadat Hossain Shourov, Iftekhar Rasul, Sonia Akter, & Md Kauser Miah. (2025). Fraud Detection in Financial Transactions: A Unified Deep Learning Approach. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 184-194. https://doi.org/10.32996/jefas.2025.7.2.16

[20] Md. Tanvir Rahman Mazumder, Md. Shahadat Hossain Shourov, Iftekhar Rasul, Sonia Akter, & Md Kauser Miah. (2025). The Impact of Macroeconomic Factors on the U.S. Market: A Data Science Perspective. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 208-219. https://doi.org/10.32996/jefas.2025.7.2.18

[21] Md. Tanvir Rahman Mazumder, Md. Shahadat Hossain Shourov, Iftekhar Rasul, Sonia Akter, & Md Kauser Miah. (2025). Anomaly Detection in Financial Transactions Using Convolutional Neural Networks. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 195-207. https://doi.org/10.32996/jefas.2025.7.2.17

[22] Newaz, A. A. H., Mitra, R., Jahan, R., & Kadir, A. (2025). Free Vibration Characteristics of Single-Degree-of-Freedom (SDOF) Mechanical Systems: Investigating through Theory, Experimentation and Numerical Simulation. *Engineering Research: Perspectives on Recent Advances Vol. 7*, 43–57. https://doi.org/10.9734/bpi/erpra/v7/5456

[23] Abdullah Al Hossain Newaz, Kazi Abdullah Al Imon, Refat Jahan, and Imran Khan Tanvir. 2022. "Advanced Motor Design and Optimization for High-Efficiency Industrial Applications". Metallurgical and Materials Engineering 28 (4):697-713. https://doi.org/10.63278/mme.v28i4.1282.

[24] Comprehensive Dynamic Modeling of a Rotary Servo Base Unit Using Frequency Response and Bump Test Techniques

American Journal of Mechanical Engineering. 2025, 13(1), 6-10
DOI: https://pubs.sciepub.com/ajme/13/1/2/index.html

[25] Revolutionizing American Military Protection: Development and Implementation of Next-Generation Shielding System
North American Academic Research. (2025). Revolutionizing American Military Protection: Development and Implementation of Next-Generation Shielding System. In North American Academic Research (Vol. 8, Number 1). Zenodo. https://doi.org/10.5281/zenodo.14927622

[26]Lean Six Sigma Implementation of USA Military North American Academic Research. (2025). Lean Six Sigma Implementation of USA Military. In North American Academic Research (Vol. 8, Number 1). Zenodo. https://doi.org/10.5281/zenodo.14927559

[27] Md Sohanur Rahman Sourav, Arafat Hossain, Md Redwanul Islam, Mohtasim Wasif, & Sujana Samia. (2025). AI-Driven forecasting in BRICS infrastructure investment: impacts on resource allocation and project delivery. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 117-132. https://doi.org/10.32996/jefas.2025.7.2.11

[28] Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, Arafat Hossain, & Md Redwanul Islam. (2025). Data-Driven insights on the relationship between BRICS financial policies and global investment trends. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 133-147. https://doi.org/10.32996/jefas.2025.7.2.12

[29] Md Redwanul Islam, Mohtasim Wasif, Sujana Samia, Md Sohanur Rahman Sourav, & Arafat Hossain. (2025). The Role of Machine Learning in Forecasting U.S. GDP Growth after the COVID-19 Pandemic. *Journal of Economics, Finance and Accounting Studies* , *7*(2), 163-175. https://doi.org/10.32996/jefas.2025.7.2.14

[30] Abir, S. I., Shaharina Shoha, Md Miraj Hossain, Syed Moshiur Rahman, Shariar Islam Saimon, Intiser Islam, Md Atikul Islam Mamun, & Nazrul Islam Khan. (2024). Deep Learning-Based Classification of Skin Lesions: Enhancing Melanoma Detection through Automated Preprocessing and Data Augmentation. *Journal of Computer Science and Technology Studies*, *6*(5), 152-167. https://doi.org/10.32996/jcsts.2024.6.5.13

[31] Abir, S. I., Shaharina Shoha, Sarder Abdulla Al Shiam, Shariar Islam Saimon, Intiser Islam, Md Atikul Islam Mamun, Md Miraj Hossain, Syed Moshiur Rahman, & Nazrul Islam Khan. (2024). Precision Lesion Analysis and Classification in Dermatological Imaging through Advanced Convolutional Architectures. *Journal of Computer Science and Technology Studies*, *6*(5), 168-180. https://doi.org/10.32996/jcsts.2024.6.5.14

[32] Nigar Sultana, Shariar Islam Saimon, Intiser Islam, Abir, S. I., Md Sanjit Hossain, Sarder Abdulla Al Shiam, & Nazrul Islam Khan. (2025). Artificial Intelligence in Multi-Disease Medical Diagnostics: An Integrative Approach. *Journal of Computer Science and Technology Studies*, *7*(1), 157-175. https://doi.org/10.32996/jcsts.2025.7.1.12

[33] Abir, S. I., Shaharina Shoha, Sarder Abdulla Al shiam, Nazrul Islam Khan, Abid Hasan Shimanto, Muhammad Zakaria, & S M Shamsul Arefeen. (2024). Deep Learning Application of LSTM(P) to predict the risk factors of etiology cardiovascular disease. *Journal of Computer Science and Technology Studies*, *6*(5), 181-200. https://doi.org/10.32996/jcsts.2024.6.5.15