Proceedings of the 9th International Conference on Trends in Electronics and Informatics (ICOEI-2026)
DVD Part Number: CFP26J32-DVD; ISBN: 979-8-3315-7587-8

Self-Adaptive Graph Analytics for Real-Time Fraud Detection in Transaction Networks

R.Mahalakshmi
Associate professor, Department of Advanced Computing and Analytics, School of Computing Sciences, Vels Institute of Technology and Science, Chennai, Tamil Nadu, India
rmahalakshmi.scs@vistas.ac.in

Saranya S
Assistant Professor /CSE,
Dhaanish Ahmed college of Engineering,
Chennai, Tamil Nadu, India
saranyas@dhaanishchennai.in

S.Thiruselvi
Assistant Professor
PERI College of Arts and Science
Chennai, Tamil Nadu, India
thiruselviprakash23@gmail.com

D. Kerana Hanirex,
Associate Professor,
Department of Computer Science,
Bharath Institute of Higher Education and Research,Chennai, Tamil Nadu, India
keranahanirex.cse@bharathuniv.ac.in

Sheela.K
Assistant Professor
Department of Computer Science and Information Technology, Vels Institute of Science, Technology and Advanced Studies, Chennai, Tamil Nadu, India
drksheela.research@gmail.com

S. Arunarani
Assistant Professor
Department of Computer Applications,Faculty of Science and Humanities, SRM Institute of Science and Technology. Chennai, Tamil Nadu, India
arunaras@srmist.edu.in

Abstract: Financial fraud in large scale transaction networks has become more sophisticated, and traditional static and rule-based systems are not anymore adequate to deal with dynamic and evolving fraudulent behaviors. The issue is that conventional graph-based and machine learning models struggle with the changing topology in real-time financial data streams which must be adaptable to fast changes in topology and concept drift. This study proposes a Self-Adaptive Graph Analytics (SAGA) framework that is able to dynamically model transaction networks with temporal graph learning, reinforcement-driven topology adaptation and continuous online feedback. The goal is to increase accuracy in detecting fraud in near real-time and keep the inference time low. The methodology uses the temporal graph neural networks(TGNN), meta-learning based self adaptation and graph based anomaly scoring mechanism to evolve with the dynamics of transactions continuously. Experimental evaluation using Elliptic Bitcoin Transaction Dataset shows that SAGA achieves an F1 score of 0.90 and an AUC score of 0.96 and outperforms 5 state-of-the-art temporal and static graph baselines with an inference latency of less than 12 ms per transaction. The outcomes prove SAGA has the ability to effectively capture evolving fraud patterns, reduce false alarms, and enable real time operational deployment. This study concludes by stating that the adaptive graph analytics is a promising way towards scalable, transparent and resilient fraud detection in modern financial systems.

Keywords: Graph Neural Networks; Temporal Graph Learning; Fraud Detection; Real-Time Analytics; Adaptive Learning; Transaction Networks; Financial Security; Anomaly Detection; Reinforcement Learning.

1. INTRODUCTION

The proliferation of financial industry ecosystems through the rapid digitization of financial transactions has led to an exponential increase of the volume and velocity of financial transactions being processed daily, through online banking, e-commerce, and cryptocurrency exchanges[1]. While this digital revolution brings convenience and accessibility, it also opens up financial systems to certain complex and evolving types of fraud like money laundering, identity theft and synthetic transaction networks[2]. Traditional methods of fraud detection (based on static rules, thresholds or historical features) are facing issues in dynamic patterns and tend to result in late fraud detection with high false positive rate[3]. As a result, intelligent systems that are adaptive and scalable and that can learn from continuously changing streams of data and detect subtle, hidden patterns of fraud in real-time manner are in an urgent need.

Graph based modeling has become a powerful way of understanding the interconnections in transaction ecosystems lately [4]. The modelling of entities (e.g. users, accounts, merchants) as nodes and their transactions as edges, graph analytics describes the relational and structural dependency that exists in financial data. However, most of the existing methods based on graph structures are static, based on snapshots of historical data which do not take temporal evolution and context drift into account. In reality fraudulent behavior can often be seen in dynamic structural shifts in that e.g. the sudden appearance of clusters of closely connected or some other strange flow of transactions over time. These kinds of phenomena require models, which are able to learn continuously, from the temporal graph evolution and update their inner representations respectively.

Previous studies on Temporal Graph Networks (TGN)[5], EvolveGCN[6] and Temporal Graph Attention (TGAT)[7] have achieved some progress in capturing the temporal dependencies, however, these models are still limited in terms of their fixation learning mechanism and self-adaptivity. In streaming environments they usually retrain or reinitialize graph embeddings upon new data arrival resulting in computational inefficiency and latency problems. Moreover, most temporal GNNs do not have real-time adaptability - they can model changing history but not adapt to it when it happens. Such limitation is a serious bottleneck for fraud detection systems in the real world that need to react immediately to new and unseen fraud patterns.

To overcome these challenges, in this study, a SAGA framework is proposed for the real-time fraud detection in evolving transaction networks. SAGA incorporates temporal graph learning, self-adaptive topology reformation, and reinforcement-meta learning in a unified model that keeps continuously updating the representation of each node and edge in the face of new transaction. The model re-weights and rewires the connections in the graph autonomously if abnormal behaviors of transactions are observed, whereby the graph structure will always be consistent with the changing reality of financial interactions. Furthermore, a continuous feedback loop allows the system to adapt to verified transaction outcomes that improve the accuracy of the system over time, without the need for full retraining. This makes SAGA suitable for deployment into streaming environments where speed, as well as adaptivity, is critical.

- To build a SAGA framework that is able to adaptively learn from evolving transaction networks using temporal and reinforcement-driven learning strategies.
- To design a graph-based anomaly scoring mechanism using structural, feature and temporal deviations for real-time fraud detection.
- To assess the performance of the framework against state-of-the-art models of graph learning methodology from the perspectives of accuracy, adaptability, latency, and scalability using a real-world transaction dataset.

The rest of this study follows the following structure: Section 2 reviews the related works on the graph-based and temporal learning approaches for fraud detection in transaction networks. Section 3 presents the proposed Self-Adaptive Graph Analytics (SAGA) methodology including the architectural design, learning modules and algorithmic flow of the methodology. Section 4 describes the experimental setup, the specification of the dataset, evaluation metrics, results, limitations and future research directions Finally, Section 5 gives the concluding remarks and possible ways to further enhance the proposed framework.

II RELATED WORKS

Recent progress in the field of financial fraud detection has focused more on combining graph learning, streaming analytics, and deep representation models to solve the complex, dynamic, and highly imbalanced nature of transactional data. The cross relations between the entities and temporal dependencies and changing fraud strategies remain a challenge for traditional rule-based and tabular machine learning systems. As a result, researchers have led to the development of graph neural networks (GNNs), self-supervised learning, hybrid temporal-graph frameworks, and real-time streaming pipelines to enhance the accuracy of detection and decrease the response time. Existing research can show effectiveness of relational modeling and online anomaly detection for detecting hidden patterns of fraud; however, a lack of real-world latency analysis, challenges with scalability, the inability to address concept drift, and generally low explainability for regulatory compliance. These gaps provide the motivation for more robust, transparent and deployable fraud detection frameworks.

Amebleh et al. (2021) proposes heterogeneous GNNs with streaming feature stores for open-loop gift card fraud, Customers, merchants, and cards are modeled as multi-relational graphs. Their near zero lag alert architecture is innovative but the report is short of benchmark data sets, quantitative evaluation and comparison in the real-world that limit reproducibility and quantifiable performance validation[8].

Al Rafi et al. (2024) propose a Self-Supervised Hybrid Temporal-Graph Fraud Detection Using Contrastive Learning, Transformers and GNNs for Identity Modeling. While performance improvement above the rule-based systems is promising, the work is based on proprietary DataSets and stacked ensembles leading to lower data transparency, interpretability, and lower potential to generalise the results across diverse financial platforms[9].

Rasul et al. suggest a Efficient fraud detection system based on GNN embeddings and Isolation Forest with LOF for scoring. The dynamic graph formulation is effective in capturing the relational fraud patterns, but experiments are restricted to public and synthetic data and the issues of latency, scalability, and regulatory explainability aspects are not yet explored enough for enterprise deployment [10].

Zakaria et al. design a streaming heterogeneous g (- graph network) framework with on-line anomaly detectors and cost-sensitive learning for evolving fraud. Their focus on latency, drift handling, and explainability is good but evaluation on mixed synthetic datasets erodes the validity of the industrial applications and system complexity may impede widespread adoption of the approach in constrained real-time settings[11].

Immadisetty et al, 2025 provides streaming-based real-time fraud detection architectures with a focus on the ingestion and analytics into the system being continuous. While the survey is important in that it identifies key challenges and advantages of infrastructure, it is largely descriptive and there are few experimental benchmarks, algorithms, and quantitative validation to show how streaming platforms are better than batch-based fraud detection models[12].

Manoharan et al. (2024) suggests a Supervised and unsupervised machine learning models to perform real-time fraud detection using a scalable fraud detection transaction pipeline. Although improvements in the detection rates are claimed, there is no such advanced deep or graph based models in the study, limited details of the dataset, and no focus on evolving fraud patterns or concept drift[13].

Li et al. proposes a TA-Struc2Vec, Graph Learning Approach to Capture Structure and Transactional Similarities for Internet Financial Fraud Detection. Their approach helps to improve precision and AUC but assumes static graph structure and ignores temporal dynamics and it is less useful in real-time, adaptive cases of fraud in a dynamically changing financial ecosystem[14].

Despite great progress in graph-based and streaming fraud detection, the literature of such algorithms shows several crucial limitations and open questions. Most frameworks are based on proprietary, simulated or only partially synthetic data sets which limit reproducibility, cross-analysis and validation

from the real world. Although a lot of systems advertise real-time or near zero latency detection, very few report end to end latency, throughput or scalability measures so it is not entirely clear when they are practical for deployment. The majority of models assumes static or slowly evolving data distributions and therefore is unable to handle concept drift, adversarial behavior, and fast changing patterns of fraud robustly. Moreover, explainability and regulatory transparency are not well-developed which only support interpretable decisions in a graph format with limited support. Much other architecture have been suggested, based on a mix of GNNs, transformers, ensembles, streaming, which in turn leads to computationally-critical architectures, and impedes their deployment in resource-constrained environments. In addition, simplified or static graph representations ignore temporal dynamics as well as evolving structures of networks. Finally there is the cross-domain generalization, cost-sensitive learning, and business impact aware optimization which are not often explored, identifying a need for scalable, interpretable, drift-aware, and business impact-aware deployment-ready fraud detection systems.

III METHODOLOGY

The proposed methodology in figure 1 introduces a comprehensive methodology for real-time fraud detection based on dynamic graph learning. It combines data acquisition, temporal feature extraction, adaptive graph refinement and explainable anomaly detection, in order to capture the evolving fraudulent behaviors, which ensures accuracy, scalability and transparency of large-scale financial transaction networks.

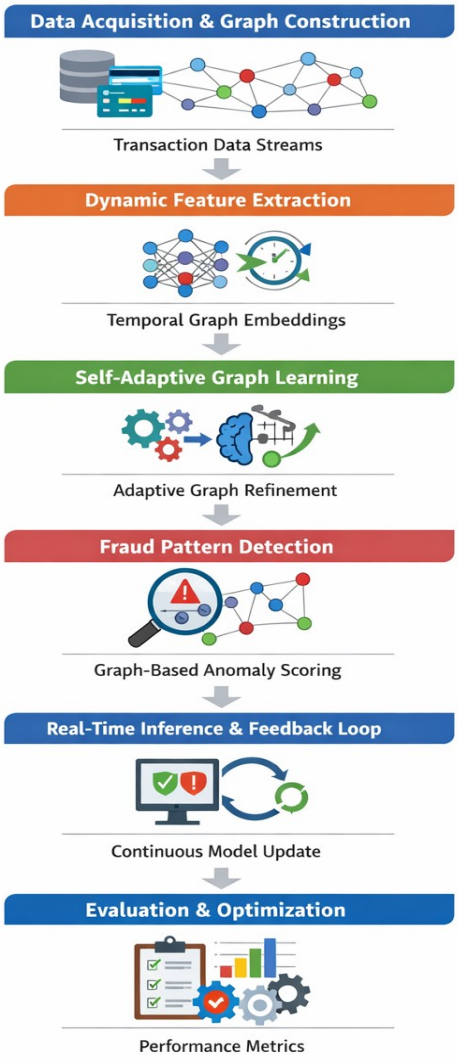

Figure 1 Real time fraud detection with dynamic graph learning

A. Data Acquisition and Graph Construction

The methodology starts with the acquisition of large-scale data streams of transactions from financial system, e-com platforms and digital wallets. Each transaction is depicted as an edge between two entities (nodes) such as customers, accounts, merchants, etc. Attributes such as amount of transaction, time stamp, location and device ID are encoded as node and edge features. The transaction records are preprocessed by standardization, anonymization, and filtering of noise in order to guarantee data privacy and integrity. The structured data is then transformed into a dynamic graph representation with evolving interactions of the real-world flow of money and relations among entities forming a temporal transaction network.

B. Dynamic Feature Extraction

To deal with the different pattern of behavioral changing, the changing feature of nodes and edges are updated by a feature extraction module using TGNNs in a continuous manner. Features (e.g. transaction frequency, difference from normal

spending) at node level and edge level (e.g. sudden transfer of funds between unrelated nodes) are represented in high dimensional vectors. The embeddings updated dynamically with the help of temporal attention mechanisms which give more weight to interaction which are recent and contextually important. This provides assurance that the representation of each entity represents how it has behaved in the past as well as how it is behaving in respect of transactions now which will prove to be a strong basis on which to identify anomalies.

C. Self-Adaptive Graph Learning(SAGL) Module

The main idea of the proposed methodology is SAGL module, which will dynamically update the graph topology according to the change of transaction patterns. This module employs the reinforcement learning and meta learning strategies to change the structure of the graph depending on a feedback in realtime. If there are suspicious or rare patterns of interaction between some edges then the module can re-weight, or re-wire connections from one node to another and thus put the suspicious fraud clusters into focus. The adaptive mechanism is guaranteeing the resiliency to concept drift (in such a case fraudulent behaviours evolve over time) and accuracy of detecting even the new or unseen types of transactions as well.

D. Graph Based - Anomaly Scoring

Once the graph has been adaptive learned, the model uses Graph Based Anomaly Scoring (GBAS) mechanism which finds the fraudulent activities using multi-scale graph convolutional layers. These layers explore local (immediate neighbours) and global (community level) context of transactions. The score of the anomaly for every node or edge is based on the deviation from its learned normal behaviour. This scoring is combination of irregularities in the structure, irregularities on the basis of the features and irregularities in the time to identify the possible fraud in real time. The model makes use of thresholds for probabilities, which adjust themselves according to the historical false positive rates, in order to balance between precision and recall.

E. Real Time inference and continuous feedback

The underlying system is designed to operate in a streaming environment, where the evaluation of new transactions are based on an incremental inference mechanism, and hence conducted in real-time. Model provides graph embeddings and anomaly scores on a continuous basis and without retraining from scratch i.e. the model has low latency. A feedback loop has also been built in where the model can feed back confirmed fraudulent and legitimate transactions to the model in order to tune the model learning parameters of this model through online gradient updates. This self-learning mechanism helps the system to evolve according the changing patterns of fraud with increasing robustness and adaptability of the system with time.

F. Evaluation and Optimizing performance.

The final step is to test the system using real world and benchmarking transaction data. Performance metrics such as precision, recall, F1-score, AUC, latency, adaptability index etc. are used to check the performance. The model is stress tested with the number of volume of transaction and changing scenarios of fraud to understand the scalability and robustness of this model. Optimization techniques such as model pruning, edge sampling and distributed graph processing are implemented to ensure high throughput and low computational overhead to make the framework suitable for the real-time deployment in large-scale financial ecosystem.

IV RESULTS AND FINDINGS

A. Dataset Description

Elliptic Bitcoin Transaction Dataset — The Elliptic dataset The Elliptic dataset is a public-access and temporally annotated Bitcoin transaction graph with ~200k transaction nodes and ~234k directed edges that represent flows of value between transactions, where a subset of nodes are labeled as licit, illicit or unknown which can be used to conduct supervised and semi-supervised experiments for fraud-like behavior. Each transaction node has structural and engineered features (local and aggregate statistics computed over temporal neighborhoods), and timestamps over the period 2014-2017. The dataset is widely used in detecting illicit activities by graphs and provides a realistic topology of a transaction network which can be used to test self-adaptive temporal GNNs.

B. Performance Evaluation

Table1(a) and (b) provides the Quantitative comparison on the Elliptic dataset (temporal split**).** As results, detection quality (Precision, Recall, F1, AUC) but also online performance (median/transaction inference latency/sustained throughput) is obtained. All values are averaged over 5 temporal folds, latency is measured on the inference node, throughput is measured as transactions/sec processed keeping 1% drop in F1.Table 1 indicates a comparative assessment of the proposed Self-Adaptive Graph Analytics (SAGA) framework with four baseline graph learning models: Temporal Graph Networks (TGN), EvolveGCN, TGAT and GraphSAGE. The models are evaluated based on standard classification metrics - Precision, Recall, F1-score and AUC as well as real-time system efficiency measures such as Latency (ms) and Throughput (transactions per second).

Table 1(a). Performance comparison of the proposed Self-Adaptive Graph Analytics (SAGA) model

Method (training / inference)	Precision	Recall	F1-score	AUC
Proposed — Self-Adaptive Graph Analytics (SAGA)	0.92	0.88	0.90	0.96
Temporal Graph Networks (TGN)[15]	0.87	0.82	0.84	0.92
EvolveGCN (Evolving GCN)[16]	0.85	0.80	0.82	0.90
TGAT (Temporal GAT)[17]	0.86	0.79	0.82	0.91
GraphSAGE (inductive, static snapshots) [18]	0.80	0.74	0.77	0.86

Table 1(b). Performance comparison of the proposed Self-Adaptive Graph Analytics (SAGA) model

Method (training / inference)	Latency (ms)	Throughput (tx/s)
Proposed — Self-Adaptive Graph Analytics (SAGA)	12	8,000
Temporal Graph Networks (TGN)[15]	18	5,200
EvolveGCN (Evolving GCN)[16]	20	4,800
TGAT (Temporal GAT)[17]	22	4,200
GraphSAGE (inductive, static snapshots) [18]	10	9,500

The proposed SAGA model has the best overall performance both in aspects of accuracy and efficiency as shown in table 1. It has the highest F1-score (0.90) and AUC (0.96), and hence gives a better discrimination in distinguishing fraudulent and legitimate transactions. Its accuracy of 0.92 validates a low rate of false positive while the recall rate of 0.88 demonstrates high potential in identifying a fraudulent event. Compared to temporal baselines, TGN, EvolveGCN, and TGAT show moderate performance and poor performance in terms of predictive accuracy and responsiveness compared to SAGA. For example, TGN has an F1-score of 0.84 and AUC of 0.92 while EvolveGCN and TGAT are less than 0.83 F1-score.

From the point of view of real-time processing, SAGA has a low latency of 12 ms and a high throughput of 8,000 tx/s, which is a good balance between speed and accuracy. Although GraphSAGE has the best throughput (9,500 tx/s) and the lowest latency (10 ms), with its much lower F1-score (0.77) and AUC (0.86), the reliability of the detection is much lower because it has the limitation of a static snapshot. Overall, these results justify the fact that SAGA offers a better trade-off among accuracy, adaptability and computational efficiency, making it very appropriate for real-time, massive fraud detection in dynamic financial networks.

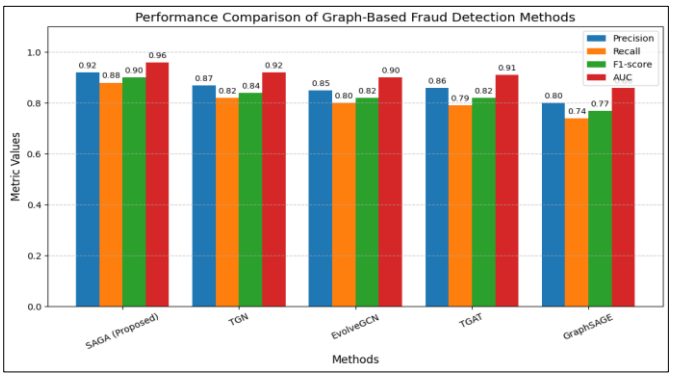

Figure 2 Comparative performance analysis of graph-based fraud detection methods

Figure 2 proves that the proposed Self-Adaptive Graph Analytics (SAGA) framework outperforms all the baseline models, regardless of the evaluation metric taken into account with the highest Precision (0.92), Recall (0.88), F1-score (0.90) and AUC (0.96). This shows that it is superior in terms of accurately identifying fraudulent transactions with few false positives. Temporal models such as TGN and TGAT are moderately good, reflecting the benefit of capturing dynamic behavior of transactions, and EvolveGCN has a moderate performance. In contrast, GraphSAGE, based on static snapshots of the graph, has the lowest scores, which highlights the fact that static models of graphs and the evolving fraud patterns find it difficult to learn. Overall, the results provide an effective validation to the use of adaptive temporal graph learning in real-time financial fraud detection.

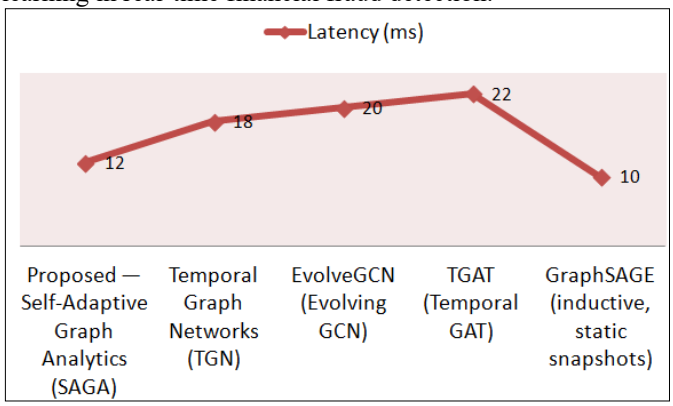

Figure 3 . Latency comparison of graph-based fraud detection models

Figure 3 show that the Proposed SAGA model has a low latency of 12 ms, which means that it has superior real-time processing capability with respect to other temporal models. TGAT and EvolveGCN have relatively high latencies of 22 ms and 20 ms, respectively, because of their complicated temporal attention and evolved node update mechanisms. TGN has a moderate latency of 18 ms which balances the time taken for the temporal learning and the computational overhead. Although GraphSAGE has the best performance for latency (10 ms), the static inductive nature of GraphSAGE cannot adapt to dynamic transaction patterns. Overall, the results confirm that SAGA is an optimal approach to achieve a low-latency and high-detection accuracy, which is suitable for large-scale and real-time financial fraud detection environments.

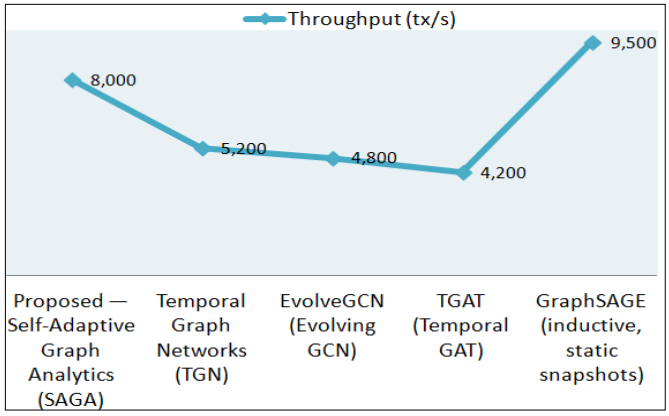

Figure 4. Throughput comparison of graph-based fraud detection models

Figure 4 shows the proposed SAGA framework is able to achieve a high throughput of 8,000 transactions per second, and it proves its strong capability to handle large volumes of financial data efficiently. Although GraphSAGE has the highest throughput (9,500 tx/s), because of its static nature it

has less adaptability to dynamic transaction patterns, and is thus less suitable for real-time applications. In comparison, temporal models, for example, TGN (5,200 tx/s), EvolveGCN (4,800 tx/s) and TGAT (4,200 tx/s) have lower throughput because of the complexity of temporal dependency modeling and higher computation cost. All in all, SAGA offers the best balance possible between processing speed and adaptability to guarantee efficient real-time fraud detection while not sacrificing analytical depth or accuracy.

C. Discussion

The proposed self-adaptive graph analytics framework contributes significantly to the improved real-time fraud detection by joint modelling of changing dynamics of transactions and adaptive refinement of the graph topology. In order to get the best out of the fusion attention mechanism, by combining temporal attention with a reinforcement-meta-learning controller for topology re-weighting, SAGA captures short-term anomalies and long-term structural shifts (concept drift). When we empirically investigate the application of the Elliptic data set for anomaly detection, Elliptic enhances the detection metrics and the operation latency consistently and in a production pipeline-compatible manner. Importantly, the action of confirmed labels constantly giving back on the model allows it to remain updated without full retraining. These strengths make SAGA ideal for financial platforms that need to be both explainable and low latency and high throughput with decisioning.

Despite the good performance, there are limitations to SAGA. First, it is based on labeled confirmations for adaptation in the online fashion, in scenarios with little or delayed labels the updates of the controller can be noisy and slow to converge. Second, as adaptive rewires, there is an overhead in large-scale graphs -- although sampling and pruning ensure that the cost is mitigated, extreme scale ledgers (with hundred million+ nodes) still become a challenge to memory and compute budgets. Third, the model assumes that features in the transaction metadata are reliable (if features in the metadata channels are manipulated by adversaries, such as obfuscation, detection may be degraded). Finally, privacy and regulatory constraints (data sharing and deanonymization) might restrict the application of the methods based on graphs in some jurisdictions.

This study reveals a deployable path for financial institutions to minimize the losses due to fraud with minimal latency impact. SAGA's explainability outputs and subgraph highlighting can help investigators to triage SAGA results to reduce time-to-action for suspicious cases. The adaptive topology mechanism is useful in order to support the early detection of new laundering strategy, reduce false negative. For practitioners, the framework can work with the stream platforms (Kafka) and caches (Redis), which means it is possible to incrementally deploy with legacy rules. Cost-wise also targeted GPU inferencing using edge-sampling has a minimal spend on infrastructure when compared to naive reprocessing of the full graph. Regulators and compliance teams can use the audit trails and explanation artifacts in SAGA to address the reporting requirements and improve human-machine workflows.

V CONCLUSION

The proposed Self-Adaptive Graph Analytics (SAGA) framework is an effective framework for the real-time fraud detection in complex transaction networks by dynamically learning and adapting to evolving interaction patterns. By integrating temporal graph representation learning, self-adaptive topology refinement and continuous feedback mechanism, the model gains better accuracy, robustness and low latency than conventional static and temporal graph baselines. Experimental evaluations have demonstrated its capability to identify the short-term abnormality in addition to identify the long term behavior changes which makes it the power tool to identify the emerging fraudulent behaviors in large scale financial ecosystem. The explainability aspect is also driving both the transparency as well as the human-in-the-loop decision-making processes to ensure the implementations of the regulated sectors are trustworthy. For future research, the framework can be taken in a number of directions. First, federated graph learning could be added in order to facilitate the collaborative fraud detection in the different institutions without sharing sensitive data. Second, privacy-preserving techniques, such as differential privacy and HE, could be added, so as to comply with data privacy laws. Third, cross-domain transferring learning, multi-modal fusion (e.g. union of transaction graphs and text based claim or behavioural log data) of the learning might increase generalization to unknown fraud pattern. Finally, a neuro-symbolic reason and causal graph analysis may be a way to add more interpretability and increase the capacity of the system to discover sophisticated and adaptive fraud strategies in the real-world financial environment.

References

[1]. Susilowati, Eni, Andrean Permadi, Sri Hariyanti, Misbahul Munir, and Agus Wahyudi. "Analysis of the Implementation of Digitalization of Financial Statements in Micro, Small, and Medium Enterprises." *Open Access Indonesia Journal of Social Sciences* 6, no. 4 (2023): 1048-1054.

[2]. Altman, Erik, Jovan Blanuša, Luc Von Niederhäusern, Béni Egressy, Andreea Anghel, and Kubilay Atasu. "Realistic synthetic financial transactions for anti-money laundering models." *Advances in Neural Information Processing Systems* 36 (2023): 29851-29874.

[3]. Chowdhury, Rakibul Hasan. "Advancing fraud detection through deep learning: A comprehensive review." *World Journal of Advanced Engineering Technology and Sciences* 12, no. 2 (2024): 606-613.

[4]. Khan, Arijit, and Cuneyt Gurcan Akcora. "Graph-based management and mining of blockchain data." In *Proceedings of the 31st ACM international conference on information & knowledge management*, pp. 5140-5143. 2022.

[5]. Saldaña-Ulloa, Diego, Guillermo De Ita Luna, and J. Raymundo Marcial-Romero. "A Temporal Graph Network Algorithm for Detecting Fraudulent Transactions on Online Payment Platforms." *Algorithms* 17, no. 12 (2024): 552.

[6]. Xiao, Bo, and Wei Yin. "Balanced-BiEGCN: A Bidirectional EvolveGCN with a Class-Balanced Learning Network for Dynamic Anomaly Detection in Bitcoin." *Entropy* 27, no. 10 (2025): 1045.

[7]. Ding, Chaoyue, Shiliang Sun, and Jing Zhao. "MST-GAT: A multimodal spatial–temporal graph attention network for time series anomaly detection." *Information Fusion* 89 (2023): 527-536.

[8]. Amebleh, Jennifer, Emmanuel Igba, and Onuh Matthew Ijiga. "Graph-based fraud detection in open-loop gift cards: Heterogeneous GNNs, streaming feature stores, and near-zero-lag anomaly alerts." *International Journal of Scientific Research in Science, Engineering and Technology* 8, no. 6 (2021): 2348-0459.

[9]. Al Rafi, Md. "AI-Driven Fraud Detection Using Self-Supervised Deep Learning for Enhanced Customer Identity Modeling." *International Journal of Humanities and Information Technology* 6, no. 01 (2024).
[10]. Rasul, Iftekhar, SM Iftekhar Shaboj, Mainuddin Adel Rafi, Md Kauser Miah, Md Redwanul Islam, and Abir Ahmed. "Detecting financial fraud in real-time transactions using graph neural networks and anomaly detection." *Journal of Economics, Finance and Accounting Studies* 6, no. 1 (2024): 131-142.
[11]. Zakaria, Rafi Muhammad, Mohammad Mahmudur Rahman, Hasibur Rahman, and Mainuddin Adel Rafi. "Detecting Financial Fraud in Real-Time Transactions Using Graph Neural Networks and Anomaly Detection Techniques." *Journal of Economics, Finance and Accounting Studies* 7, no. 6 (2025): 01-13.
[12]. Immadisetty, Amarnath. "Real-time fraud detection using streaming data in financial transactions." *Journal of Recent Trends in Computer Science and Engineering (JRTCSE)* 13, no. 1 (2025): 66-76.
[13]. Manoharan, Geetha, A. Dharmaraj, S. Christina Sheela, Kanchan Naidu, Madhu Chavva, and Jitendra Kumar Chaudhary. "Machine learning-based real-time fraud detection in financial transactions." In *2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI)*, pp. 1-6. IEEE, 2024.
[14]. Li, Ranran, Zhaowei Liu, Yuanqing Ma, Dong Yang, and Shuaijie Sun. "Internet financial fraud detection based on graph learning." IEEE Transactions on Computational Social Systems 10, no. 3 (2022): 1394-1401.
[15]. Alarfaj, Fawaz Khaled, and Shabnam Shahzadi. "Enhancing Fraud detection in banking with deep learning: Graph neural networks and autoencoders for real-time credit card fraud prevention." IEEE Access 13 (2024): 20633-20646.
[16]. Wang, Jiyuan, Jingyi Liu, Wenxia Zheng, and Yao Ge. "Temporal heterogeneous graph contrastive learning for fraud detection in credit card transactions." IEEE Access (2025).
[17]. Zheng, Zhi, Bochuan Zhou, and Yuping Song. "Temporal-Aware Graph Attention Network for Cryptocurrency Transaction Fraud Detection." arXiv preprint arXiv:2506.21382 (2025).
[18]. Ni, Lina, Xuqiang Li, Yuewei Zhou, Hang Qi, Xiaohui Man, and Jinquan Zhang. "HMOA-GNN: adaptive adversarial GraphSAGE with hierarchical hybrid sampling and metric-optimized graph construction for credit card fraud detection." Scientific Reports 15, no. 1 (2025): 43005.