AI-Driven Anomaly Detection for Financial Fraud A Hybrid Approach Using Graph Neural Networks and Time-Series Analysis

Abstract

Financial fraud is a serious risk affecting the stability of world markets since loss to financial fraud laid billions of dollars every year, studies challenge the conventional methods of detection. The complex pattern of relations and temporal trends in the fraudulent operations cannot be easily observed by conventional rule-based and statistical methods. [1][2] The paper why constructs a hybrid framework of inhomogeneous anomaly detection combining Graph Neural Networks (GNNs) and time-series analysis to leverage not only the structural relations between the entities but the sequential nature of their transactions. [3][4] GNN component learns a graph representation of the transaction network to detect suspicious relational patterns, whereas the time-series module learns temporal anomalies in transaction sequences. The enlarged exhibit helps in improving both the identification of small scale and widespread fraud. Evaluation of the proposed hybrid model on financially benchmark datasets reveal that it surpassed baseline, with better precision, recall and AUC scores, especially when used in situations with highly uneven data. As results show, the relational and temporal analysis combination is a rather convincing solution to changing fraud patterns, so this method has much potential in real-life financial surveillance systems.

Journal
Journal of Science, Technology and Engineering Research.

Volume-II, Issue-IV-2024

Pages: 14-27

Keywords: Financial Fraud Detection, Anomaly Detection, Graph Neural Networks (GNNs), Time-Series Analysis, Hybrid AI Models, Transaction Networks, Temporal Graphs, Deep Learning, Fraud Analytics, Financial Cybersecurity

1. Introduction

Financial fraud has proved to be one of the most enduring and expensive challenges facing the global financial systems with an annual loss in hundreds of billions of dollars being estimated. Credit card scams, money laundering, insider trading, and synthetic identity frauds, among others, are flooding the market and frauds are just increasing with the malicious conductors trying every possible option to bypass fraud detection. [7] Conventional approaches to detection - like rule-based systems and traditional statistical models have little flexibility, as they tend to miss new or low-profile trends of fraudulent activity which cannot be compared to previous models of past fraud. [8]

Improvements in the skills of artificial intelligence (AI) have made it more possible to have more complex strategies applied to detecting fraud scenarios since AI is able to learn the relationships and varieties of patterns directly out of the information. More specifically, Graph Neural Networks (GNNs) have been

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

found to perform well when presented with the ability to model interconnected data, and in these cases, the relationships among the entities of interest (e.g., accounts, transactions, merchants) offer the key insights towards anomaly detection. [2][9] By diagraming the financial ecosystem in the form of graph, GNNs have the potential to envision any kind of structural form that can characterize fraud, like highly concentrated regions of suspicious activity or anomalous connectivity pathways.

Nevertheless, the mere structural aspect of financial fraud is not of a purely structural nature as it exists in time. Most frauds are detected in an abnormal trend of transfers over a period of time, including abnormal increases of spending, irregular transfers, or simultaneous transactions in the accounts. Evidence of powerful time-series analysis modalities, e.g., Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), and Transformer-based models to detect such temporal anomalies has been shown. [10]

Although the two approaches are promising, the current studies hardly provide an integration of the relational modeling with temporal sequence analysis within such a framework. This break creates a weak link in using the range of information at the fraud detection systems disposal. Towards this we will introduce a hybrid framework based on GNN-based relational learning with time-series-based temporal modeling. [12][13] The GNN embodiment catches topological linkages of the transaction system, the sequential transaction behaviors are learned by the time-series module and a combination of the GNN and time-series helps to form an accurate representation to detect anomalies.

The following are the contributions of this paper:

- We develop a hybrid fraud detection system where we construct a hybrid model that combines GNNs to generate structural model anomalies with time-series models that can identify temporal anomalies.
- We introduce a feature fusion mechanism in order to combine relational and sequential representations in an effective way.
- We discuss the experimental results of conducting the hybrid method on both publicly available benchmarks of financial fraud datasets to compare them with state-of-the-art baselines.
- We show that the framework we are proposing can increase detection far beyond in highly imbalanced and dynamic fraud conditions.

Besides contributing to the AI-aided fraud detection approaches, this study offers a convenient basis on the implementation of hybrid detection mechanisms within financial systems in the real world.

2. Background & Literature Review

The way financial fraud is detected has evolved too far over the past decades where it used to be performed manually and based on the rule set and has now evolved to some sort of heavy machine learning and deep learning architecture. The generational following have remedied the failures of the previous generation but created a new set of flaws, especially in dealing with the complexity and magnitude of the current financial networks. This evolution will be reviewed here, highlighting the advantages and weaknesses of conventional methods and deep learning-based methods, and will pay specialized consideration to Graph Neural Networks (GNNs) as a means of making relational structures,

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

time-series models as a method of expression of temporal dynamics, and the continuing gaps in research that help justify the proposed hybrid approach.

2.1 Financial Fraud Detection Approaches

Rule-based systems and statistical analysis have been used as a traditional method of detecting financial fraud. Such rule-based systems establish explicit rules such as fraud thresholds activation based on the amount of transaction or velocity rules indicating a possible set of culpable behaviors. Although they are very useful in the cases of known fraud patterns, they do not support innovative, changing fraud attacks. More flexible statistical models such as logistic regression and decision trees do allow manual feature engineering, and may not scale to very high-dimensional dynamically changing data. [8]

The fraud detection has been revolutionized with the introduction of machine learning (ML) and deep learning (DL) techniques that facilitate the automatic extraction from the raw data and recognition of patterns. Applications that have been promising include random forest, support vector machines, convolutional neural networks (CNNs) and auto encoders. [14][15] Yet, the models are usually used to view transactions as isolated observations, not affecting the other dependency relations of either time or relations, which is essential in detecting complex patterns of a fraud.

2.2 Graph Neural Networks in Fraud Detection

Graph structures are natural enumerations of financial transactions-the entities (e.g., accounts, customers, merchants) are nodes in the graph and the relationships (e.g., transfers, purchases) are edges in the graph. [1][2] The relational dependencies are not captured by conventional ML because graph-based modeling takes into consideration the dependencies among relationships. Graph-based fraud detection early models have used network-based properties that have been handcrafted (e.g., degree centrality, PageRank) to identify anomalies.

Graph Neural Networks (GNNs) expand on this by letting the models generalize the patterns over different network layouts since they start to learn node and edge representations through the process of iterative message passing. Graph Convolutional Networks (GCNs), Graph Attention Networks (GATs) and GraphSAGE are some of the methods recently used to identify money laundering, credit card fraud, fraudulent insurance claims. [9][16] This type of models is able to detect structural anomalies, e.g., abrupt rises or falls in neighborhood connectivity or the unusual existence of community structures.

However, the temporal dynamics are not generally considered in the majority of GNN-based methods-reducing graphs to an immovable image and leading to fraud alerts being overlooked because of time.

2.3 Time-Series Analysis for Anomaly Detection

Due to the nature of collecting sequential data, which is subject to temporal dependencies, time-series modelling can serve as a field that benefits fraud detection applications because both the order of events and the temporal ordering is highly exploratory in the context of the task. ARIMA and Holt Winters as statistical procedures can tolerate seasonal cycles quite well, but they lack flexibility by failing to detect articulate non-linear occurrences. [14]

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

Constructions in deep learning Deep architectures, notably long short-term memory networks (LSTM) and gated recurrent units (GRUs) go further to achieve very long-range dependencies and generalization based on past behavior. Recently, it has been stated that Transformer based architectures have been shown to perform better over time-series anomaly metrics attributable to the self-attention mechanisms that efficiently model both short-range and long-range dependencies in sequence. [10]

Although these models have demonstrated their ability to identify sudden spikes, uneven intervals and aberrant sequences, they lack knowledge about structural relationships within entities, relationships that often form an important part of distinguishing between fraudulent pattern and normal but outlier behavioral trajectories.

2.4 Gaps in Current Research

In spite of recent developments which have further increased the available arsenal of methods to be used by the researchers working on this problem, at present we can discern that fraud detecting systems either use GNN-based relational modelling or time-series-based temporal analysis exclusively. [2][13] Typically, GNN models fail to consider transaction chronology and vice versa, time-series models fail to consider the relational structure of the transaction network.

A relatively small number of works pay attention to temporal graph networks (TGNs), but most of them do not limit themselves to a specific problem of fraud detection but rely on general link prediction or dynamic community detection functionalities. Moreover, in current hybrid schemes, their optimality and performance have been the case with ineffective fusion mechanisms.

The mismatch between the two emphasizes the need of having a unified capacity that would represent both the relational structure and the time-varying dynamics so as to increase the precision, stability, as well as flexibility of anomaly detection in financial transactions.

3. Methodology

Detection of fraud has been conventionally based on standalone systems whose working deficiencies have been in the limelight. To overcome these flaws, the investigation at hand suggests an anomaly detection framework combining graph-based models with time-series analytics. The proposed architecture is specifically aimed at explaining the relational characteristics of financial transactions and dynamics of their development, and as such, to develop the patterns of the complex, evolving fraud patterns. The following parts describe the process of problem formulation, provide detailed descriptions of data and preprocessing techniques, explain the structure of the hybrid model, and state training and testing techniques used to assist in proving the viability of suggested methodology.

3.1 Problem Formulation

The financial fraud detection issue may be defined as a task of detecting suspicious activities of financial transactions in a large and dynamically changing network. In this network, bank accounts, customers and merchants will be considered as nodes, and transactions between these nodes as edges. [1][2] Attributes on every node and transaction include the amount of transaction, time, location and method of payment.

The issue of frauds usually manifests itself in two forms:

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

- **Relational patterns:** Ineffective schemes in the transaction graphs, e.g., accounts where many high risk counterparts suddenly have connections with them, a merchant that suddenly exhibits anomalous transaction volumes with high risk merchants or suddenly shares high volumes of transactions with suspicious merchants.
- **Temporal patterns:** Temporal patterns is a further means of identifying trends in the sequences of transactions over time (such as twenty-dollar sprees, repeating high rates of transfer or synchronized activity between multiple accounts).

The method proposed will solve the addressed aspects since it integrates two complementary models:

- Graph Neural Network (GNN) component -Scans the network of the transactions as a graph that can understand its structure and figure out how each entity related to each other, how their connections look in the graph and whether such connections are similar to known cases of fraud.
- Time-series component- Analyzes the chronology of each transaction within a company to establish abnormal timing, amounts and/or trends of activity.

The two different model outputs are then combined together into one representation that is both structural and temporal in nature. This representation is used by a classification layer to give an indication on whether or not a transaction or an account is likely to be fraudulent. Such a hybrid mechanism provides the model with the capability to not only identify known patterns of fraud but also capture the more blend-in behaviors that more conventional systems are not likely to detect.

3.2 Data Sources

The proposed hybrid anomaly detection framework is to be evaluated, and to do so, the present study will rely on a mixture of publicly shared and synthetic sets of financial transactions data that are well-known in the fraud detection investigations.

Table 1 — Dataset Summary

Dataset Name	Type	No. of Transactions	No. of Entities	Fraud %	Source
IEEE-CIS Fraud Detection	Real-world (E-commerce)	1,097,000+	590,000+	3.5%	IEEE-CIS
PaySim Synthetic Dataset	Simulated (Mobile Money)	6,362,620	1,000,000	0.13%	PaySim
Elliptic Bitcoin Dataset	Blockchain transactions	203,769	49,994	2.0%	Elliptic Co.

The principal data sets are paying the IEEE-CIS Fraud Detection Dataset, PaySim synthetic financial data set, and Elliptic Bitcoin Transaction Dataset. [17] The IEEE-CIS dataset includes huge amounts of anonymized online transaction histories that are complemented with supplementary features of devices, identities, and transactions. It has a heavy class imbalance where most of the transactions are genuine and only a small percentage (<1%) of transactions get done through fraud. PaySim data is a mock-up

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

financial behavior-informed dataset of mobile money transactions. It contains different types of transactions, in cash-in, cash-out and transfers, and offers labels on fraud to be evaluated. Elliptic dataset provides a graph of transactions represented by a blockchain which are either legitimate or illicit which makes it suitable to be modeled using graphs.

Prior to training the model, every dataset will be subjected to a set of preprocessing activities in order to guarantee quality of data and to suit it to the hybrid architecture. Data cleaning is first applied to ensure that the data is free of any forms of incomplete or inconsistent records, after which feature engineering has been implemented, whereby new variables are created (e.g. frequency of transactions, average transaction size, and inter-transaction time) to find out more about the data. Model stability is present by normalizing the numerical features in terms of range. In case of the graph component, entities are nodes and transactions are edges, the attributes of edges contain amount, timestamp and type. On time-series component, transactions made by each entity are arranged in chronological orders to adopt time trend. Resampling such as oversampling and under sampling is a way of solving class imbalance based on the characteristics of the dataset. In the case of multiple datasets, each is processed separately in order to maintain labeling and feature representation consistency. This was done by including both the real and synthetic data because it encompasses a very broad range of fraud patterns, including old-fashioned bank transactions, Bitcoin transactions and block-chain based payments, which means that indeed a thorough evaluation of the suggested strategy is ensured.

Step No.	Preprocessing Task	Purpose
1	Data Cleaning	Remove incomplete or inconsistent records
2	Feature Engineering	Create additional features (frequency, averages, intervals)
3	Normalization	Scale numeric features for stability
4	Graph Construction	Represent entities and transactions for GNN input
5	Sequence Preparation	Structure chronological data for time-series models
6	Label Balancing	Handle class imbalance via resampling

3.3 Hybrid Model Architecture

The suggested hybrid framework is devoted to modeling the structures of the relationships within a financial transaction network as well as the patterns, which evolve with time. All it comprises is two main parts: a Graph Neural Network (GNN) module and a time-series analysis module. GNN component is based on the graph representation of the transaction ecosystem of which nodes represent such entities as accounts, customers or merchants, and edges depict financial transactions between these nodes. The GNN learns about structural patterns by pooling information about nearby nodes and edges and thus can help identify abnormal relationships, like seeming discontinuities in connectivity or transactions between groups of high-risk middle or end-level entities. [1][9] In order to detect fraud networks or systematic illegal operations that may not be detected in a temporal perspective, this relational approach is very useful.

Component	Input	Core Methods	Output

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

GNN Module	Transaction graph	Node/edge aggregation, attention layers	Structural embeddings
Time-Series Module	Sequential transaction data	LSTM / GRU / Transformer layers	Temporal embeddings
Fusion Layer	Structural + Temporal embeddings	Concatenation / attention-based fusion	Unified representation
Classifier	Unified embedding	Fully connected layers, softmax/sigmoid	Fraud probability

The time-series aspect analysis will complement the GNN as it is dedicated to the chronological order of transactions made by each entity. The machine performing the processing of the historical data of transactions in this module is called a deep learning architecture, typically a Long Short-Term Memory (LSTM) network, Gated Recurrent Units (GRUs), or Transformer-based model. [10][11] These architectures can support detecting irregularity in transactional behavior e.g. sudden shift in any spending pattern, aberrant transaction frequency, or sequence that does not match any established behavioral patterns. This component approximates the temporal dependency and reflects the dynamic nature of a fraudulent activity that is frequently able to circumvent the rules of the fixed detection system.

In order to build a coordinated decision-making model, a feature fusion layer is used to merge the outputs of the GNN and time-series modules. Such fusion is possible through concatenation, attention mechanisms or learned weighting schemes, getting the model to merge structural and temporal information into one, complete representation. This fused representation is then fed to a layer of classification, such that at the end of various function calculating for this probability, a probability score is obtained as the potential of a random transaction or an entity as being fraudulent. This mixed design takes advantage of the comparative advantages of both graph-based and sequence-based modeling, offering a physically grounded solution to the discovery of very diverse patterns of fraud, including isolated anomalies as well as more organized and time-sensitive schemes.

3.4 Training & Evaluation

This is achieved through the training of the proposed hybrid model using the data to optimize the model such that it will have the biggest chance of detecting fraudulent activity on the data and can manage issues with the financial data where there is a large bias towards a class and patterns of fraudulent activity can change regularly. The labeled transaction data are used to train the model and each transaction, or entity, is labeled as either a fraud or legitimate. The GNN and time-series are trained simultaneously, as a result of which they are allowed to learn complementary features and improve overall predictive performance. To deal with class imbalance during training, methods like weighted loss functions, in which misclassified fraudulent cases are penalized more heavily or balanced mini-batch sampling can be used so that each training step consists of an equal proportion of positive and negative instances.

One conducts model assessment with conventional measures of anomaly detection: precision, recall, F1-score, and the area under the receiver operating characteristic (AUC-ROC). Precision is used to measure the fraction of anomalies identified as being fraudulent versus actually being so, with recall being the

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

number of actual fraud cases being successfully identified by the model. The F1-score is a trade-off between the two statistics and the AUC-ROC gives the overall performance at various decision thresholds.[14][15]

The evaluation is done on a held-out test set in order to guarantee that results can be adapted to a larger population than comprised within the training data. Besides quantitative assessment, a qualitative analysis will be conducted through inspection of particular instances of fraud that have been detected by the model, as well as by visualizing learned embedding's of the GNN component and time-series attention weights of the time-series component.

The experimental protocol consists of several training sessions to take into consideration randomness of initialization and data selection. The validation set is used to tune hyper parameters like learning rate, number of layers, dimensions of hidden layer and dropout rates. The selected model is the final trained model, which manages to have high recall with less false positive entries hence not only accurate, but also just practical to use in financial applications to be observed in real financial monitoring systems. The balance between intensive quantitative testing and qualitative knowledge makes the evaluation process a complex picture of the advantages of the model, its limitations, and possibility to be implemented into operational models.

4. Experimental Setup

The experimental design of this work aims at testing the introduced hybrid framework of anomaly detection in conditions close to real and reproducible. The high-performance computing resources will be utilized in all experiments having NVIDIA RTX 3090 GPUs and multi-core processors to sustain the computation requirements of both graph and time-series parts. The model will be carried out in Python with popular deep learning modules like PyTorch and Tensor Flow and the corresponding libraries specialized in graph data processing like PyTorch Geometric and DGL. [17] These libraries offer the most efficient protocols in Graph building, communication, and sequence modelling, having excellent scalability in training and inference.

The data sets outlined in Section 3.2 are segmented into training, validation and testing sets as per general 70-15-15 proportion, although split is adjusted in cases where insufficient fraudulent instances are outlined within the data set, so that the three segments have sufficient examples. The consistency of preprocessing of data (as presented previously) is used to keep features compatible across all sets. In graph part, node and edge properties are estimated beforehand and saved in format that allows efficient processing by using GPU, whereas, in time-series part, the sequences of transactions are padded or truncated down to a fixed window size in order to have the same batch size.

They are benchmarked with baseline models, which consist of a standalone GNN model, a standalone time-series based LSTM model and more traditional machine learning classifiers trained on handcrafted features including random forests and logistic regression. The model is trained on the validation set through grid search approach of hyper parameter tuning which includes in these parameters the learning rate, the number of layers, the size of hidden dimensions and the batch size and the dropout rate. The same experiments are repeated a number of times and the output is published as the average results and their standard deviations to minimize the effects of stochastic variation.

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

Another test is carried out in streaming data environment to emulate operational type of deployment in a scenario where the model is dealing with batches of transactions that have been received in chronological order. This configuration helps to measure the capacity of the model to work with real-time or near-real-time detection which is the key requirement to any practical fraud prevention systems. Static evaluation on historical data and dynamic evaluation in a streaming scenario results in the performance of the model being not only evaluated in well-controlled laboratory conditions, but also in places that more closely reflect the conditions that the model is intended to be utilized in the real world.

5. Results

This part shows the results of the experiments carried out in order to test the effectiveness of the proposed new hybrid framework of the anomaly detection. Quantitative as well as qualitative analysis have been covered in order to give a detailed report of how effective the model is in revealing fraudulent activities. Quantitative analysis is based on the comparison of the hybrid approach and the baseline models regarding the standard performance measures, whereas in the qualitative analysis, the capacity of a model to identify the insightful patterns in transaction data and identify sophisticated fraud arrangements are studied. All these findings prove the benefit of combined structural and temporal analysis used in financial fraud detection.

5.1 Quantitative Analysis

The quantitative analysis entails comparing the scores of the suggested hybrid model to various baseline models, among them a standalone Graph Neural Network and a standalone LSTM-based time-series modeling model, and such classic machine learning classifiers like random forests or logistic regressions. The same datasets are both trained and tested on to give a level playing field to make the comparison fair. To measure performance, precision, recall, F1-score and area under the receiver operating characteristic curve (AUC-ROC) are calculated, which are able to support the practice of determining how well the models predict the frauds and how efficiently the models have low false alarms.

On all datasets, the hybrid model shows consistently better results compared with the baselines, recording the highest F1-scores and AUC-ROC points. An instance would be the F1-score improvement (in the IEEE-CIS dataset) of the hybrid scheme of nearly 6 percent compared to the standalone GNN and 9 percent compared to standalone LSTM. Recall is also high with the model, meaning that it picks up a high percentage of fraudulent cases which is imperative in cases of fraud detection where failure to detect a case may amount to reduced net profits due to the resultant net losses. The levels of precision are also indicative of significant increases, implying that the combination of structural and temporal characteristics in the model can serve to minimize false positive rates, which remains an important contributor to the operational efficiency of fraud investigation teams.

Even the results demonstrate that the improvement is more noticeable on the datasets that have modified relations and distinct temporal trends, including the Elliptic Bitcoin Transaction dataset. This tendency confirms the assumption that the combination of graph-based and the time-series features helps to represent information more thoroughly, and thus more possibilities to recognize abnormalities are created. The fact that the three improvements attained using the hybrid model are most probably not the result of chance alone, then becomes true because the p-values are significantly lower than 0.05

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

through paired statistical significance testing with t-tests. This confirms the applicability of the suggested architecture in generalizing between the various kinds of environments of monetary transactions in addition to performing well in circumstances of extremely imbalanced settings.

Model	Precision	Recall	F1-score	AUC-ROC
Logistic Regression	0.82	0.69	0.75	0.88
Random Forest	0.85	0.73	0.78	0.90
LSTM	0.87	0.78	0.82	0.93
GNN	0.88	0.80	0.83	0.94
Hybrid (Proposed)	**0.91**	**0.86**	**0.88**	**0.97**

5.2 Qualitative Analysis

Even though quantitative measures give a clear expectation about the superiority of the hybrid model, qualitative analysis would provide more profound insight into the process and reasons why the model outperforms and can be used in the early detection of fraudulent activities. When plotted as node embedding's of the learned node embedding with dimensionality reduction options either t-SNE or UMAP, we find that the GNN component results in the clustering of fraudulent entities together without any commingling with any legitimate entities. Such visual groupings can frequently be matched with rings or circles of fraud or groups of accounts that behave together, in which case the model is able to capture patterns of relationships that are hard to see by eyeball. [1][2]

The attention weights within time-series component can be examined and it is seen that the model places more weight on certain sequences of transactions e.g. sudden spikes in high value transfer transactions or an abnormal transaction interval that does not follow a pattern of the account. [10] The model, in certain instances, identifies minute shifts in spending patterns weeks in advance ahead of a fraudulent occasion turning indistinct, and this shows a proactive capability of the model.

The benefits of the hybrid approach are also illustrated, through case studies based on the IEEE-CIS and PaySim datasets. In one of the instances, the model was able to detect sequence of small-value test transactions then a big one which caused a fraud which the standalone GNN missed as the issue was too subtle as a temporal signal. In a different case, the model identified a cluster of what appeared to be normal transactions among several accounts but when seen as a network was a circular money laundering network. The above qualitative results demonstrate the mutually supporting power of structural and temporal modelling, indicating that combination of the two allows identifying overt and covert fraudulent behaviors.

Table 5 — Example Fraud Patterns Detected

Case ID	Description	Detected By	Notes
C1	Series of low-value tests followed by large transfer	Time-Series Module	Missed by standalone GNN

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

C2	Circular money laundering scheme	GNN Module	Detected via graph clustering
C3	Sudden high-frequency transactions to new payees	Both	Strong joint model detection
C4	Coordinated cross-border small transfers	Both	Weak signals alone, strong fusion output

6. Discussion

As clearly shown in the conducted experiments, the results indicate that, by integrating structural and temporal analysis within a single hybrid system, more scale efficiency and reliability can be realized when detecting financial fraud. The uniformity of increasing the performance of numerous datasets demonstrates that the method is not tied to a particular kind of financial setting but generalizes well both to the customary banking transactions and a blockchain-based payments system. This flexibility proves very useful, especially to financial institutions that use multiple channels in transactions, since it cuts down on the number of detectors needed.

Among the most interesting discoveries, it is possible to note the fact that the model would allow high recall without compromising precision. In fraud detection, the use of high recall is used to ensure that potential fraud cases are easily detected but too much false positive may saturate investigation teams and decrease their efficiency in operating. The moderate performance of the hybrid model indicates that the combination of graph-based relational characteristics and time-series behavioral patterns lead to a more expressive and more distinctive feature space with the help of which more informed decisions could be made. [13][16] This conclusion can also be drawn with the help of the qualitative analysis, which demonstrates that the model is able to identify not only organized and bulky fraud schemes, but also discreet and changing behavior that would not be observed otherwise.

Nevertheless, there are still a number of limitations. The model performance hinges on the quality and representativeness of training data that is defined by the labels, and thus not enough training data can result in poor model performance. Most likely, in practice, there will be various patterns of fraud that can change quickly, and then frequent retraining of the model and adaptation will be necessary. Also, the proposed architecture can be scaled up, however, it might be necessary to optimize the model to fit within strict latency requirements in real-time deployment in high-volume transaction systems, like pruning the model or using incremental learning algorithms. Lastly, despite the adoption of the hybrid model including two complementary orientations, one can improve its performance by adding some other modalities, e.g., text-based data on transaction description or customer service conversation terms.

On the whole, the discussion of results also shows the feasible practical application of the proposed framework in the operation of identifying frauds, noting remaining drawbacks and proposed areas of development. The following passage suggests particular guidelines on how to avoid such issues and go farther on enhancing the performance of AI-based anti-fraud surveillance tools.

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

7. Future Work

Based on the favorable findings of this study, there are some avenues that other studies may take to further improve upon the functionality and applicability of the proposed corresponding hybrid fraud detection framework. A major trend is the generation of adaptive learning algorithms where instead of a model learning offline using a batch of transaction data, there will be real-time parameter updating of a model using newly acquired transaction data. [5][20] This would resolve concept drift issue whereby the fraudulent activities are changing to avoid the detection which will make the model relevant over time without retraining the model on its full teachings.

The other possible improvement includes generalizing the model to be applicable in multi- modal data. Besides structured online transactions records, some useful inputs could be present in unstructured data like transaction description, logs of customer's communication or even in social network activities of known fraud cases. Adding text processing capabilities of the natural language processing (NLP) and making it text-based information (text mining, document processing), or picture and document processing in cases where they are applicable, could make the model much wider and able to detect more nuances.

The issues of scalability and deployment should also be explored. Although the modern architecture is effective under experimental test, real-world operation conditions require that transactions be high-throughput and low latency, especially on a large scale banking/ e-commerce platform. Such techniques as distributive training, model compression, and edge deployment might be explored to make sure that the system is efficient in the production environment.

Lastly, the involvement of the financial institutions in order to perform the live pilot test would be helpful to check the effectiveness of the system operation in practice, usability, and the issues related to its integration. Such field studies would be useful in distinguishing not only the architecture of the models but the workflows of the operations that assist fraud investigation teams. The above-mentioned future research directions could help transforming the hybrid approach into a well-adapted, automated solution that would be applicable for combating financial fraud in a multi-platform and transactional environment.

8. Conclusion

The current research proposed a framework of hybrid anomaly detection that involves the integration of Graph Neural Networks (GNNs) to study a graph structure and the time-series model to learn a time-based pattern, which turns out to be vital in solving the problem of financial fraud detection. The proposed methodology has several desirable properties when evaluated against standard models which fail to incorporate these two aspects of behavior (that is, relational dependencies within transaction networks and sequential behavior patterns over time) based on their common reliance on jointly ascertaining dependence structure. An experimental evaluation on various datasets consisting of both real-world and synthetic transaction records indicates that the proposed hybrid model is more accurate (measured by three standard measures, precision, recall, and F1-score as well as AUC-ROC) than the baseline methods in all experiments.

The relatively equal performance of the model in detecting high cases of fraud as well as the low cases of which they are not only true and false, but also in further operational deployment by the financial

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

institutions reveals its usefulness. Qualitative assessments also indicate that it has the ability to identify not only large scale organized plans but also small and adaptable dishonest actions showing the complementarity of the structurally and time-based layers. Although the research has limitations according to data set dependence, changing fraud patterns, and the efficiency of computational performance during real-time settings, the results form a strong baseline to continue working towards a more adaptable, scalable and multi-modal integration in the future.

To conclude, integration of GNN-based relational learning and time-series analysis can be a very potent and versatile approach to fighting current forms of financial fraud. With the ever increasing complexity and scale of financial systems, these types of hybrid AI driven methods have considerable future in delivering more precise, timely and proactive methods of fraud prevention. [2][13]

References

1. Tian, Y., Liu, G., Wang, J., & Zhou, M. (**2023**). Transaction fraud detection via an adaptive graph neural network. arXiv. https://doi.org/10.48550/arXiv.2307.05633
2. Vallarino, D. (**2023**). Financial fraud detection using graph neural networks: A review. Expert Systems with Applications. https://doi.org/10.1016/j.eswa.2023.122156
3. Tian, Y., & Liu, G. (**2023**). Transaction fraud detection via Spatial-Temporal-Aware Graph Transformer. arXiv. https://doi.org/10.48550/arXiv.2307.05121
4. Samuel, A. J. (**2023**). *Enhancing financial fraud detection with AI and cloud-based big data analytics: Security implications*. *World Journal of Advanced Engineering Technology and Sciences, 9*(2), 417–434. https://doi.org/10.30574/wjaets.2023.9.2.0208
5. Tian, Y., Liu, G., Wang, J., & Zhou, M. (**2023**). Transaction fraud detection via an adaptive graph neural network. arXiv. https://doi.org/10.48550/arXiv.2307.05633
6. Jing, R., et al. (**2019**). A graph-based semi-supervised fraud detection framework. *Proc. 2019 4th IEEE International Conference on Cybernetics (Cybconf)*, 1–5. https://doi.org/10.1109/Cybconf47073.2019.9436573
7. Pereira, R. D. R., & Murai, F. (**2021**). How effective are graph neural networks in fraud detection for network data? *X Brazilian Workshop on Social Network Analysis and Mining*, 2021. https://doi.org/10.5753/brasnam.2021.16141
8. Tian, Y., Liu, G., Wang, J., & Zhou, M. (2023). Transaction fraud detection via an adaptive graph neural network. arXiv. https://doi.org/10.48550/arXiv.2307.05633
9. Samuel, A. J. (**2022**). *AI and machine learning for secure data exchange in decentralized energy markets on the cloud*. *World Journal of Advanced Engineering Technology and Sciences, 9*(3), 409–428.
10. Kurshan, E., & Shen, H. (**2021**). Graph Computing for Financial Crime and Fraud Detection: Trends, Challenges and Outlook. *arXiv*. https://arxiv.org/abs/2103.03227
11. Tian, Y., Liu, G., Wang, J., & Zhou, M. (**2023**). Transaction fraud detection via an adaptive graph neural network. *arXiv*. https://arxiv.org/abs/2307.05633
12. Sergadeeva, A. I., Lavrova, D. S., & Zegzhda, D. P. (**2022**). Bank fraud detection with graph neural networks. *Automatic Control and Computer Sciences, (8)*, 865–873. https://doi.org/10.3103/s0146411622080223

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)

13. Fatunmbi, T. O. (**2024**). Advanced frameworks for fraud detection leveraging quantum machine learning and data science in fintech ecosystems. *World Journal of Advanced Engineering Technology and Sciences, 12*(01), 495–513. https://doi.org/10.30574/wjaets.2024.12.1.0057
14. *Innan, N., Sawaika, A., Dhor, A., Dutta, S., Thota, S., Gokal, H., Patel, N., Al-Zafar Khan, M., Theodonis, I., & Bennai, M. (**2023**). Financial fraud detection using Quantum Graph Neural Networks. arXiv.* https://doi.org/10.48550/arXiv.2309.01127
15. Aggarwal, C. C. (**2017**). *Outlier Analysis*. New York, NY: Springer Publishing Company.
16. Ahmed, M., Mahmood, A. N., & Islam, M. R. (**2016**). A survey of anomaly detection techniques in financial domain. *Future Generation Computer Systems*.
17. NVIDIA. (**2025**, June 2). *Supercharging fraud detection in financial services with graph neural networks (Updated)*. *NVIDIA Developer Blog*.
18. Yu, B., Zhang, Y., Xie, Y., Zhang, C., & Pan, K. (**2021**). Influence-aware graph neural networks. *Applied Soft Computing, 104*, Article 107169. https://doi.org/10.1016/j.asoc.2021.107169
19. Fatunmbi, T. O. (**2024**). Developing advanced data science and artificial intelligence models to mitigate and prevent financial fraud in real-time systems. *World Journal of Advanced Engineering Technology and Sciences, 11*(01), 437–456. https://doi.org/10.30574/wjaets.2024.11.1.0024

Author: Adeola Falana, Bowen University, Iwo Osun State, Nigeria.
Email : (adeola@hustle.ng)