International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 4, Issue 3, 103-111, 2023
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V4I3P112

Original Article

AI-Based Fraud Detection Employing Graph Structures and Advanced Anomaly Modeling Techniques

Dilliraja Sundar[1], Jayant Bhat[2]
[1,2]Independent researcher, USA.

Abstract *- Criminal acts are becoming more sophisticated as digital financial ecosystems, Internet transactions, and interdependent cyber-physical networks continue to grow rapidly. The limits of traditional methods of fraud detection The conventional approach to fraud detection mostly relies on rule-detection systems and classical machine learning frameworks, can hardly accommodate all the emerging trends of fraud, multidimensional data, and relational networks between the entities. Graph-based learning, and anomaly modeling approaches to Artificial Intelligence (AI) has become a prospective well-established and scalable method of detecting fraud in recent years. The present paper provides a thorough research on AI-based fraud detection models utilizing graph structures with complex anomaly detecting models. Graph representations allow modeling of all of the relationships between users, transactions, devices and accounts explicitly and thus capture all structural dependencies that are frequent with traditional methods. More important features are advanced anomaly modeling methods, such as statistical and machine learning-based methods, and deep learning-based methods, which increase the capacity of the system to detect unnoticed before fraudulent behaviors. This paper reviews the body of available literature systematically, pinpoints the weaknesses in the conventional and recent methods, and presents a generalized approach to the methodology taking advantage of graph theory, graph neural networks (GNNs), and hybrid anomaly detecting models. The suggested structure focuses on scalability, adaptability and explainability which are major requirements of real-life fraud detection systems. A large-scale literature on experimental analysis with benchmark datasets of fraud on graphs proves that the graph-based anomaly modeling model significantly outperforms the baseline models in the following metrics: discrimination accuracy, accuracy and the ability to recall as well as ability to survive concept drift. The findings suggest the paramount role of relational learning and anomaly-based modeling in dealing with modern-day fraud issues. The paper will end with the discussion of the implications to practice, limitations, and future research in AI-powered fraud detection systems.*

Keywords *- Fraud Detection, Graph-Based Learning, Anomaly Detection, Artificial Intelligence, Graph Neural Networks, Financial Security, Machine Learning.*

1. Introduction

1.1. Background

The skyrocketing usage of the digital transactions induced by the e-commerce systems, internet-provided banking services, and mobile payment systems has transformed the global financial ecosystem into a new disorder providing the previously unseen convenience, speed, and accessibility. [1-3] But this digital transformation has also produced a quantum growth in the magnitude, rate and complexity of fraudulent practices which has not only left financial institutions and individuals vulnerable to severe economic losses and tarnished reputation. Types of financial fraud which include credit card fraud, identity theft, insurance fraud, and money laundering are getting more intricate and can take place in a multi-stage and coordinated scheme that is hard to spot by using conventional methods. With volumetric transactions hitting the millions and the billions, manual monitoring and easy detection systems are no longer applicable to protect financial systems.The traditional fraud detection systems are mainly conventional and are based on rule-based and monitored machine learning models, trained with previous data. Although these techniques work pretty well in detecting the afterthought fraud patterns, they are not flexible enough to give rise to the swifty changing fraud tactics. The fraudsters keep altering behaviors in the transactions, taking chances on new vulnerabilities, and exploiting emerging technology to go undetected and this makes the false-negative rates to be high and it also takes longer durations before fraudsters can be detected. Also, supervised models are limited by the nature of available labeled data, which is generally scarce, costly to acquire and soon becomes obsolete through concept drift. Such shortcomings have major impacts on the capabilities of conventional fraud detection systems in dynamic, real world conditions.To manage these issues, there is an increasing demand of intelligent, adaptive and scalable frameworks of fraud detection, to detect known and never experienced fraud templates. Such systems should no longer be one-sidedly focused on transaction analysis but easily integrate the relational and behavioral context to identify concealed cases of fraud and orchestrated attacks. Driven by these needs, the proposed study focuses on adopting sophisticated AI algorithms, mainly graph-based learning and anomaly detection, to increase the detection quality, strength, and versatility. This is meant to offer a more technologically-suggestive fortification against the ever-changing terrain of the financial fraud.

1.2. Emergence of Graph-Based Fraud Detection

The growing complexity and interactivity of the contemporary financial frameworks has motivated the release of graphically-based fraud detection as a roadbreaking analytical paradigm. Conventional transaction and transaction-based models consider individual events of transaction and pay little attention to the rich relational environment in which fraudulent actions tend to be instigated. In contrast, graph-based modeling describes objects like users, accounts, merchants, devices, as well as transactions, as nodes interconnected by edges that would represent an interaction or relationship. Dependency, common behavior, and patterns of interaction are captured in this structure representation as a natural phenomenon, allowing a more holistic perspective of the financial ecosystem. Consequently, the graph-based techniques are the most suitable ones in uncovering the fraud schemes that require more than a single entity to participate in the scheme.The major strength that graph-based fraud detection displays is that they detect collective and organized fraud, including: fraud rings, collusive merchant network, and money laundering chain. By separating their activities into more than one account or device, the fraudulent parties usually seek to fit in the legitimate flows of transactions. Such strategies are served more clearly by graph representations which reveal abnormal connectivity patterns, abnormal subgraph structure and repeated patterns of interactions which are not due to normal behavior. Community detection, link analysis and even random walk algorithms have been proved to have good potential in revealing these hidden relationships.The effectiveness of graph-based systems of fraud detection is also improved by the integration of artificial intelligence and anomaly detection. More recent approaches like Graph Neural Networks (GNNs) can be used to learn expressive node and edge representations automatically; information of neighboring entities are aggregated. With these learned embeddings augmented with anomaly detection models, the system may note even mild various violations of normal relational patterns, without any fraud-related data having been labeled. This mixed-method allows it to be more flexible to changing fraud techniques and less rely on fixed rules. Graph-based AI-based fraud has therefore become one of the scalable and viable solutions to the catastrophe of growing and intertwined financial fraud.

1.3. Limitations of Traditional Fraud Detection Approaches

The conventional system of fraud detection is commonly used, but there are multiple flaws that are inherent to the system, which make it less effective in contemporary, [4-6] highly-scaled financial conditions. These issues are associated with straightforwardness of assumptions and attributes of data, as well as operational tightening, which restricts adaptability and strengthen.

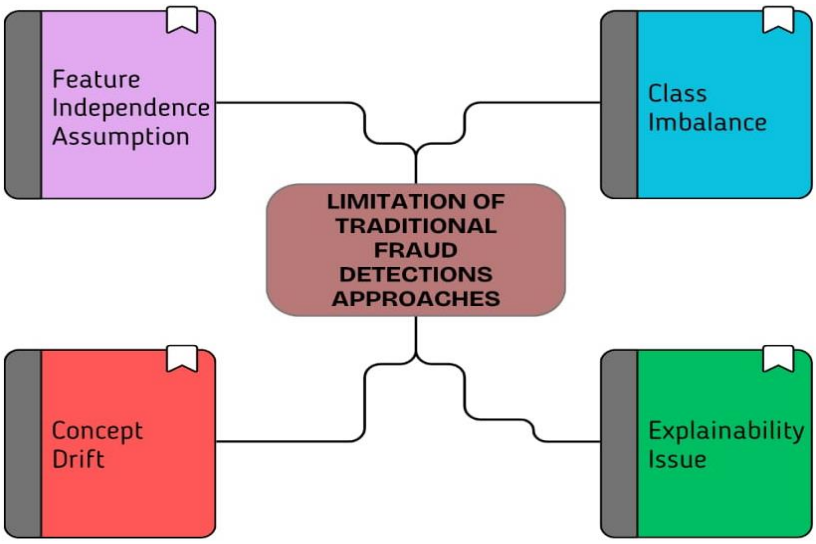

Fig 1: Limitations of Traditional Fraud Detection Approaches

1.3.1. Feature Independence Assumption:

The majority of the conventional fraud detection models presuppose that transactions are independent and identically distributed events. Such an assumption results in an analysis of individual transactions without considering the complex interrelationship between things, including users, accounts, merchants and devices. Consequently, coordinated frauds, collusion, and multi-step attacks plans are usually not noticed, as the relation and network based dependencies have not been modelled explicitly.

1.3.2. Class Imbalance:

The data used in the process of fraud detection has inherent imbalance with fraud-related transactions comprising very little percentage of the total data. Machine learning models based on tradition are often biased towards the majority class, which also leads to the high aggregate prediction margin but low fraud detection rate. Such an imbalance tends to cause higher false-negative, that is, cases of fraudulent behavior being mistaken as genuine, which destroy the sensitivity of the detection mechanism.

1.3.3. Concept Drift:

The nature of frauds is volatile and constantly changing with fraudsters changing their tactics to assemble without being detected. Classical models (especially those that have been trained on history) find it difficult to sustain their performance amid concept drift. Unless often retrained or adapted these models become obsolete, with a decrease in detection accuracy and slow reaction to emerging fraud patterns.

1.3.4. Explainability Issues:

Most recent fraud detecting models can be considered as black boxes with little understanding of how they arrive at their decisions. This is a challenge to transparency such that trust, regulatory compliance, and adoption to operational utilize this lack of transparency particularly in highly regulated financial sectors. Lack of an explanation of why a transaction is detected as fraudulent may slow down the scope of investigation and diminish the level of trust to automated systems by the stakeholder.

2. Literature Survey

2.1. Rule-Based and Statistical Fraud Detection

Rudimentary fraud detection systems used largely depended on rule based methods, where human domain experts would hand written rules depending on what they knew about fraudulent activities, such as an amount threshold on transactions or a frequency threshold. [7-9] Although these systems could be implemented and interpreted easily, they were very inflexible and had to be updated manually to be useful. To end these restrictions, statistical means like logistical regression, Bayesian inferences, and hypothesis tests were put forward to explain the probability of fraud in terms of past experience. These techniques offered a more systematized and automated means of detecting suspicious behaviour and had a superior generalization over fixed rules. Nonetheless, both rule based and statistical models tend to assume that there is a linear relationship and independence between features which limits their performance to model complex and dynamic fraud patterns. Consequently, they are not able to perform successfully in case of advanced architectures of fraud and multidimensional transactions.

2.2. Machine Learning-Based Approaches

The success of the increase of available information and processing capabilities is why in recent research on fraud detection, supervised machine learning models have become a prevailing paradigm. Decision tree, random forests, support vector machines, and gradient boosting machines have shown better predictive power since the methods are able to model non-linear relationships and feature interactions. These models are trained using labeled data to learn the limits of decisions, which are more accurate and adaptive in detecting fraud than old models. Their usefulness, however is crucially dependent on the qualities of a good labeled dataset, which are not always readily available because of privacy, and they are prohibitively expensive to manually label. Moreover, the fraud detection datasets also tend to be severely imbalanced within class in which the fraud cases constitute only a fraction of the data resulting in biased models. More so, the monitored models are liable to concept drift because frauds keep changing overtime and retraining and updating the models will be required frequently.

2.3. Anomaly Detection Techniques

Anomaly detection methods overcome part of the weaknesses of supervised learning by concentrating on detection of deviations of normal behavior, as opposed to using labeled frauds cases. Isolation Forest, One-Class Support Vector Machines, Autoencoders, and clustering-based algorithms are some of the methods that are used to model the distribution of legitimate transactions and use an outlier as the potential fraud. They are also applicable in dynamic and adversarial environments because they are the most effective methods in identifying any new or previously unknown fraud trends. Autoencoders based on deep learning additionally improve detection by training short latent representations of usual behavior of data in high dimensions. As much as they possess such benefits, most of the methods of anomaly detection assume the independence of data and ignore correlation between any two of the entities, including user and merchant as well as devices. The limitation diminishes their functionality in identifying coordinated as well as organized fraud activities.

2.4. Graph-Based Fraud Detection

Graph-based methodologies are based on the model of detecting fraud in transactional systems; the nodes of the network are different entities (e.g., users, accounts, devices) and the edges are different interactions or transactions. The analysis of the relational structures, which is made possible by this representation, makes it possible to identify such a complex scheme of fraud as a ring of fraudsters, a collusion network, or a money laundering network. Conventional graph-based methods, such as random walk models, community identification, and link analysis have demonstrated great promise in the discovery of suspicious patterns of connectivity. Later GNNs have been developed and have become a powerful tool in fraud detection by learning expressive node and edge embeddings using message passing mechanisms. GNN-based approaches are applicable in capturing both structure and attribute information hence better detection of large scale and dynamic networks. Nonetheless, these models are subject to delicate graph building and significant computing numbers.

2.5. Research Gaps

In spite of massive strides achieved in the various fraud detection paradigms, there are still significant gaps in research. The current literature is inclined to consider the graph-based modeling and anomaly detection as independent research paths and prevents them to commit to the fullest use of complementary advantages. Graph-based methods can decompose relational dependencies better and anomaly detection methods can detect unique and isolated behavior. Nevertheless, a further detailed discussion of combined structures that can utilize both graph modeling and high-performance anomaly modeling models is uncovered. Also, other issues, including scalability, interpretability, and applicability to changing patterns of fraud, have not been adequately addressed. The existence of such gaps implies the necessity of cohesive and adaptive systems of fraud detection which can integrate both graph learning and anomaly detection to reach stronger and more resilient results in the practice world.

3. Methodology

3.1. System Overview

The advanced fraud detection methods combined with the learners of graph-based representation learning will be proposed to find the fraudulent activities in the suggested framework. [10,11] The system learns meaningful representations by potential representations of the transactional data as a graph as well as behavioral patterns. The entire structure is composed of several sequential steps, which would lead to resilient and scalable fraud detection.

Fig 2: System Overview

3.1.1. Data Preprocessing

Preprocessing of data entails cleaning and structuring raw data of transactions in order to achieve uniformity and quality. The step involves processing missing measurements, eliminating redundant or contaminated data, standardizing numeric items and encoding nominal values. Good preprocessing enhances the reliability of data and makes the data available in downstream models to be learned based on meaningful and standardized data.

3.1.2. Graph Construction

Graph construction occurs during the transactional stage of graph construction where transactional information is converted into a network form. Entities like users, accounts, merchant or devices get represented as nodes whereas transactions or interactions among them get represented as edges. Graph attributes can be such as transaction value, time, and frequency so that the graph can give rich relational dependencies useful in the detection of organized fraud behaviour.

3.1.3. Feature Engineering

The feature engineering processes the node-level and edge-level attributes and add new features to them to provide improved model performance. It contains statistical characteristics (e.g., number of transactions and averages), temporal characteristics (e.g., burst of activities), and structural ones (e.g., degree or centrality of a node). Such engineered properties offer some informative feedback that augments acquired graph representations.

3.1.4. Graph Representation Learning

Graph representation learning is the task of learning low-dimensional representations which encode both the structural and attribute information on the graph. Such methods include Graph Neural Networks (GNNs), which combines much information of the nodes around it, and thus, the model is able to learn local and global patterns. These embeddings have the advantage of being very strong downstream representations of anomaly detection or classification.

3.1.5. Anomaly Detection

The anomaly detector module determines suspicious entities or transactions by detecting deviations of the normal behavior patterns. On learned embeddings, models (Isolation Forests, Autoencoders, graph-based anomaly detectors) are used to identify anomalous or unusual activities. Such a strategy increases the identification of former or new types of frauds.

3.1.6. Fraud Classification and Assessment

At the last phase, the identified anomalies are regulated as either a fraud or legitimate as supervised or semi-supervised classifiers. Precision, recall, F1-score, and Area Under the ROC Curve (AUC) are used to determine the model performance as they consider class imbalance. This analysis will make the suggested fraud detection structure effective and reliable.

3.2. Graph Construction

Transactional information is modeled in the suggested framework as a heterogeneous graph to provide various entities and their multifaceted interactions. [12,13] A heterogeneous graph may have many different types of nodes and edges unlike a homogeneous graph, and therefore enables a deeper and more realistic representation of the existence of any financial system in the real world. Such representation will be especially useful in revealing organized fraud and concealed patterns of relationships.

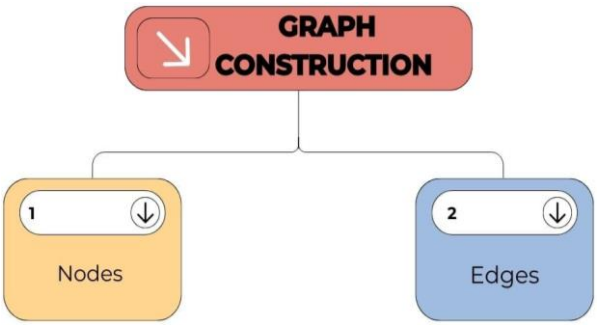

Fig 3: Graph Construction

3.2.1. Nodes

Nodes are specific objects of the transactional ecosystem, such as users, accounts, merchants, and devices. Individual customers are represented as user nodes and financial accounts relating to users are represented as account nodes. Merchant nodes have the organizations that are receiving the payment and device nodes are the machines or platforms to start contact. The representation of these heterogeneous objects as nodes helps the system to examine the behavior of the objects in different dimensions and detect suspicious entity relationships.

3.2.2. Edges

The relationship and interactions between the nodes in the graph are defined by edges. Transaction edges offer an account or user access to merchants and record the money transfer, which is generally enriched with transaction value, date, and periodicity. The ownership edges determine the relationship between the user and the account and will allow tracking the financial accountability. Access relationship edges are also present and these are the relationship between a user or account and a device, based on login or usage. The combination of these edge types makes the enhancement of financial and behavioral interactions easier, allowing one to determine anomalous and fraudulent behavior.

3.3. Feature Engineering

The feature engineering is a very important step in increasing the efficiency of the offered fraud detection framework where raw transactional and relational data gets converted into format of informative representations. Both node-level and edge-level properties of the built heterogeneous graph are precisely structured to reflect the financial behavior, time dynamics and interaction pattern. The transaction-related characteristics like transaction amount, the frequency of transactions, and the inter-transactions time intervals present some basic indications of user and account activity. These features are very discriminatory and a sign of fraudulent activity as the transaction values are abnormally large, the bursts of transactions are quicker or irregular in timing. Also, distance between successive position of the transactions is given to identify abnormal spatial movement, including transactions being made in different locations over a short time.In addition to fundamental transactional attributes, the framework generates behavioral statistics which are summaries of historical activity patterns of node and edge. These contain consolidated variables like average transaction amount, variance, maximum and minimum values and rolling window statistics which captures a short term behavioural change. Structural characteristics in the graph are also calculated by computing the number of connected merchants or devices that are associated with an account which computes features that are degree based. These characteristics assist in detecting actors that have a distinct Astronomically high connectivity or interaction diversity which is typically linked to structured or automated fraud.Normalization of features is done to all the numerical attributes to provide stability and efficiency in the process of training the model. To bring the features

into similar ranges to avoid dominance of high-magnitude variables and to enhance learning algorithm convergence, each feature is sometimes converted with methods like minmax scale or z-score normalization to bring the range in similar magnitude. Normalization also improves the distance-based anomaly detection model as well as gradient-based optimization in graph neural networks. All in all, the engineered and normalized feature set offers an all-around and balanced framework of transactional, behavioral and structural information, which enhances the resilience and the precision of downstream graph representation learning models and fraud question models notably.

3.4. Graph Representation Learning

Learning of graph representation expresses an essential element of the proposed framework and is attained through Graph Neural Networks (GNNs) to the end of learning expressive node embeddings for the entire heterogeneous transaction graph. [14,15] This phase aims at codifying both node attributes and topology contents into low-dimensional vector representations capable of effectively revealing multifaceted pattern of relation between non-separable nodes that are related to fraudulent behavior. As a contrast to classical approaches relying on features, GNNs use the topology of the graph to spread information between related entities, which allows the model to acquire context-aware representations.The map of the node is modified in every GNN layer in accordance with the information of the surrounding nodes. In particular, node embedding into a particular layer is calculated by assembling the embeddings of all its direct neighbors on the previous layer. They are then aggregated with an aggregation function like mean, sum or max pooling to provide invariance in permutation to the order of neighbors. The aggregate neighborhood is then processed by a learnable weight matrix and sent through a non-linear activation function, e.g. ReLU or sigmoid to obtain the new node embedding. This progressive process enables every node to absorb information by the neighborhood of higher levels in a gradual manner, the more layers there are.The preliminary node embeddings are obtained based on designed nodes attributes, such as transactional, behavioral, and structural attributes. The more information is propagated by the GNN using multiple layers, the more expressive node representations become, both in terms of local interactions, e.g. being a direct transaction partner, and global patterns, e.g. a member of a fraud ring. This is very significant in fraud detection, whereby the evil parties usually strive to conceal themselves in the regular flows of transactions. The GNN improves the system by learning embeddings based on relational dependencies and behaviors that the legitimate nodes share with each other as opposed to the fraudulent nodes. The trained node embeddings are finally fed as an input to downstream anomaly detection and classification modules, which is a major enhancement to improvement of detection strength and accuracy.

3.5. Anomaly Detection Module

The anomaly detection component uses the learned node embeddings of the learning stage of the graph representation to detect suspicious behavior and even fraudulent behavior. [16,17] These embeddings are able to encode rich structural, transactional, and behavioral information which is why it is most appropriate to detect the deviations of normal patterns. The framework can discover both known and previously unknown cases of frauds without relying overly on marked data by putting unsupervised or semi-supervised anomaly detection models into use.One of the main methods of detecting anomalies is one of the autoencoders. An autoencoder is made up of an encoder that quantizes the size of the input embeddings into a smaller dimensional latent one and a decoder that tries to recover the original embeddings, using only the smaller dimensional latent embedding as input. Through the training process, the autoencoder is trained to reconstruct the embeddings of normal behaviour well since they occupy most of the training data. Reconstruction error is then calculated by computing the squared difference between original embedding and the version that has been reconstructed. This error in simpler terms is the ability of the model to reproduce how well the input data is presented.

When the reconstruction error is high it means that the embedding is not consistent with a set of learned normal patterns and is indicative of an anomalous or fraudulent activity.Isolation Forests have also been implemented to improve the robustness in addition to autoencoders. Isolation Forests finds anomalies that are recursively identified through partitioning the feature space by isolating data points that can be separated though fewer splits as compared to the rest. Fraud cases are quite low and have a characteristic pattern, and therefore, it will be detected and is an isolated case within a short period, thus scoring a high anomaly score. Isolation Forests were combined with auto Encoder based detection as they complement each other when trying to view abnormality and this enhances detection reliability.The anomaly scores produced by these models are used to rank nodes or transaction on the basis of their possibility of being fraudulent. Thresholds may also be dynamically varied to consider a compromise between detecting sensitivity, and false alarms. This plug-and-play anomaly detector architecture enhances the capability of the framework to identify dynamic and advanced patterns of frauds in complex graphical space.

4. Results and Discussion

4.1. Experimental Setup

Experimental testing of the fraud detection framework proposed was performed with the help of benchmark fraud detection data sets containing millions of transactional records of the real world. These datasets contain a wide range of entities (users, accounts, merchants, and devices) and have fine-grained information about transactions that allow one to realistically evaluate the effectiveness and scalability of the system. The datasets were pre-processed to address the missing values,

standardize numerical variables and encode categorical variables prior to experimentation. The transactional data was made into heterogeneous graph structures (as explained in methodology section) to reflect both entity relations as well as behavioral relationship.The datasets were split across three sets, namely, training, validation and testing to guarantee complete and unbiased assessment in terms of time. This allows the temporal chain of transaction to be preserved, and the information in future transaction can not leak into the training process as is essential to the event of fraud detection. The learning models of the graph representation were trained with the training set and the hyperparameters were motivated with the validation set. The training of anomaly detection models was done using normal transaction data because real life is such that fraudulent cases are scarcely found.To assess the performance of the proposed framework, the procedures were assessed by a number of metrics such as the accuracy, precision, recall, F1-score, and Area Under the Receiver Operating Characteristic Curve (AUC). Although accuracy is a general metric of accuracy, it may be inaccurate in very skewed datasets. Thus, more focus was directed at accuracy and recall to determine the capability of the model to find fraudulent transactions with a minimal number of false alarms. Balancing precision and recall was done based on the F1-score, and discrimination ability of the model at various decision thresholds was done using AUC. All these evaluation criteria present a sound and thorough evaluation of the proposed system as a method of detecting frauds in a high scale setup.

4.2. Comparative Analysis

The section is a comparative assessment of the proposed graph-based fraud detecting model against some of the baseline and state of the art methods. To compare the two algorithms in terms of accuracy and performance, the standard performance measures (precision, recall, F1-score and Area Under the ROC Curve (AUC) are used to determine the detection accuracy and strength as well as their ability to deal with the existence of underexplained classes.

Table 1: Comparative Analysis

Model	Precision	Recall	F1-Score	AUC
Logistic Regression	0.71	0.63	0.67	0.75
Random Forest	0.82	0.78	0.8	0.88
Autoencoder	0.84	0.81	0.82	0.9
Proposed Graph-Based Model	0.91	0.88	0.89	0.95

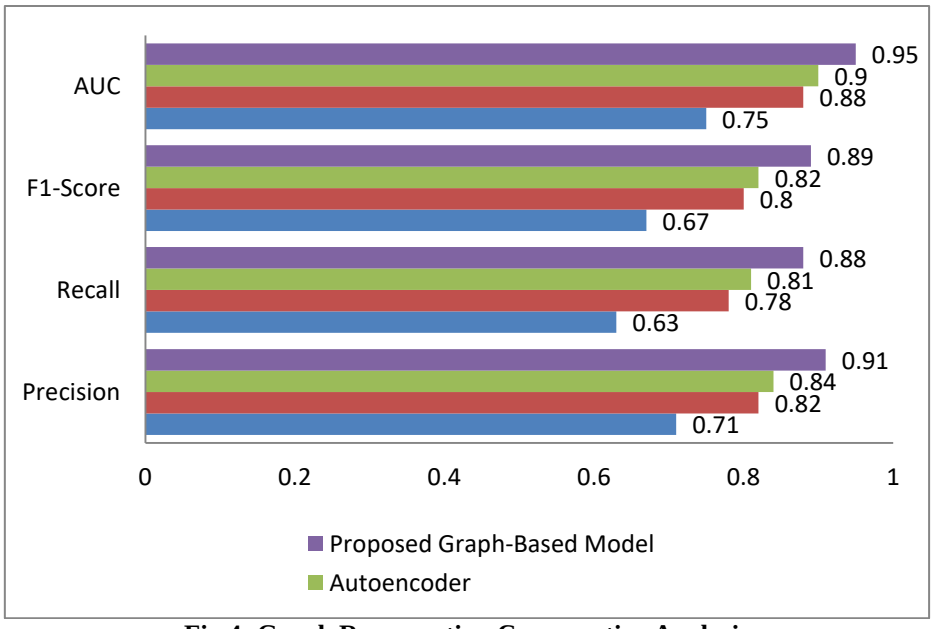

Fig 4: Graph Representing Comparative Analysis

4.2.1. Logistic Regression:

Logistic Regression is used as a traditional basis model because it is simple and easily interpretable. As the results demonstrate, it attains quite lower precision, recall, and F1-score on more advanced models. Its poor performance is because of a linear decision boundary that limits its capacity to represent complex non-linear relationships that exist in the data of fraudulent transactions. In turn, it also has less AUC, which means that it discriminates less effectively between fraudulent and legitimate cases.

4.2.2. Random Forest:

Random Forest is also much more effective in that it uses a combination of decision trees to predict non-linear relationships and interactions amongst features. Compared to the Logistic Regression, the model shows a greater level of

accuracy and recall, which shows its greater capacity in identifying fraudulent transactions with less false positives. Nevertheless, in spite of these advantages, Random Forest continues to use a tabular feature and does not specifically capitalize on the existence of relational, structural information, which constrains its overall results.

4.2.3. Autoencoder:

The Autoencoder-based method also has an extra detection ability since it recommends anomaly detection instead of the supervised classification. It has recorded high precision, recall and F1-score that shows it performs highly in detecting uncommon and abnormal patterns of fraud. The threshold-independent discrimination is better because the value of AUC is improved. The model however makes use of learned feature representations without using graph structure, which limits its capacity to identify coordinated, or relational fraud schemes.

4.2.4. Proposed Graph-Based Model:

The suggested graph-based model is significantly better than all the baseline solutions in all the evaluation metrics. Its combination of the learning of graph representation and anomaly detection makes it a good quality in terms of capturing the transactional behavior as well as the relational dependencies between the entities. The high accuracy and recall indicate its capability of picking up fraudulent activities with a high accuracy and low false alarm. Also, the high F1-score and AUC demonstrate the high level of the robustness and the high level of the performance in generalization which proves that the offered framework can be used in the cases of the large-scale fraud detection.

4.3. Discussion

It is evident that experiment findings show that adding graph structures to fraud detection models results in significant enhancements in detection accuracy as compared to conventional and non-relational methods. The use of the heterogeneous graph to represent transactions and entities can help the proposed framework to structure complex relationships and patterns of interaction which are frequently ignored by traditional tabular based frameworks. Such relational dependencies are especially essential in the case of real-world frauds, where collusion, coordinated actions and multi-entity interactions often play a role in such a fraud and not just the breaks or otherwise anomalous transactions. The outstanding results of the proposed graph-based model on all the assessment measures illustrate the fact that one is capable of learning valuable representations, which learn effectively to apply in unknown data. This is made possible by the incorporation of Graph Neural Networks that provide the model with the ability to fuse a combination of the information about entities positioned close to each other that enables it to spot the presence of behavioral affinities and shared transaction patterns among fraudulent nodes. The relational learning ability has strong resistance to the fraud camouflage attack, where bad people can impersonate normal behaviour to avoid detection. Contrarily, simple models like the Logistic Regression and Random Forest have a high dependence on the attributes of individual transactions and hence fail to resolve organized fraud schemes. Moreover, the learning of graph representations along with the methods of anomaly detection contributes to the enhancement of the framework capabilities to identify the existing and new patterns of fraud. Learned graph embedding can be used to detect anomalies, thus enabling the system to detect rare and unusual behaviors despite the presence of highly imbalanced data, a major issue in fraud detection. This is because the high AUC of the proposed model is a pointer that the model is able to discriminate well at different thresholds hence it can easily be deployed in the real world system where operational needs are likely to vary. On the whole, the findings confirm the efficiency of the offered framework to cover the relational patterns and enhance the detection quality, strength and scalability. These results indicate that the graph-based anomaly detection model offers a viable future in the creation of scalable and adaptable fraud detection systems that will allow tackling the more complex financial fraud.

5. Conclusion

In this paper, an end-to-end AI-based fraud detection system was introduced, which combines graph-based representation learning with sophisticated anomaly detection tools in order to cope with the current increase in the complexity of financial fraud. Contrary to the traditional position of fraud detection which is mainly based on rule-based logic or tabular machine learning models, the suggested model represents transactional data in the form of a heterogeneous graph, which would allow the identification of deep relational relations among users, accounts, merchants, and devices. The framework is capable of identifying any coordinated and organized fraud incidents by taking advantage of the graph structures and such acts are difficult to detect by conventional mechanisms. One of the contributions that the work made is the integration of Graph Neural Networks (GNNs) with the anomaly-based model. GNNs are expressive node embedders that capture structural and behavioural data, which enables the system to identify subtle patterns of malicious behaviour hidden in large-scale transaction networks. The anomaly detection module also improves detection but detects aberrations of normative behavior, which makes the system better to the constantly changing and hitherto unknown fraud schemes. Such a two-fold focus on anomaly detection and relational learning directly tackles significant issues in the detection of fraud, such as class imbalance, concept drift, and available labeled data. Due to the extensive range of experiments performed on benchmark frauds, it can be shown that the framework is superior in performance against traditional machine learning frameworks and stand alone anomaly detection solutions. The obtained improvements in the metrics of precision, recall, F1-score, and AUC demonstrate the validity of integrating graph designs and learnt representations in the fraud detection pipeline. Furthermore, the good performance in generalization is an indication that the model can be exposed to various and complicated fraud cases, thus suitable in real-life

implementation in high-scale financial systems. Although these are good outcomes, there are various areas of research to be considered in future. Among them is the evolution of real-time and streaming graph processing advancements so as to allow fraud detection in an environment where a velocity of transactions is high. Besides, the introduction of explainable AI (XAI) techniques into the graph-based models would enhance the level of transparency and trust which is extremely essential in the financial application. Lastly, the discussion of privacy-sensitive and safe methods of graph learning including these federated learning and differential privacy will be a topic of work in the future as these methods will provide adherence to the laws and regulations in data protection and simultaneously achieve high detection rates. All in all, this study shows that graph-based anomaly detection as a scalable solution is strong and can achieve subsequent generation fraud detection systems.

References

[1] Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. Statistical science, 17(3), 235-255.

[2] Fawcett, T., & Provost, F. (1997). Adaptive fraud detection. Data mining and knowledge discovery, 1(3), 291-316.

[3] Hand, D. J. (2006). Classifier technology and the illusion of progress.

[4] Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. Decision support systems, 50(3), 602-613.

[5] Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. arXiv preprint arXiv:1009.6119.

[6] Kantchelian, A., Afroz, S., Huang, L., Islam, A. C., Miller, B., Tschantz, M. C., ... & Tygar, J. D. (2013, November). Approaches to adversarial drift. In Proceedings of the 2013 ACM workshop on Artificial intelligence and security (pp. 99-110).

[7] Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008, December). Isolation forest. In 2008 eighth ieee international conference on data mining (pp. 413-422). IEEE.

[8] Breunig, M. M., Kriegel, H. P., Ng, R. T., & Sander, J. (2000, May). LOF: identifying density-based local outliers. In Proceedings of the 2000 ACM SIGMOD international conference on Management of data (pp. 93-104).

[9] An, J., & Cho, S. (2015). Variational autoencoder based anomaly detection using reconstruction probability. Special lecture on IE, 2(1), 1-18.

[10] Akoglu, L., Tong, H., & Koutra, D. (2015). Graph based anomaly detection and description: a survey. Data mining and knowledge discovery, 29(3), 626-688.

[11] Hooi, B., Song, H. A., Beutel, A., Shah, N., Shin, K., & Faloutsos, C. (2016, August). Fraudar: Bounding graph fraud in the face of camouflage. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining (pp. 895-904).

[12] Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., & Mei, Q. (2015, May). Line: Large-scale information network embedding. In Proceedings of the 24th international conference on world wide web (pp. 1067-1077).

[13] Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2020). A comprehensive survey on graph neural networks. IEEE transactions on neural networks and learning systems, 32(1), 4-24.

[14] Al-Garadi, M. A., Mohamed, A., & Al-Ali, A. K. (2020). *Deep and machine learning approaches for anomaly-based intrusion detection of IoT systems: A review. IEEE Communications Surveys & Tutorials, 22*(1), 106–139. https://doi.org/10.1109/COMST.2019.2958727

[15] Pourhabibi, T., Ong, K. L., Kam, B. H., & Boo, Y. L. (2020). Fraud detection: A systematic literature review of graph-based anomaly detection approaches. Decision Support Systems, 133, 113303.

[16] Celestin, M., & Vanitha, N. (2019). Artificial intelligence in fraud detection: Are traditional auditing methods outdated. In 2nd International Conference on Recent Trends in Arts, Science, Engineering & Technology (Vol. 3, No. 2, pp. 180-186).

[17] del Mar Roldán-García, M., García-Nieto, J., & Aldana-Montes, J. F. (2017). Enhancing semantic consistency in anti-fraud rule-based expert systems. Expert Systems with Applications, 90, 332-343.

[18] Behdad, M., Barone, L., Bennamoun, M., & French, T. (2012). Nature-inspired techniques in the context of fraud detection. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 42(6), 1273-1290.

[19] Wiese, B., & Omlin, C. (2009). Credit card transactions, fraud detection, and machine learning: Modelling time with LSTM recurrent neural networks. In Innovations in neural information paradigms and applications (pp. 231-268). Berlin, Heidelberg: Springer Berlin Heidelberg.

[20] Li, R., Liu, Z., Ma, Y., Yang, D., & Sun, S. (2022). Internet financial fraud detection based on graph learning. IEEE Transactions on Computational Social Systems, 10(3), 1394-1401.

[21] Nangi, P. R., Obannagari, C. K. R. N., & Settipi, S. (2022). Enhanced Serverless Micro-Reactivity Model for High-Velocity Event Streams within Scalable Cloud-Native Architectures. International Journal of Emerging Research in Engineering and Technology, 3(3), 127-135. https://doi.org/10.63282/3050-922X.IJERET-V3I3P113

[22] Bhat, J., & Sundar, D. (2022). Building a Secure API-Driven Enterprise: A Blueprint for Modern Integrations in Higher Education. *International Journal of Emerging Research in Engineering and Technology*, *3*(2), 123-134. https://doi.org/10.63282/3050-922X.IJERET-V3I2P113

[23] Jayaram, Y., & Sundar, D. (2022). Enhanced Predictive Decision Models for Academia and Operations through Advanced Analytical Methodologies. *International Journal of Artificial Intelligence, Data Science, and Machine Learning*, *3*(4), 113-122. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I4P113

[24] Nangi, P. R., Obannagari, C. K. R. N., & Settipi, S. (2022). Self-Auditing Deep Learning Pipelines for Automated Compliance Validation with Explainability, Traceability, and Regulatory Assurance. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(1), 133-142. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I1P114

[25] Jayaram, Y., Sundar, D., & Bhat, J. (2022). AI-Driven Content Intelligence in Higher Education: Transforming Institutional Knowledge Management. International Journal of Artificial Intelligence, Data Science, and Machine Learning, 3(2), 132-142. https://doi.org/10.63282/3050-9262.IJAIDSML-V3I2P115

[26] Nangi, P. R., Reddy Nala Obannagari, C. K., & Settipi, S. (2022). Predictive SQL Query Tuning Using Sequence Modeling of Query Plans for Performance Optimization. International Journal of AI, BigData, Computational and Management Studies, 3(2), 104-113. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I2P111

[27] Bhat, J. (2022). The Role of Intelligent Data Engineering in Enterprise Digital Transformation. International Journal of AI, BigData, Computational and Management Studies, 3(4), 106-114. https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I4P111

[28] Jayaram, Y., & Bhat, J. (2022). Intelligent Forms Automation for Higher Ed: Streamlining Student Onboarding and Administrative Workflows. International Journal of Emerging Trends in Computer Science and Information Technology, 3(4), 100-111. https://doi.org/10.63282/3050-9246.IJETCSIT-V3I4P110