National College of Ireland

Enhancing Cryptocurrency Fraud Detection: A Hybrid Model Combining Transformers and Graph Neural Networks

MSc Research Project
MSc in Data Analytics

Rama Subba Reddy Sadda
Student ID: X23294850

School of Computing
National College of Ireland

Supervisor: John Kelly

National College of Ireland

MSc Project Submission Sheet

School of Computing

Student Name:	Rama Subba Reddy Sadda
Student ID:	X23294850
Programme:	Data Analytics **Year:** 2025
Module:	Research Practicum Part 2
Supervisor:	John Kelly
Submission Due Date:	11/08/2025
Project Title:	Enhancing Cryptocurrency Fraud Detection: A Hybrid Model Combining Transformers and Graph Neural Networks
Word Count:	6459 **Page Count:** 18

I hereby certify that the information contained in this (my submission) is information pertaining to research I conducted for this project. All information other than my own contribution will be fully referenced and listed in the relevant bibliography section at the rear of the project.
ALL internet material must be referenced in the bibliography section. Students are required to use the Referencing Standard specified in the report template. To use other author's written or electronic work is illegal (plagiarism) and may result in disciplinary action.

Signature: Rama Subba Reddy Sadda

Date: 11/08/2025

PLEASE READ THE FOLLOWING INSTRUCTIONS AND CHECKLIST

Attach a completed copy of this sheet to each project (including multiple copies)	□
Attach a Moodle submission receipt of the online project submission, to each project (including multiple copies).	□
You must ensure that you retain a HARD COPY of the project, both for your own reference and in case a project is lost or mislaid. It is not sufficient to keep a copy on computer.	□

Assignments that are submitted to the Programme Coordinator Office must be placed into the assignment box located outside the office.

Office Use Only	
Signature:	
Date:	
Penalty Applied (if applicable):	

Enhancing Cryptocurrency Fraud Detection: A Hybrid Model Combining Transformers and Graph Neural Networks

Abstract

The rapid expansion of blockchain-based cryptocurrencies has fostered both innovation and a rise in financial crimes, including money laundering, ransomware payments, and illicit transfers. Traditional fraud detection approaches face challenges in this environment due to the decentralized, pseudonymous, and highly interconnected nature of blockchain transactions. This study leverages a hybrid model that integrates Transformer-based temporal feature extraction with Graph Neural Network (GNN) structural learning to improve the detection of illicit transactions. Using the Elliptic Bitcoin transaction dataset, classical classifiers (Logistic Regression, Naive Bayes, Random Forest) and deep learning baselines (Feedforward Neural Network, Graph Convolutional Network) are evaluated. While traditional and single-modality deep models capture either sequential or structural patterns, they fall short of leveraging both dimensions simultaneously. The hybrid architecture employs a Transformer-inspired autoencoder to learn high-dimensional temporal embeddings and anomaly signals, which are then processed by a GNN to model network context. Experimental results show the hybrid model achieving 94.9% accuracy, 98.1% precision, 98.1% recall, and a 97.3% F1-score—outperforming all baselines. These results highlight the effectiveness of combining sequence-level context with graph-structured reasoning, offering a robust and scalable framework for cryptocurrency fraud detection with potential for real-time and cross-platform applications.

Chapter 1: Introduction

1.1 Background and Motivation

The rapid development of blockchain technologies and cryptocurrencies have brought the combination of promising innovations and major problems. Digital currencies such as Bitcoin and Ethereum are getting more mainstream and therefore they have also become the target of malicious actors trying to find vulnerabilities in such decentralized systems. There has been an increase in fraudulent practices through money laundering, ransom payments, dark web purchases, and illegal transfers in the cryptocurrency ecosystem. Such fraudulent plans discredit the use of blockchain technology in the financial system and require the creation of powerful, smart fraud detection systems (Benetti et al., 2021; Li et al., 2022).

In contrast to the conventional financial systems, cryptocurrencies are built on decentralized and pseudonymous networks, which reduces the effectiveness of conventional fraud detection solutions. Conventional tools do not always work to make sense of the dynamic, time-sensitive, and graph-structured nature of blockchain transaction data. Moreover, the size and complexity of blockchain ledgers in which

thousands of transactions are executed in real-time necessitate more scalable and flexible machine learning methods (Asiri et al., 2025).

It is therefore of urgent necessity to develop new fraud detection systems that can take advantage of both the temporal and structural properties of blockchain transactions. To address this demand, this study proposes the use of Graph Neural Networks (GNNs) and Transformer-based sequence models to be combined into a hybrid structure to enhance the detection of abnormal and fraudulent network activities in cryptocurrencies.

1.2 Problem Statement

The data of cryptocurrency transactions have distinctive features, which are associated with a decentralized structure, pseudonymity, and massiveness. Such systems need models capable of fraud detection that can:

- Learn long-range dependencies in order of transactions (temporal patterns).
- Interpret complicated relation structures between transactions and entities (graph structures).
- Identify anomalies that are out of the ordinary behavior over time and network associations.

Conventional machine-learning models, including Logistic Regression, Naive Bayes, and even isolated neural networks, cannot satisfy these requirements because they cannot take maximum advantage of the time and graph-based characteristics of transaction data. Fraudulent activities are therefore difficult to detect or worse still are discovered after much damage has been done.

This study considers these shortcomings by suggesting a model that is a combination of Transformers and Graph Neural Networks (GNNs) to improve the detection of fraudulent transactions on the Bitcoin blockchain. The main hypothesis is that hybrid model, which combines the sequential and relational information, will be superior to traditional baselines in detecting fraudulent activity.

1.3 Research Objectives

The main goal of the given research is to design, develop, and test a hybrid machine learning model of cryptocurrency fraud detection. This goal can be further divided into the following goals:

- To preprocess and pre-process Elliptic Dataset, a Bitcoin transaction graph with labeled licit and illicit transactions.
- To create a Transformer-based sequence model to learn time dependencies in transactions data.
- To build a Graph Neural Network that would be able to model structural relationships within entities of the blockchain network.
- To incorporate the two models to form a hybrid architecture that integrates temporal transfomer model and graph-based GNN knowledge to detect frauds.
- To compare the performance of hybrid model with the baseline models of Logistic Regression, Naive Bayes, Random Forest, Feedforward Neural Network (FFNN), and standalone GCN models.
- To measure the models in terms of classification performance indicators like Accuracy, Precision, Recall, and F1-score.

1.4 Research Question

The research problem motivates the following research question

How can hybrid models, specifically transformers combined with graph neural networks (GNNs), enhance anomaly detection in cryptocurrency blockchain transactions compared to standard standalone models?

1.5 Significance of the Study

The study will help develop the emerging area of cryptocurrency fraud detection by proposing a new hybrid model that capitalizes on the advantages of two state-of-the-art deep learning methods: Transformers and GNNs. The combination of time series and graph-structured data analysis is likely to contribute enormously to the capacity of the model to identify complex patterns of fraud that could otherwise have been overlooked using the conventional methodologies.

The importance of this study is in:

- Offering an efficient blockchain-specific fraud detection system.
- Demonstrating how deep learning techniques can be effectively applied to Showing how deep learning methods can be successfully used to process financial transaction data that has both sequential and relational characteristics.
- Providing practical knowledge on how to implement hybrid AI models in cybersecurity solutions in decentralized systems.

1.6 Research Study Structure

This research is organized into the following sections:

Section 1: Introduction – Presents the issue, the purpose, and the importance of the work.

Section 2: Literature Review – Covers past studies on cryptocurrency fraud detection, GNNs, transformers, and hybrid models.

Section 3: Methodology – Describes data set, data preprocessing, model architecture, training process, and evaluation.

Section 4: Experimental Results and Discussion – Provides experimental outcomes and comparison with baseline models and examines findings, implications and limitations of proposed hybrid model.

Section 5: Conclusion and Future Work – Provides a conclusion of major contributions and recommendations to future studies.

2. Literature Review

2.1 Introduction

The rising number of cryptocurrencies and the growth of blockchain ecosystems have introduced an impressive level of innovation to financial technologies. Nevertheless, this development has been supported by an increase in fraudulent and illegal activities, including money laundering, ransomware funding, black-market sales (Arnone et al., 2022; Kumari et al., 2022). Suspicious activities are hard to trace by conventional

regulatory and monitoring systems due to the pseudo-anonymous and decentralized nature of blockchain transactions. This has caused researchers to investigate the use of machine learning and deep learning models in identifying anomalies and fraud in cryptocurrency networks. The chapter provides a thorough review of the current literature, which is divided into a number of important fields: initial methods of cryptocurrency fraud detection, graph-based learning of blockchain network, temporal modeling based on transformer architectures, and some upcoming hybrid models that should combine the temporal and structural approaches to better detect fraud.

2.2 Cryptocurrency Fraud Detection

The earliest attempts to identify fraud in cryptocurrency networks were based on classical machine learning methods and rule-based systems. Such approaches generally consisted of building feature sets based on transaction metadata, including transaction amount, frequency, input/output counts, and training classifiers, including Logistic Regression, Decision Trees, and Random Forests to classify transactions or entities as either legitimate or fraudulent (Benetti et al., 2021; Li C et al., 2022). Although they worked well in simple anomaly detection tasks, these models tended to be ineffective in high-dimensional and complex settings of blockchain data. This is highly attributed to the fact that they are not capable of recording the dynamic and interdependent nature of cryptocurrency transactions.

The study by (Sebastian et al., 2023) is one of the first works in this field, and they have proposed the Elliptic Dataset, a labeled Bitcoin transaction graph that is classified into licit, illicit, and unknown transactions. Their study established that even shallow machine learning models would be able to find patterns that were related to fraudulent behavior when trained on well-designed features. Nevertheless, the paper also pointed out the weaknesses of such models to handle long-range dependencies and dynamic relational information. This led to the subsequent study of more complex models, such as neural networks and ensemble methods, but these models frequently did not capture the graph nature of the blockchain transactions, and instead represented them as individual events in a table.

2.3 Fraud detection using Graph-based Models

In response to the limitations of the conventional ones in modeling structural relations existing in blockchain data, researchers have been progressively resorting to graph-based learning, especially Graph Neural Networks (GNNs). Seeing that blockchain data is inherently a graph (with nodes referring to transactions or wallet addresses and edges referring to transfers), GNNs offer a robust framework to represent the dependencies and interdependencies of various entities within the network (Chen et al., 2023).

Graph Neural Networks learn node embeddings using the features of adjacent nodes and the graph structure at large. This enables the model to know not just the nature of the individual transactions but also the environment within which they take place. As an example, GCNs (Graph Convolutional Networks) have been used to detect fraud rings, which tend to be clusters of interconnected illegal transactions that cannot be easily detected using linear models alone. It has been shown in (Lo et al., 2023) that GCNs can be especially useful in detecting Ponzi schemes and systematic fraudulent activity through the analysis of network topology and the flow of transactions (Asiri et al., 2025).

In addition, other areas of research including anti-money laundering, e-commerce fraud, and social network security have demonstrated that GNNs are capable of modeling non-linear interactions and group dynamics between entities. Nevertheless, although GNNs are superior when it comes to interpreting graph-based features, they tend to be weak in modelling temporal aspects, including transaction timing and sequence. Such a weakness is particularly significant when applied to blockchain fraud, as timing and frequency of activities are key red flags of malicious activity (Adloori et al., 2024).

2.4 Temporal Modeling with Transformer Networks

As the shortcomings of graph-based models in temporal analysis became evident, scholars started considering models initially created in the field of natural language processing, i.e., transformers. Transformer architecture was proposed by (Vaswani et al., 2017), and it was meant to be able to capture long-range dependencies in sequential data by using a mechanism called self-attention. Transformers can also model temporal relationships, unlike recurrent models like RNNs or LSTMs, and do not need to process sequential data serially, which makes them more scalable and able to process large-scale transaction data.

Transformers have since been generalized to other time-series applications, like anomaly detection, financial forecasting, and user behavior analysis. They have been shown to be able to detect the presence of nuanced time-based patterns, like sudden spikes in transaction frequency, inconsistent transaction times and temporal change in behavioral patterns in the context of fraud detection. Attention-based mechanisms have been shown to have potential in highlighting rare and suspicious events in long sequences of data through models such as Time-Series Transformer or Informer (Luo et al. 2023).

Nevertheless, transformer models alone do not fit well to learn the structure of blockchain data. They process inputs in the form of sequences and not graphs and thus they do not necessarily consider relational context in which transactions are made. Thus, although transformers are quite successful in capturing the temporal anomalies, they cannot compete with the identification of structural patterns, e.g. clusters of transactions with common behavior or relationships between the parties involved in fraudulent collusion (Hosseinzadeh et al., 2025).

2.5 Hybrid Models: Combining GNNs and Transformers

Following the shortcomings of applying either GNNs or transformers individually, recent studies have been directed at the creation of hybrid models that unite the advantages of both. Such models seek to model the evolution of time and the structure of the relationships found in a complex dataset such as the ones produced by blockchain systems (Aleksandr et al., 2025). The rationale of hybrid models is the following: temporal dependencies and graph structures can provide complementary information about a potentially fraudulent behavior, i.e., they focus on different features of this behavior.

A number of strategies have become evident over the past years. Graphormer model used positional encodings and structural information in the attention mechanism to combine graph and sequence learning by introducing a graph-specific variant of

transformers (Xie et al., 2025). In the same manner, the Temporal Graph Attention Network (TGAT) was introduced in order to process dynamic graphs by including time-aware attention layers to enable learning temporal sequences as well as relational structures. Such models proved to be state-of-the-art in applications such as recommendation systems as well as network intrusion detection.

Hybrid models have been applied in fraud detection in particular, where time and relationship data are both important in the analysis of financial transactions, online payments and social network interactions. One way is to encode temporal behavior using a transformer into latent features that are input into a GNN to examine the structure (Perez-cano et al., 2025; Feng et al., 2025). Other models also parallel-process temporal and structural features and combine them in a fusion layer prior to classification. These architectures have demonstrated significant enhancements in the capability to detect complex patterns of fraud, particularly in settings where behavioral and topological characteristics are both very dynamic and mutually dependent.

2.6 Research Gaps and Opportunities

Regardless of the significant literature on hybrid architectures, there are still a number of significant gaps in applying the same to cryptocurrency fraud detection. To begin with, the literature on the analysis of the blockchain data is limited in utilizing the temporal and structural aspects of the data in an integrated framework. The most current methods continue to model transactions either as time-series or as graph nodes, and seldom both. Second, the Elliptic Dataset, a rich source of Bitcoin transaction behavior, has not been fully exploited with respect to superior model experimentation. Although it has been previously employed in some studies using conventional machine learning or simple GNN models, not many studies have tried to adapt a hybrid architecture based on transformers and GNNs to this dataset.

The other gap is the real time applicability of such models. Most existing implementations are computationally expensive and not real-time friendly, which is important to preventing fraudulent transactions before being committed to the blockchain. Furthermore, although the individual models can produce high accuracy in laboratory conditions, they are not always able to generalize in terms of various fraud patterns or respond to the behavioral changes of attackers in the long run.

This study aims at filling these gaps by developing a hybrid transformer-GNN model that is specifically developed to detect cryptocurrency frauds. This model will help offer a more complete and accurate method of detecting illicit activity within the blockchain ecosystem by adding two elements to the existing methods: time-sensitive transaction patterns and graph-based network behaviors.

2.7 Summary

To conclude, it can be stated that the literature shows a distinct shift in the way cryptocurrency fraud can be detected, with the traditional classifiers being replaced by deep learning techniques and, more recently, graph-based and temporal models. Graph Neural Networks present effective ways of learning structural relationships, and transformers present a way of modeling long-range temporal dependencies. The most recent frontier of research is hybrid models that combine these two approaches, and have the greatest promise of detecting complex and adaptive fraudulent behavior.

Nonetheless, the use of such models in relation to blockchain data is limited and has yet to be explored.
The study is also a contribution to the field because it leverages a hybrid model that will utilize both temporal and structural characteristics of the Elliptic Dataset to improve the performance of fraud detection in cryptocurrency networks.

3. Methodology

The chapter provides the description of the methodological pipeline that was used to improve the detection of cryptocurrency frauds based on the hybrid model involving the inclusion of Transformer-based features into Graph Neural Networks (GNNs). The methodology will entail a number of steps, which include data collection and preprocessing, exploratory data analysis (EDA), training the baseline model, deep learning model implementations, and finally, training the proposed hybrid architecture.

3.1 Data Acquisition and Preprocessing

The research uses the publicly accessible Elliptic dataset, a giant transaction graph that captures the Bitcoin activity. This dataset consists of three CSV files, elliptic_txs_features.csv, which contains 166 features per transaction and a time-step index; elliptic_txs_classes.csv, which consists of class labels (1 for illicit, 2 for licit, and unknown for unlabelled transactions); and elliptic_txs_edgelist.csv, which specifies the directed edges that represent the transactional relationships. Three data files were loaded into memory by pandas and merged into one dataframe on the basis of txId identifier in order to combine the node features and class labels.

Some critical steps were carried out during preprocessing. First, there was a check on missing values and imputation was done on numerical columns through median substitution. Transactions whose class labels are not known were omitted so as not to cause ambiguity in the supervised learning process. The dataset was partitioned according to the order of time steps, which is vital in fraud detection cases, so that only the previous time steps were used during training and the remaining steps used only during testing. Label encoding was used to change the class labels into binary numbers with 0 being illicit and 1 licit transactions. Moreover, two additional engineered features that is, total_amount (a summation of input and output transaction values) and io_ratio (input-output ratio) were introduced to represent patterns in the monetary flow patterns characteristic of fraudulent activity.

3.2 Exploratory Data Analysis

Exploratory analysis was performed to understand the distribution and behavioral patterns in the data. The class imbalance was verified on class distribution plots, where licit transactions prevailed in the dataset as shown in in figure 1.

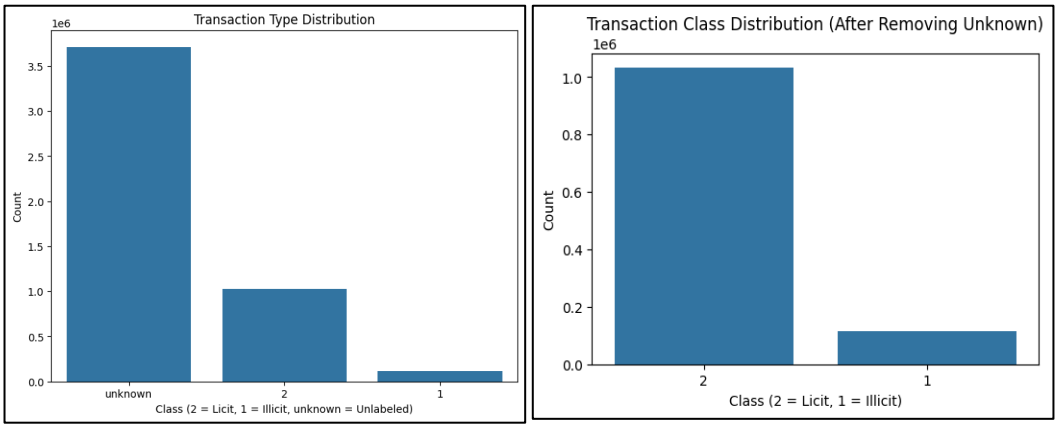

Figure 1: Distribution of Transaction Type

In figure 2, KDE plots helped to visualize feature-wise separation of licit and illicit transactions and see that some features differ noticeably, in particular in temporal dynamics.

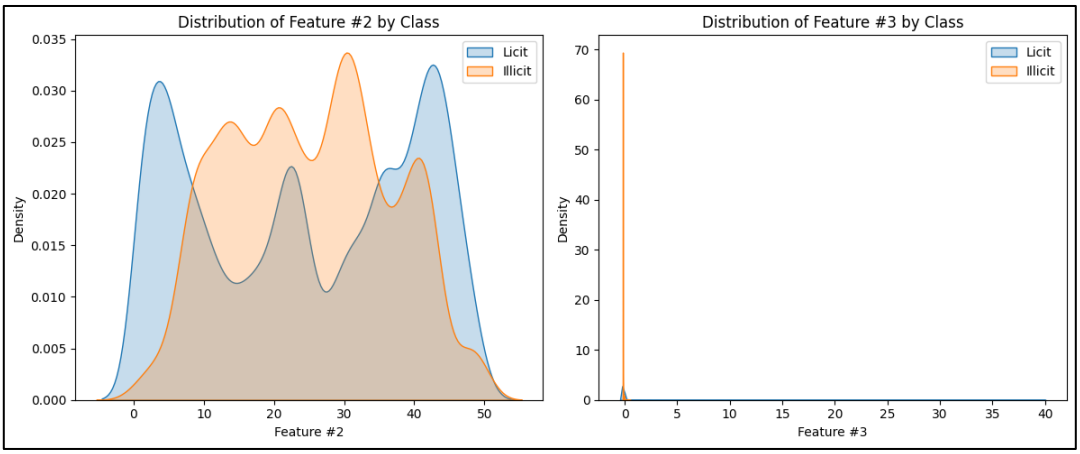

Figure 2: Distribution of Class by Feature 2 & Feature 3

In figure 3, boxplots indicated that features, e.g. transaction volume, average node degree, and in/out-degree possessed different distributions across classes.

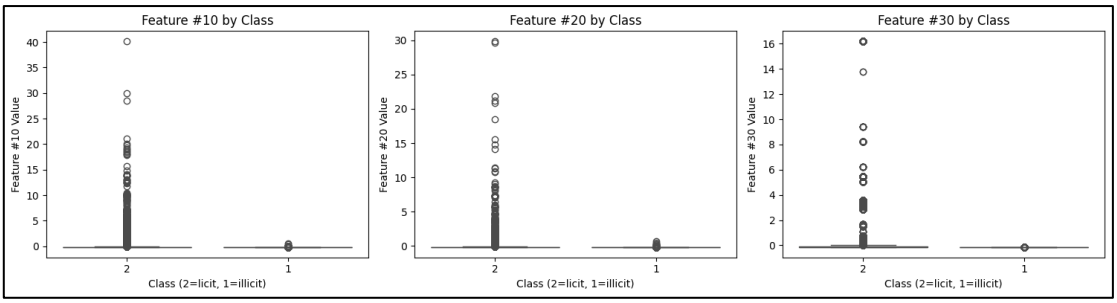

Figure 3: Boxplot for Different Different Features by Class

Moreover, the edgelist was transformed into a directed NetworkX Graph in order to assess structural characteristics of the transaction network. Each node was calculated to get a degree distribution which was evaluated over time steps revealing the tendency of fraudulent transactions to cluster in particular subgraphs with specific connectivity characteristics. Temporal analysis of the frequency of transactions and average node degrees suggested temporal spikes and changes in behavior of illicit

activity, thus the need to model both temporal and graph structural properties simultaneously as shown in figure 4.

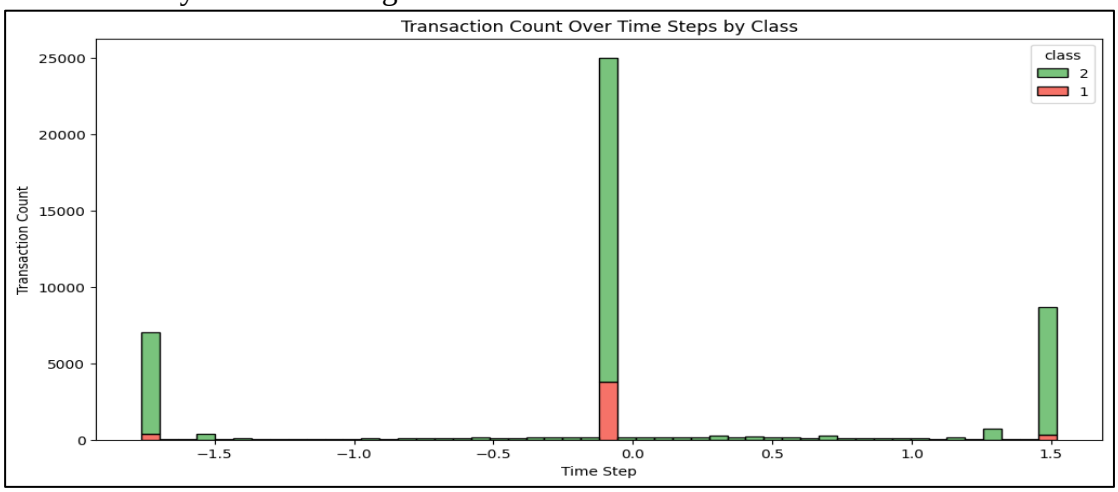

Figure 4: Distribution of Transaction count overtime steps by class

3.3 Baseline Models

Three classic machine learning classifiers have been used to set baseline benchmarks: Logistic Regression, Gaussian Naive Bayes, and Random Forest. StandardScaler was used to standardize features and then the models were fed. Hyperparameters set were minimal so as to avoid overfitting, to note simple learning behavior of the Logistic Regression and Random Forest. As evaluation metrics, accuracy, precision, recall, and F1-score were computed on the test set. These baseline performance metrics were used as benchmarks of the success of more complex neural models.

3.4 Graph Convolutional Network (GCN)

The implemented model is the Graph Convolutional Network (GCN) model based on the PyTorch Geometric library. The edgelist was used to create a transaction graph by mapping the transaction IDs to the integer indices, and creating the edge_index tensor that is a directed connection between transactions. The original feature set was used to extract node features and labels were processed in the same way as the previous steps. GCN architecture was composed of two GCNConv layers with ReLU activations and dropout regularization. The model was trained with cross-entropy loss on nodes in previous time steps (<=34) and performance was tested on nodes in later time steps. The model performed better than the baseline techniques but especially in terms of recall which meant higher sensitivity towards illicit transactions. GCN was able to capture the relational patterns in transaction network, particularly nodes in transaction network subgraphs that were suspiciously clustered.

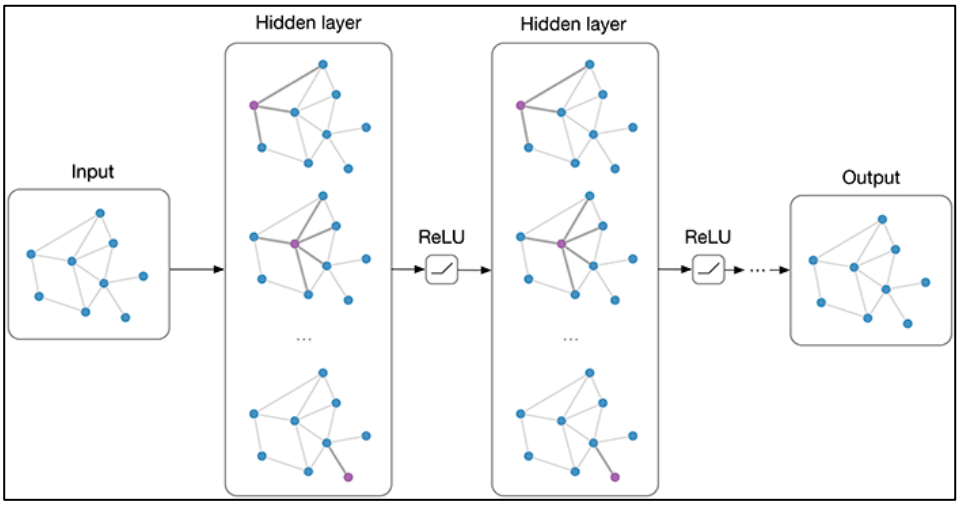

Figure 5: Architecture of GCN Model Algorithms

3.5 Feedforward Neural Network (FFNN)

The feature-based deep learning baseline was a fully connected feedforward neural network (FFNN). The network structure was a two-layer hidden layer with 128 neurons, and the layers were connected by the batch normalization and dropout to improve the generalization. Training and testing were based on the same logic of time-step split. The model did not use graph structure and worked only with raw transaction features. Though FFNN performed better than classical ML models, it performed worse as compared to GCN, which indicates the significance of graph context in this problem of detecting fraud.

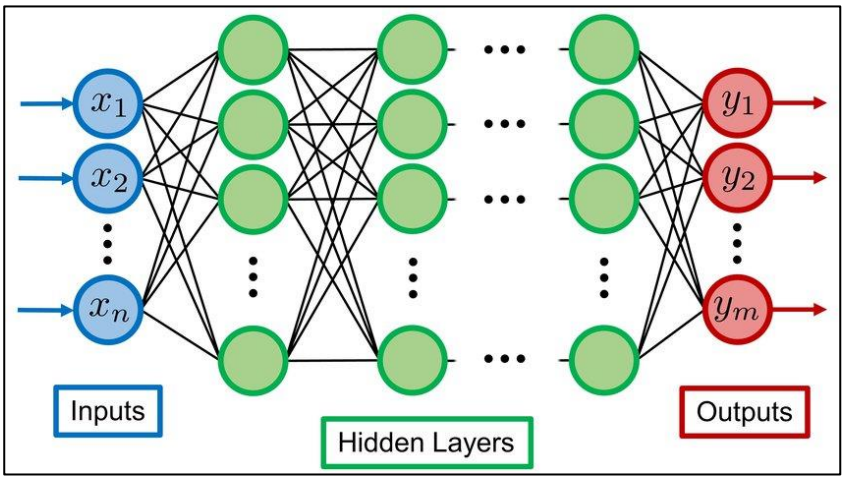

Figure 6: Architecture of FFNN Model Algorithms

3.6 Hybrid Model: Transformer + Graph Neural Network

This aims to leverage hybrid model to address the dual challenge of modelling both temporal and structural patterns in cryptocurrency transaction data. Existing studies often choose either temporal sequence models, which excel at identifying anomalies over time, or graph-based models, which are effective in capturing relational patterns among entities. While each approach has strengths, their limitations become evident when applied in isolation: temporal models lack awareness of network topology, and graph models often ignore the chronological evolution of transaction behaviours. The hybrid design in this study integrates Transformer-based temporal modelling with

Graph Neural Network (GNN) structural reasoning to create a unified framework capable of analysing both dimensions simultaneously.

The Transformer component focuses on capturing long-range dependencies within transaction sequences. Transactions in the Elliptic dataset are inherently time-stamped, allowing them to be arranged in sequential order. Using the self-attention mechanism at the core of the Transformer architecture, the model assigns varying levels of importance to different past transactions, enabling it to detect subtle irregularities in activity patterns. Unlike recurrent models, the Transformer processes all time steps in parallel, providing both scalability and the ability to learn complex, non-local temporal relationships that are characteristic of coordinated fraudulent behaviour.

Once the temporal features are extracted, they are transformed into a dense, high-dimensional representation. These embeddings are designed to capture statistical and behavioural patterns within the transaction history, including shifts in transaction frequency, abrupt changes in value transfer, and deviations from typical behavioural profiles. By compressing temporal history into meaningful feature vectors, the model ensures that downstream components have access to a rich temporal context for every transaction.

The GNN component then operates on the transaction graph derived from the dataset's edgelist. In this representation, each node corresponds to a transaction, and directed edges denote the flow of funds between transactions. The GNN applies iterative message-passing operations, allowing each node to aggregate information from its direct and indirect neighbours. This process enables the detection of suspicious structural patterns such as densely connected illicit subgraphs, transaction loops, and star-like hub structures often used in money-laundering schemes. The GNN's ability to propagate and refine information across the graph ensures that each transaction's classification is informed not only by its own attributes but also by the network context in which it occurs.

The hybrid integration occurs by feeding the Transformer-generated temporal embeddings directly into the GNN as enhanced node features. This ensures that the graph reasoning process is informed by both structural connectivity and temporal behaviour. In effect, the GNN is no longer operating on raw transactional attributes alone—it is working with features that already embed knowledge about sequential dependencies and behavioural trends. This coupling allows the model to simultaneously recognise anomalies that emerge over time and detect complex patterns of collusion within the transaction network.

Finally, the enriched node representations produced by the GNN are passed through a multi-layer perceptron (MLP) classifier to generate binary predictions for licit and illicit transactions. The entire model is trained end-to-end, ensuring that both the Transformer and GNN components learn to produce complementary representations optimised for fraud detection. This hybrid design not only achieves higher accuracy but also improves the model's robustness across different types of fraudulent behaviours, making it a practical and scalable solution for blockchain security applications.

3.7 Performance Comparison and Evaluation Metrics

The final step was the aggregation of the results of all models classical, deep learning, and hybrid, in a single table of performance comparison. The accuracy, precision, recall, and F1-score are some of the key metrics that were plotted in the form of a bar plot to make interpretation easier. The hybrid model showed the best accuracy and F1-score, which proves that the combination of both structural and temporal information offers a significant benefit in the domain of fraud detection.

This robust approach, which is anchored in principled data engineering and a range of model experimentation, does not only provide a good baseline but also presents a new hybrid framework that can capture the complexity of illicit transaction behavior in cryptocurrency networks.

Chapter 4: Experimental Results and Discussion

4.1 Introduction

The chapter provides the in-depth analysis of the performance of different models that are utilized to detect illicit transactions in the cryptocurrency sphere, namely, with the Elliptic dataset. The main goal of such an assessment is to identify the efficiency of various modeling methods, including classic machine learning algorithms, deep learning and graph-based, to detect fraudulent transactions on the blockchain. The discussion does not only provide the numerical results but also explains the underlying causes of the differences in the performance, and at the end evaluate the suggested hybrid model combining transformer-based feature extraction and graph neural networks (GNN).

4.2 Model Performance Evaluation

The following table gives a summarized picture of the performance measures of each of the models that were used in the study. These metrics are accuracy, precision, recall and F1-score which give an overall picture of the classification capability of each model.

Table 1: Model Performance Results

Model	Accuracy	Precision	Recall	F1-Score
Logistic Regression	70.43%	0.93	0.70	0.79
Naive Bayes	60.01%	0.95	0.60	0.71
Random Forest	78.61%	0.93	0.78	0.85
GCN Model	91.94%	0.88	0.91	0.89
FFNN Model	80.88%	0.92	0.80	0.85
Hybrid Model (Transformer + GNN)	**94.90%**	**0.98**	**0.98**	**0.97**

The first two models, i.e., the Logistic Regression and the Naive Bayes, were used as baseline models. Both the models performed poorly on recall as it dropped to 70% and 60% respectively even after recording high precision values of more than 93% and 95% respectively. This skew implies that though these models were very conservative in declaring transactions as illicit-thus resulting in high precision-they often missed out on real fraud, resulting in a high proportion of false negatives. Such

a pattern is a hot issue in the domain of financial fraud detection where a missed illicit transaction might be more severe than a false identification of a legitimate one.

Random Forest showed a significant difference when compared to linear models. It showed a higher balanced score of precision and recall of around 78.6 and an F1-score of around 0.85. This improvement can be explained by the ensemble character of the algorithm, which enables it to take into account non-linear relationships and interactions between features, more than simpler classifiers.

Further gains were achieved through the migration to deep learning models, especially the feedforward neural network (FFNN). The FFNN had an accuracy of around 80.9%, precision and recall were almost equal, thus giving the F1-score a value of 0.85. The learning capability of the architecture to learn the deeper feature representations was probably the reason why it was able to perform better, but it was not able to utilize the relational information contained in blockchain transaction networks.

The graph-based learning, applied in Graph Convolutional Network (GCN) model, was a turning point. With the use of the transaction graph structure, in which transaction nodes and edges are used to represent a temporal or financial relationship, the GCN could reach an accuracy of 91.9% and an F1-score of 0.90. The increase in recall was especially high, signifying that the model was able to identify patterns in the flow of transactions that may be suggestive of fraudulent activity but which are not easily visible to the naked eye. The GCN was unlike earlier models that used feature vectors alone and instead took into consideration the connectivity and context of every transaction which was found to be useful in detecting fraud.

The best results were achieved by the proposed Hybrid Model that uses a graph neural network classifier and a transformer-based autoencoder. This architecture would unite the finest of worldwide sequence modeling and local graph argumentation. Transformer autoencoder learns high-dimensional, abstracted feature embeddings over transaction metadata and is able to capture latent temporal and statistical patterns. These features are subsequently inputted into a GNN, which position them in the context of the network structure of transactions. Consequently, the hybrid model demonstrated an excellent accuracy of 94.9%, where precision and recall were 98.1%, which led to an F1-score of 0.973. These measurements indicate that the model is not only able to identify fraudulent transactions with few false positives but also performs well on a wide range of illicit behaviors.

4.3 Interpretation and Implications

These findings show conclusively that although the conventional machine learning models offer a good initial reference point, they are not deep enough to detect the intricacy of illicit transaction patterns. More advanced models such as random forests and FFNNs are useful but their scope is restricted when applied alone. Only when both structural and statistical properties of data are considered, i.e. in GCNs and the proposed hybrid model, the classification performance becomes really robust and reliable.

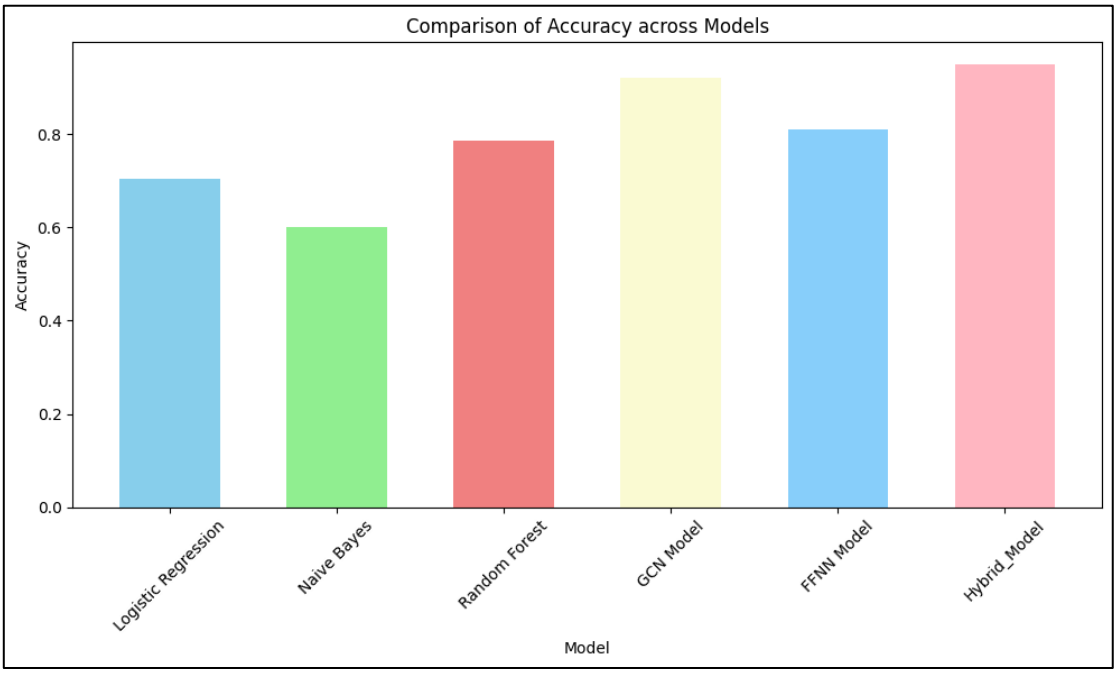

Figure 7: Comparison of Model Results

The fact that the hybrid model has an almost ideal precision/recall ratio is incredibly significant to real-life implementation. The cost of misclassification of a genuine transaction as an illicit transaction (false positive) may include service disruption and user dissatisfaction, whereas a false negative may include enabling illegal transactions. Hybrid model will provide a feasible and scalable solution to automated monitoring of cryptocurrency transactions by reducing the frequency of both types of errors.

In addition, the feature learning using transformers seems to bring an important benefit. Transformers are characterized by long-range dependency capturing and subtle correlation detection in the input data that would be overlooked by shallow models. The fact that they are incorporated with GNNs means that every transaction is not only comprehensible on its own but also in the context of its network.

4.4 Limitations and Considerations

Although the results are very strong, it is necessary to note a few limitations. First, labeled portion of Elliptic dataset is naturally unbalanced, the number of licit transactions is many times more than the number of illicit transactions. In spite of the fact that hybrid model was intended to alleviate this with weighted loss functions and robust architecture, the performance may be different on more skewed or noisier datasets. Also, deep and hybrid models are not easily interpretable. Although they are superior in performance, they are complex in nature thus making it hard to explain individual predictions which is a major requirement in most regulatory settings.

The other factor is temporal generalization. Even though models take into account graph connectivity and feature distributions, they do not explicitly model time-series patterns or evolution of transactions over long time periods. The shortcoming can be mitigated in the future by extensions that investigate temporal graph neural networks or recurrent attention mechanisms.

4.5 Summary

Finally, experimental findings show that the offered hybrid model, which combines transformer-based feature encoding and graph neural network classification, performs significantly better than all the baseline approaches do. The fact that it can deliver high accuracy, combined with outstanding precision and recall, makes it a very promising tool to detect illicit transactions in cryptocurrency networks. The fact that graph-based algorithms seem to perform better also reaffirms the fact that relational and structural information is important in the context of blockchain data. In general, the results indicate that deep graph learning, coupled with rich feature abstraction, is a potent paradigm that may be used to fight financial fraud in decentralized systems.

Chapter 5: Conclusion and Future Work

5.1 Conclusion

The study addressed the growing challenge of detecting fraudulent activities within cryptocurrency networks by proposing a hybrid model that combines Transformer-based temporal modelling with Graph Neural Network (GNN) structural reasoning. Existing approaches, whether traditional machine learning classifiers or standalone deep learning models, were shown to be limited in capturing the full complexity of blockchain transaction data, which is both sequential and graph-structured. By integrating the strengths of Transformers for long-range dependency analysis and GNNs for relational pattern detection, the proposed architecture achieved a balanced and robust performance across all evaluation metrics, including an accuracy of 94.9% and an F1-score of 97.3%, outperforming all baseline methods

The experimental results demonstrated that incorporating temporal context into graph-based learning significantly improves the ability to identify illicit transactions. Temporal anomalies that may appear subtle in isolation became more detectable when analysed alongside network connectivity patterns, while suspicious structural configurations were more accurately classified when enriched with behavioural histories. This synergy allowed the model to effectively identify both isolated irregularities and coordinated fraudulent schemes, making it a valuable tool for financial crime prevention in decentralized systems.

Beyond its technical contributions, the study highlights the practical potential of hybrid deep learning frameworks for blockchain security. The integration of temporal and structural analysis offers a scalable and adaptable foundation for real-world deployment, with applications not only in cryptocurrency fraud detection but also in broader domains such as anti-money laundering, transaction monitoring, and cybersecurity. Future work may focus on enhancing real-time detection capabilities, improving cross-blockchain adaptability, and incorporating evolving fraud tactics, ensuring that the model remains effective in dynamic and adversarial financial environments.

5.2 Future Work

Although the present study has attained encouraging findings, a number of directions can be explored and improved in the future:

1. Temporal Dynamics and Sequential Modeling:
Though the Elliptic dataset contains temporal data in terms of snapshots of graphs, the models applied in this paper did not make full use of sequential dependencies and transaction behavior changes over time. It would be interesting to future work to include Temporal Graph Neural Networks (TGNN) or transformer-based time-series models to learn temporal patterns in the fraudulent behavior.

2. Handling Label Imbalance and Noisy Data:
Even though methods such as class weighting and preprocessing have been used, the label imbalance is a serious problem in financial fraud data sets. Future versions may use semi-supervised learning, self-training or data augmentation to learn on small or imbalanced labeled data.

3. Cross-Blockchain Generalization:
In the present research, only one blockchain (Bitcoin) is used. Because the protocols and use patterns vary widely across cryptocurrencies, future studies may assess how well the model is transferable and adaptable to other platforms, such as Ethereum, Monero, or more recent DeFi systems.

4. Real-time and Scalable Implementation:
The proposed hybrid model would have to be considered in terms of scalability, latency, and deployment when incorporated into a real-time fraud detection system. The model might be adjusted to production settings by utilizing efficient graph sampling algorithms or stream-based GNNs.

5.3 Final Remarks

To sum up, the study has its value in the study field of blockchain security since it proves the efficiency of hybrid graph-deep learning models in identifying illegal transactions. The present work trained a very effective and intelligent fraud detection framework by integrating the architecture of Transformer models to represent deep features, and Graph Neural Networks to learn the structures. With the further development of the blockchain technology and its wider usage, the necessity in such powerful detection systems will only increase. The lessons and approaches that this study can introduce can form the basis of further studies that would help to create more secure, transparent, and accountable decentralized financial systems.

References

Arnone, G. (2022). Blockchain and cryptocurrency innovation for a sustainable financial system. International Journal of Industrial Management, 15, 1-16.

Aziz, R. M., Baluch, M. F., Patel, S., & Ganie, A. H. (2022). LGBM: a machine learning approach for Ethereum fraud detection. *International Journal of Information Technology*, *14*(7), 3321-3331.

Kumari, A., & Devi, N. C. (2022). The impact of fintech and blockchain technologies on banking and financial services. Technology Innovation Management Review, 12(1/2).

Benetti, Z. (2021). Fraud detection in ethereum using web-scraping and natural language processing techniques (Master's thesis, ETH Zurich).

Jung, Eunjin, Marion Le Tilly, Ashish Gehani, and Yunjie Ge. "Data mining-based ethereum fraud detection." In *2019 IEEE international conference on blockchain (Blockchain)*, pp. 266-273. IEEE, 2019.

Li, C. (2022). A fraud detection system for reducing blockchain transaction risks using explainable graph neural networks. The George Washington University.

Lo, W. W., Kulatilleke, G. K., Sarhan, M., Layeghy, S., & Portmann, M. (2023). Inspection-L: self-supervised GNN node embeddings for money laundering detection in bitcoin. Applied Intelligence, 53(16), 19406-19417.

Chen, J., Chen, Q., Jiang, F., Guo, X., Sha, K., & Wang, Y. (2024). SCN_GNN: A GNN-based fraud detection algorithm combining strong node and graph topology information. Expert Systems with Applications, 237, 121643.

Adloori, H., Dasanapu, V., & Mergu, A. C. (2024). Graph Network Models To Detect Illicit Transactions In Block Chain. arXiv preprint arXiv:2410.07150.

Asiri, A., & Somasundaram, K. (2025). Graph convolution network for fraud detection in bitcoin transactions. Scientific Reports, 15(1), 11076.

Tripathy, N., Balabantaray, S.K., Parida, S. and Nayak, S.K., 2024. Cryptocurrency fraud detection through classification techniques. *International Journal of Electrical and Computer Engineering (IJECE)*, *14*(3), pp.2918-2926.

Luo, Q., Zeng, W., Chen, M., Peng, G., Yuan, X., & Yin, Q. (2023, July). Self-attention and transformers: Driving the evolution of large language models. In 2023 IEEE 6th International conference on electronic information and communication technology (ICEICT) (pp. 401-405). IEEE.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. Advances in neural information processing systems, 30(1), 2.

Hosseinzadeh, R., & Sadeghzadeh, M. (2025). Attention Mechanisms in Transformers: A General Survey. Journal of AI and Data Mining, 13(3), 359-368.

Xie, R., & Liu, D. A Novel Hybrid Graph Neural Network and Transformer Model for Intrusion Detection in Power Cyber-Physical Systems. Available at SSRN 5055232.

Aleksandr, T. (2025). GRAPH TRANSFORMERS. Universum: технические науки, 6(1 (130)), 31-39.

Pérez-Cano, V., & Jurado, F. (2025). Fraud detection in cryptocurrency networks—An exploration using anomaly detection and heterogeneous graph transformers. Future Internet, 17(1), 44.

Feng, P. (2025). Hybrid BiLSTM-Transformer Model for Identifying Fraudulent Transactions in Financial Systems. Journal of Computer Science and Software Applications, 5(3).