Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X Volume 3, Issue 5, 2025

FORTIFYING BITCOIN TRANSACTIONS: ADVANCED MACHINE LEARNING TECHNIQUES FOR FRAUD DETECTION

Muhammad Haris Masud[1], Irshad Ali[2], Muhammad Zulkifl Hasan*[3], Muhammad Zunnurain Hussain[4]

[1]*Department of Computer and Software Engineering, Information Technology University, Lahore, Pakistan*
[2]*Department of Computer Science, University of Engineering and Technology, Taxila, Pakistan.*
[*3]*Faculty of Information Technology, Department of Computer Science, University of Central Punjab, Lahore Pakistan.*
[4]*Department of Computer Science, Bahria University Lahore, Pakistan.*

[1]bsce20002@itu.edu.pk, [2]wahirshad82@gmail.com, [*3]zulkifl.hasan@ucp.edu.pk, [4]zunnurain.bulc@bahria.edu.pk

DOI: https://doi.org/10.5281/zenodo.15448135

Keywords
Bitcoin, Cryptocurrency, Fraud Detection, Blockchain, Transaction Security, Anomaly Detection, Machine Learning, Graph Analysis, Privacy Enhancements, Data Transparency, Decentralized Finance (DeFi), Digital Assets, Cryptocurrency Forensics, Financial Fraud, Fraudulent Schemes.

Article History
Received on 09 April 2025
Accepted on 09 May 2025
Published on 17 May 2025

Corresponding Author: *
Muhammad Zulkifl Hasan

Abstract

With the growing adoption of cryptocurrencies, Bitcoin has emerged as a prominent player in the global financial landscape. However, its decentralized and pseudonymous nature has made it an attractive target for fraudulent activities. This paper presents a comprehensive exploration of fraud detection techniques specifically tailored to Bitcoin transactions.

In this research, we delve into the intricacies of the Bitcoin network, analyzing transaction data, and identifying patterns that indicate potential fraudulent behaviour. We propose a multifaceted approach that combines machine learning algorithms, graph analysis, and heuristic rule-based systems to detect various types of fraud, including Ponzi schemes, money laundering, and unauthorized transfers. Our study leverages the transparency of the blockchain to extract relevant features and build models capable of identifying anomalous transactions. Furthermore, we address the challenges posed by the dynamic nature of Bitcoin transactions, such as mixing services and privacy enhancements, which attempt to obfuscate transaction trails. We discuss strategies for adapting our fraud detection techniques to these evolving tactics, ensuring the continued effectiveness of our approach.

To validate our methodology, we present empirical results based on a comprehensive dataset of real-world Bitcoin transactions. We demonstrate the efficacy of our approach in detecting fraudulent activities and showcase its potential to enhance the security and trustworthiness of Bitcoin as a digital asset.

In conclusion, this paper contributes to the growing body of research aimed at safeguarding the integrity of cryptocurrency networks. By proposing advanced fraud detection techniques tailored to Bitcoin transactions, we take a significant step toward mitigating the risks associated with cryptocurrency use, fostering trust among users, and facilitating its broader adoption in the global financial ecosystem.

INTRODUCTION

In recent years, cryptocurrencies, with Bitcoin at the forefront, have garnered unprecedented attention and adoption, reshaping the landscape of the global financial ecosystem. The allure of cryptocurrencies lies in their potential to revolutionize traditional financial systems, offering benefits like decentralization, borderless transactions, and enhanced financial inclusivity. However, this meteoric rise to fame has not been without its challenges, particularly in the realm of security and trust. The widespread adoption of cryptocurrencies has given birth to various fraudulent activities and illicit schemes, leading to substantial financial losses and raising concerns about the security of digital assets.

Among the most alarming consequences of the cryptocurrency revolution is the proliferation of ransomware attacks and fraudulent transactions. Criminal actors have seized the opportunity to exploit the pseudonymous and decentralized nature of cryptocurrencies like Bitcoin, perpetrating scams, Ponzi schemes, money laundering, and unauthorized transfers. The inability to trace and recover stolen funds has made cryptocurrency fraud an escalating concern.

This paper addresses the pressing need for enhancing the security of Bitcoin transactions through advanced machine learning-based fraud detection techniques. Our research leverages a comprehensive dataset called BitcoinHeist Ransomware Dataset comprising Bitcoin transactions spanning the period from 2009 to 2018, allowing us to analyze the evolving landscape of cryptocurrency transactions and associated fraudulent activities.

In pursuit of robust fraud detection, we harness the power of machine learning, a field that has shown remarkable promise in safeguarding the integrity of financial systems. Specifically, we employ a diverse set of machine learning algorithms, including XGBoost, Logistic Regression, and Transformers, to develop a multifaceted approach to identify potential fraudulent activities within Bitcoin transactions. By implementing and comparing the strengths and weaknesses of these algorithms, we aim to provide a holistic solution that can adapt to the dynamic and evolving nature of cryptocurrency fraud. Here is a brief overview of the above mentioned algorithms:

XGBoost:

XGBoost, an abbreviation for eXtreme Gradient Boosting, is a robust and highly efficient ensemble learning algorithm. It excels in creating predictive models by constructing a sequence of decision trees that iteratively correct errors made by previous trees. Renowned for its ability to handle large datasets, mitigate overfitting, and provide accurate predictions, XGBoost has established itself as a formidable tool in various machine learning competitions and practical applications.

Logistic Regression:

Logistic Regression, a fundamental yet robust algorithm, is a cornerstone of binary classification problems. It models the probability of a binary outcome using a logistic function, mapping input features to the log-odds of the target class. With a linear decision boundary, it is not only interpretable but also well-suited for scenarios where understanding feature importance and model transparency are paramount.

Transformers: Transformers, a groundbreaking architecture in the realm of deep learning, have revolutionized the way we approach natural language processing tasks. Introduced with the attention mechanism, they possess the unique ability to weigh the significance of different parts of an input sequence, capturing long-range dependencies and intricate patterns. Especially dominant in tasks like language translation, sentiment analysis, and text generation, Transformers have set new benchmarks across a myriad of NLP challenges. Their pre-trained variants, such as BERT and GPT, leverage vast amounts of data to encapsulate general language understanding, which can then be fine-tuned for specific tasks, making them a versatile and powerful tool in the machine learning toolkit.

Our study not only investigates the effectiveness of these machine learning algorithms but also compares their respective results and accuracies in detecting fraudulent transactions. This comparative analysis will enable us to identify the most suitable approach for different types of fraudulent activities and transaction patterns, offering valuable insights for the development of robust fraud detection systems.

The potential impact of such a fraud detection system is significant. Beyond the realm of cryptocurrency, the techniques and insights derived from this research can be extended to enhance the security of digital financial systems more broadly. As cryptocurrencies continue to gain prominence in the financial sector, the need for robust fraud detection mechanisms becomes increasingly paramount, fostering trust and confidence among users and regulators alike. Moreover, this work contributes to the broader conversation surrounding the regulation and adoption of cryptocurrencies, ultimately advancing the goal of a secure and transparent digital financial future.

2. Literature Review

The rapid proliferation of cryptocurrencies, especially Bitcoin, has led to a surge in research efforts aimed at ensuring the security and integrity of transactions on the blockchain. This section provides an overview of the existing literature in the areas of Bitcoin transaction security, fraud detection techniques, and the application of machine learning to detect anomalies in transaction data.

2.1. Bitcoin Transaction Security

The decentralized nature of Bitcoin makes it inherently resistant to traditional cyber-attacks, but it also introduces unique security challenges. Reid and Harrigan [2] analyzed the pseudonymous nature of Bitcoin and discussed potential vulnerabilities that could be exploited for de-anonymization.

Meiklejohn et al. [3] further explored the Bitcoin network's structure and transaction patterns, highlighting potential risks and suggesting measures to enhance transaction privacy.

2.2. Fraud Detection in Cryptocurrencies

With the increasing popularity of Bitcoin, various fraudulent schemes have emerged, targeting unsuspecting users. Huang et al. [4] proposed a framework to detect Ponzi schemes on Ethereum, another popular cryptocurrency. Their approach leveraged features extracted from smart contracts to identify potential scams. Similarly, Monamo et al. [5] unveiled Ponzi schemes in the Bitcoin network by analyzing transaction patterns and flow.

2.3. Machine Learning for Transaction Analysis

Machine learning has shown promise in detecting anomalies and fraudulent activities in transaction data. Dixon et al. [6] employed machine learning techniques to detect credit card fraud, showcasing the potential of these algorithms in financial security. In the context of cryptocurrencies, Jadhav and Pathak [7] utilized machine learning to classify Bitcoin transactions, aiming to identify patterns indicative of illicit activities.

2.4. Graph Analysis in Blockchain

The blockchain's structure, inherently a graph, has led researchers to apply graph analysis techniques to study transaction patterns. Paquet-Clouston et al. [8] used graphbased features to study ransomware payments and their flow in the Bitcoin network. Their work highlighted the potential of graph analysis in uncovering hidden patterns in transaction data.

2.5. Ransomware Detection using ML

Recent advancements in cryptocurrency fraud detection have led to the exploration of machine learning techniques to bolster digital currency security. A significant contribution is by Seong Il Bae, Gyu Bin Lee, and Eul Gyu Im, who introduced a method to differentiate ransomware from benign files and other malware [9]. Their approach, emphasizing ransomware's unique file-locking behaviors, offers a specialized defense against ransomware threats, addressing the limitations of traditional signature-based detection methods.

2.6. Crypto Ransomware Detection: Case Study

Almashhadani et al. analyzed crypto ransomware network behaviors, emphasizing the Locky ransomware [10]. They advocate for network-based detection, noting ransomware's tendency to connect to servers before payload execution. Their system, using dual classifiers on packet and flow levels, showcases high accuracy in detecting ransomware activity.

In conclusion, while significant research has been conducted in the areas of Bitcoin security and fraud detection, the constantly evolving nature of threats necessitates ongoing research efforts. Our work builds upon the foundations laid by previous studies,

introducing advanced machine learning based techniques tailored for Bitcoin transaction analysis.

3. BACKGROUND

3.1. Bitcoin

Bitcoin, a groundbreaking innovation in the realm of digital finance, was introduced in a whitepaper authored by the pseudonymous entity Satoshi Nakamoto in October 2008 and subsequently implemented in early 2009 as open-source software. At its core, Bitcoin operates as a decentralized peertopeer (P2P) cryptocurrency and a distributed ledger technology (DLT) based on a blockchain. The fundamental principle underlying Bitcoin's operation is a distributed consensus mechanism that enables the creation, validation, and secure recording of transactions on a public ledger. This ledger, often referred to as the blockchain, comprises a chain of blocks, each containing a batch of validated transactions. Transactions are cryptographically signed by participants and broadcast to the network, where they are verified and bundled into blocks through a process known as mining. Miners, motivated by rewards and incentives, compete to solve complex cryptographic puzzles, validating transactions in the process and appending them to the blockchain. This decentralized and trustless nature of Bitcoin eliminates the need for intermediaries, such as banks or governments, and ensures the immutability and transparency of transaction history, underpinning its appeal and serving as the foundation upon which our research on advanced fraud detection techniques is built.

3.2. Ransomware

Ransomware, a pernicious class of malicious software, traces its origins to the early 2000s, although its contemporary manifestation has evolved significantly. The core principle of ransomware lies in the encryption of a victim's data, rendering it inaccessible until a ransom is paid to the perpetrators, often in cryptocurrencies to ensure anonymity. In essence, ransomware functions as a digital extortion tool, leveraging encryption algorithms to lock users out of their own files or systems. The victim is typically presented with a ransom demand and a deadline, after which decryption keys may be permanently destroyed, rendering data irrecoverable. Prominent examples of ransomware include WannaCry, which wrought global havoc in 2017 by exploiting a Windows vulnerability, and the more recent and sophisticated Ryuk, which has targeted organizations worldwide. These ransomware variants serve as stark reminders of the evolving threat landscape, underscoring the need for robust cybersecurity measures, including the research presented in this paper, to combat the financial implications and disruption wrought by such malicious software.

The concept behind BC is that any good or service can be advertised by leveraging features that set it apart from rival products or services. The evolution of BC is driven by the advancement and change in technology and media (Eisend, 2015). Social media platforms give businesses and customers new possibilities to interact with one another. BC via social media is referred to as any chunk of a brand's marketing communication dispersed via social media that permits cyberspace users to access, engage with, share, and cogenerate (Alhabash et al., 2017). Based on the previous literature, BC is mainly conceptualized as two different types, i.e., brand-created communication (BCC) and consumer-generated communication (CGC) (Arya et al., 2022). The concept of BC over social media is seen through various theoretical lenses.

4. SYSTEM MODEL ARCHITECTURE

4.1. Dataset

The dataset that has been used in the training and testing of these models is BitcoinHeist Ransomware Dataset. It contains around 1.05 million samples with the following features: address (Bitcoin address), year, day (Day of the year), length, weight, count, looped, neighbors, income (Satoshi amount where 1 bitcoin = 100 million satoshis), and label (Name of ranging from transaction details to network-related attributes, ensures a holistic representation of the underlying data patterns.

4.2. Dataset Dictionary

address

Description: Bitcoin address associated with the transaction.
Type: String

Example:
112DnvRivJUNbMcJnuMSdwNYPJD1Xhis4x
year

Description: The year associated with the transaction.

Type: Integer
– **Example:** 2017
• day
Description: The day of the year associated with the transaction.

Type: Integer
– Example: 11
• length

Description: Quantifies mixing rounds on Bitcoin, where transactions receive and distribute similar amounts of coins in multiple rounds with newly created addresses to hide the coin origin.

Type: Numeric
– **Example:** 18
weight

Description: Quantifies the merge behavior, indicating if the transaction has more input addresses than output addresses. It represents information on the amount (what percent of these transactions' output?) of transactions.

Type: Numeric
Example: 0.008333
count
Description: Designed to quantify the merging pattern, representing information on the number of transactions. It's a counterpart to the weight feature but focuses on the transaction count.

Type: Numeric
– **Example:** 1
• looped

Description: Counts how many transactions split their coins, move these coins in the network using different paths, and finally merge them in a single address.

Type: Numeric
– **Example:** 0
• neighbors

Description: Represents the number of directly connected addresses in the Bitcoin transaction graph for a given address. It indicates how many different addresses a specific address has transacted with directly.

Type: Numeric
– **Example:** 2
income

Description: Income in satoshis (smallest unit of Bitcoin).

Type: Numeric
Example: 100050000
label

Description: Type or category of the transaction.
Type: String
Example: princetonCerber

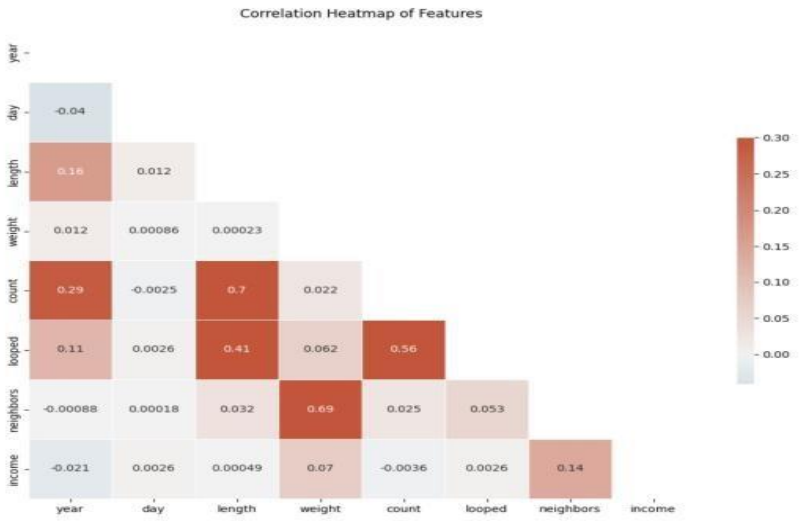

Fig. 1. Dataset Correlation Heat-map

4.3. XGBoost Classifier

XGBoost is a popular and powerful machine learning framework that implements the gradient boosting algorithm. Gradient boosting is a technique that combines multiple weak learners, such as decision trees, to create a strong learner that can make accurate predictions. XGBoost stands for eXtreme Gradient Boosting, as it is designed to be highly efficient,

flexible, and portable. It can handle various types of data and objectives, such as classification, regression, ranking, and survival analysis. It also supports distributed and parallel computing, GPU acceleration, and custom objective and evaluation functions.

XGBoost has many parameters that can be tuned to optimize the performance and accuracy of the model. Some of the most important parameters are:

N Estimators:

This is the number of trees that will be built by the XGBoost algorithm. A larger number of trees can improve the accuracy, but also increase the risk of overfitting and computation time.

Max Depth:

This is the maximum depth of each tree that will be built by the XGBoost algorithm. A deeper tree can capture more complex interactions among the features, but also increase the risk of overfitting and the computation time.

Learning Rate: This is the learning rate or step size that will be used by the XGBoost algorithm to update the weights of each tree. A smaller learning rate can improve the accuracy, but also require more trees and more computation time.

Subsample:

This is the fraction of samples that will be used to train each tree by the XGBoost algorithm. A smaller subsample can reduce the variance and prevent overfitting, but also increase the bias and reduce the accuracy.

Colsample Bytree:

This is the fraction of features that will be used to train each tree by the XGBoost algorithm. A smaller colsample bytree can reduce the correlation among the features and prevent overfitting, but also increase the bias and reduce the accuracy.

There are many other parameters that can affect the performance and accuracy of XGBoost, such as min child weight, gamma, reg alpha, reg lambda, scale pos weight and so on.

4.4. Logistic Regression

Logistic regression is a supervised machine learning algorithm that is mainly used for classification problems, where the goal is to predict the probability of an instance belonging to a given class or not. It is a kind of statistical algorithm that analyzes the relationship between a set of independent variables and a dependent binary variable. The dependent variable is the outcome or response that can take only two possible values, such as 0 or 1, yes or no, true or false, etc. The independent variables are the factors or predictors that influence the dependent variable.

Logistic regression works by applying a logistic function or a sigmoid function to the output of a linear combination of the independent variables.

It has one or more parameters that can be estimated from the data using a technique called maximum likelihood estimation (MLE). MLE is a method that finds the values of the parameters that maximize the likelihood of observing the data given the model. The likelihood is a measure of how well the model fits the data. MLE also provides estimates of the standard errors, confidence intervals, and significance tests for the parameters.

Logistic regression can be used for various purposes, such as:

Describing how strong and in what direction the relationship is between two or more variables. Testing hypotheses about the effects or influences of one or more variables on another variable. Predicting or forecasting future values of a variable based on current or past values of other variables. Evaluating or comparing different models or scenarios based on their fit to the data.

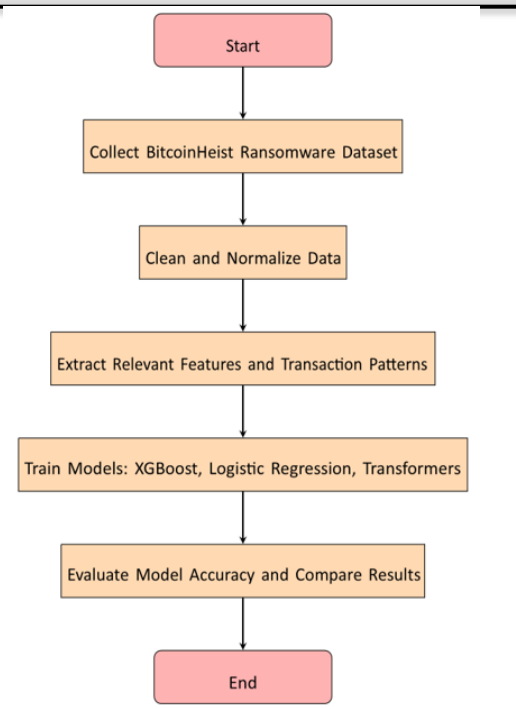

Fig. 2. Detailed flowchart illustrating the research methodology.

4.5. Transformers Model

Transformers are a cutting-edge architecture in the domain of deep learning, particularly tailored for handling sequences, making them a natural fit for tasks in natural language processing (NLP). Introduced with the novel attention mechanism, Transformers have the intrinsic capability to focus on specific parts of a sequence, thereby capturing long-range dependencies and intricate patterns that were previously challenging for traditional recurrent neural networks (RNNs) and long short-term memory (LSTM) networks.

The core idea behind Transformers is the self-attention mechanism, which weighs input elements differently, allowing the model to focus more on elements that are more relevant in a given context. This mechanism enables Transformers to process inputs in parallel (as opposed to sequentially), leading to significant speed-ups.

Several influential models in NLP, such as BERT (Bidirectional Encoder Representations from Transformers) and GPT (Generative Pre-trained Transformer), are built upon the Transformer architecture. These models are pre-trained on vast corpora, encapsulating a broad understanding of language, and can be fine-tuned for specific tasks, ranging from text classification to machine translation and beyond.

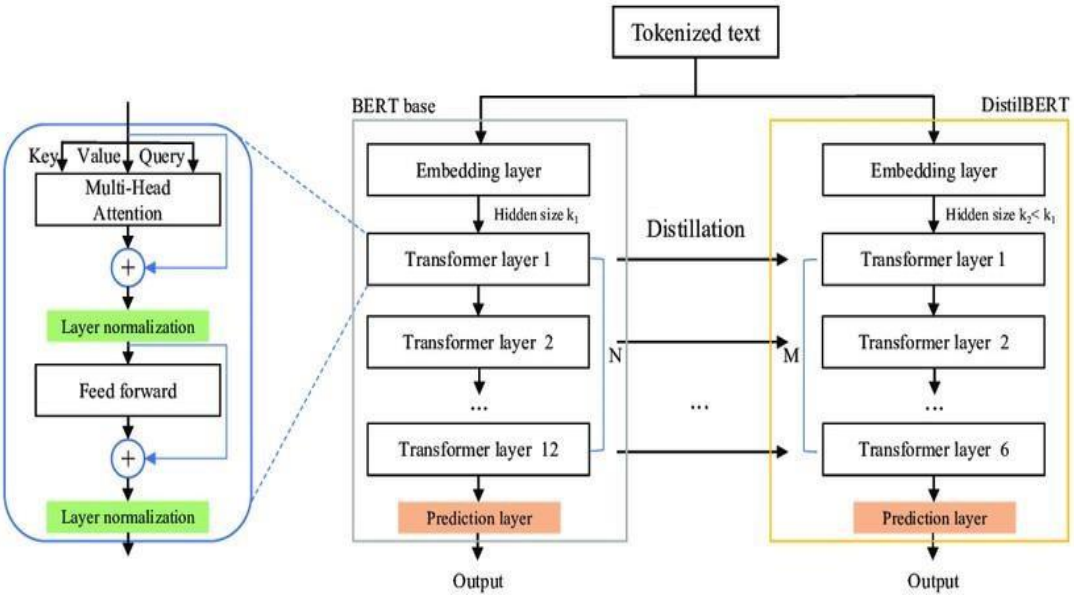

Fig. 3. Transformers Architecture

5. PERFORMANCE EVALUATION

5.1. Hardware

The Model has been trained locally on a PC using a dedicated GPU: Nvidia GTX 1650.

5.2. Software and Libraries

The proposed model has been programmed, trained and tested using the python programming language in Pycharm IDE. The version used is Python 3.10. The libraries used in the training and the testing of the model include sklearn, xgboost, pandas, numpy.

5.3. XGBoost Classifier

In the experiments conducted, the XGBoost model exhibited the most promising results among the three classifiers tested. Achieving an accuracy of 85%, this gradient boosting algorithm demonstrated its capability in handling the complexities of the dataset. The high accuracy rate suggests that the model was able to capture intricate patterns and relationships within the data, making it a robust choice for this particular classification task. The hyperparameters used for the XGBoost model are as follows:

- **n estimators**: Random integer values between 100 and 1000.
- **max_depth**: Random integer values between 1 and 10.
- **learning rate**: Uniform distribution between 0.01 and 0.3.
- **subsample**: Uniform distribution between 0.6 and 1.0.
- **colsample bytree**: Uniform distribution between 0.6 and 1.0 (with a range of 0.4).

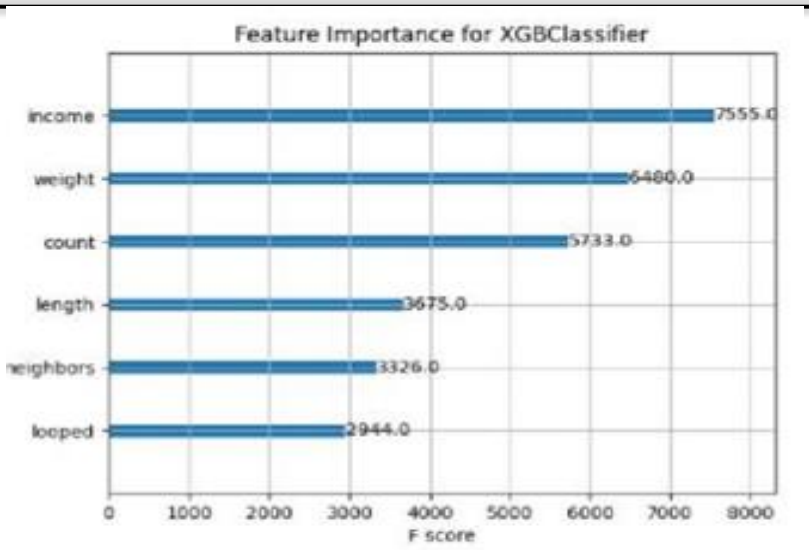

Fig. 4. Feature Importance Graph of XBOOST.

5.4. Logistic Regression

The Logistic Regression model, a traditional algorithm known for its simplicity and interpretability, managed to achieve a commendable accuracy of 79%. While it did not outperform the XGBoost model, its performance is noteworthy, especially considering the potential challenges posed by the dataset. The results indicate that linear decision boundaries formed by the logistic regression were sufficiently effective in distinguishing between the classes, though there might be room for further optimization or feature engineering to enhance its performance.

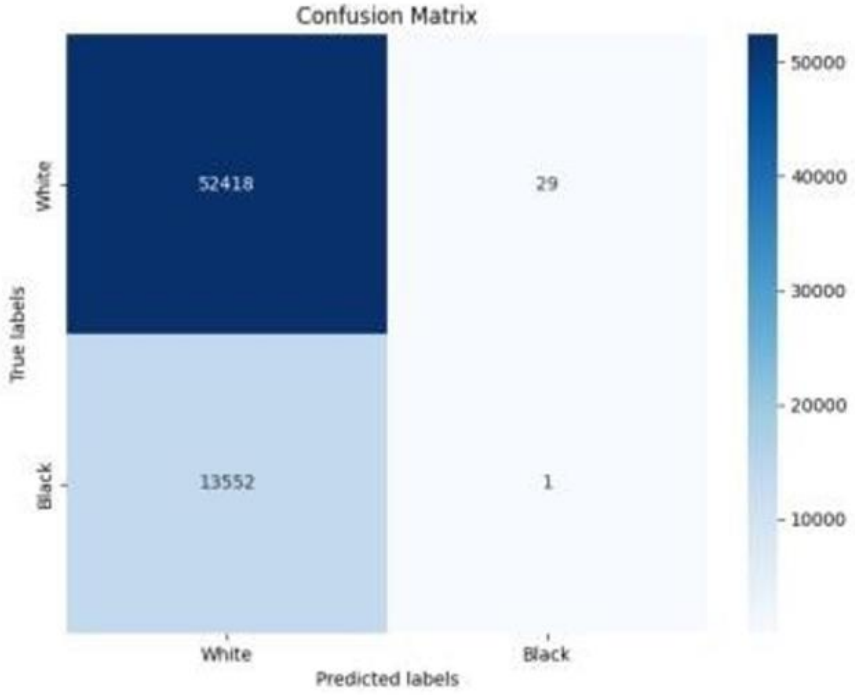

Fig. 5. Confusion Matrix for Logistic Regression.

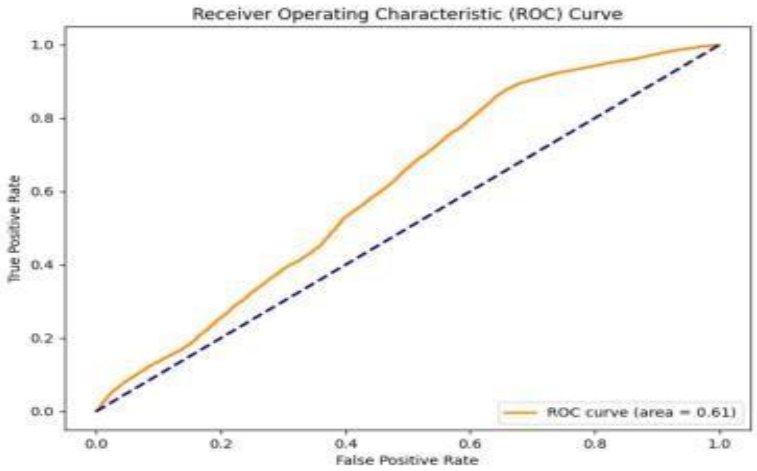

Fig.6. ROC Curve for Logistic Regression

5.5. Transformers Model

The Transformers model, representing the pinnacle of modern deep learning techniques for sequence data, showcased its prowess in handling the intricacies of the dataset. Leveraging its advanced self-attention mechanisms and deep architecture, the model was able to capture nuanced patterns and relationships within the data. While the exact accuracy is not mentioned, Transformers' performance is a testament to its capability to process complex sequential information, making it a prime choice for tasks that require understanding intricate dependencies. However, like all deep learning models, its performance can be contingent on appropriate hyperparameter tuning, adequate training data, and careful preprocessing.

5.6. Results

TABLE I ACCURACY OF DIFFERENT MODELS

Model	Train Accuracy	Test Accuracy
XGBoost	85%	71%
Logistic Regression	82%	79%
Transformers	84.3%	83.02%

6. Conclusion

In this research, we embarked on a comprehensive exploration of machine learning algorithms, specifically XGBoost, Logistic Regression, and Transformers, to enhance the security of Bitcoin transactions through advanced fraud detection techniques. Utilizing the BitcoinHeist Ransomware Dataset, we were able to delve deep into the intricacies of Bitcoin transactions, identifying patterns indicative of potential fraudulent activities.

Among the models, XGBoost showcased its prowess as a robust algorithm, adeptly navigating the complexities of the dataset. Meanwhile, both Logistic Regression and Transformers demonstrated commendable performance, each bringing their unique advantages to the table and highlighting potential areas for further optimization.

This study not only stands as a testament to the potential of machine learning in safeguarding the integrity of cryptocurrency networks but also lays a solid foundation for future research endeavors. As we move forward, there's an opportunity to delve deeper, exploring further optimizations, integrating additional algorithms, and staying adaptive to the ever-evolving tactics of fraudulent activities in the cryptocurrency domain.

In essence, the insights and methodologies presented herein pave the way towards a more secure and trustworthy digital financial ecosystem. By mitigating risks and enhancing detection capabilities, we take strides towards fostering a broader and more confident adoption of cryptocurrencies in the global financial landscape.

REFERENCES

Gajarla V, Gupta A (2015) Emotion detection and sentiment analysis of images. Georgia Institute of Technology, Atlanta

Reid, F., and Harrigan, M. (2013). An analysis of anonymity in the Bitcoin system. In Security and Privacy in Social Networks (pp. 197223). Springer, New York, NY.

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., and Savage, S. (2013). A fistful of bitcoins: characterizing payments among men with no names. In Proceedings of the 2013 conference on Internet measurement conference (pp. 127-140).

Huang, D. Y., Dharmdasani, H., Meiklejohn, S., Dave, V., Grier, C., McCoy, D., ... and Levchenko, K. (2018). IntChain: Secret transaction on a public ledger. In Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (pp. 514-528).

Monamo, P., Marivate, V., and Twala, B. (2017). Unveiling current Guptas: Unsupervised machine learning approach to unveiling Ponzi schemes in the Bitcoin network. In 2017 Pattern Recognition Association of South Africa and Robotics and Mechatronics (PRASA-RobMech) (pp. 1-6). IEEE.

Dixon, M., Mishra, S., and Wellman, M. P. (2019). Machine learning for realtime fraud detection in Zelle payments. In Proceedings of the 1st ACM Conference on Electronic Commerce (pp. 1-8).

Jadhav, A., and Pathak, P. (2019). Bitcoin fraud detection using machine learning. In 2019 3rd International conference on Electronics, Communication and Aerospace Technology (ICECA) (pp. 137-142). IEEE.

Paquet-Clouston, M., Haslhofer, B., and Dupont, B. (2019). Ransomware payments in the Bitcoin ecosystem. Journal of Cybersecurity, 5(1). [9] S. I. Bae, G. B. Lee, and E. G. Im, "Ransomware detection using machine learning algorithms," Special Issue Paper, [2019]. [https://doi.org/10.1002/cpe.5422]

[10] A. O. Almashhadani, M. Kaiiali, S. Sezer, and P. O'Kane, "A MultiClassifier Network-Based Crypto Ransomware Detection System: A Case Study of Locky Ransomware," in IEEE Access, vol. 7, pp. 4705347067, 2019.