





**Please cite the Published Version**

Obi-Okoli, Chibuzo , Jogunola, Olamide , Adebisi, Bamidele  and Hammoudeh, Mohammad  (2023) Machine Learning Algorithms to Detect Illicit Accounts on Ethereum Blockchain. In: ICFNDS '23: The International Conference on Future Networks and Distributed Systems, 21 December 2023 - 22 December 2023, Dubai, United Arab Emirates.

**DOI:** <https://doi.org/10.1145/3644713.3644838>

**Publisher:** Association for Computing Machinery (ACM)

**Version:** Published Version

**Downloaded from:** <https://e-space.mmu.ac.uk/636228/>

**Usage rights:**  [Creative Commons: Attribution 4.0](https://creativecommons.org/licenses/by/4.0/)

**Additional Information:** This is an open access conference paper of a presentation first given at ICFNDS '23: The International Conference on Future Networks and Distributed Systems

**Enquiries:**

If you have questions about this document, contact [openresearch@mmu.ac.uk](mailto:openresearch@mmu.ac.uk). Please include the URL of the record in e-space. If you believe that your, or a third party's rights have been compromised through this document please see our Take Down policy (available from <https://www.mmu.ac.uk/library/using-the-library/policies-and-guidelines>)



# Machine Learning Algorithms to Detect Illicit Accounts on Ethereum Blockchain

Chibuzo Obi-Okoli

chibuzo.d.obi-okoli@stu.mmu.ac.uk  
Manchester Metropolitan University  
Manchester, UK

Bamidele Adebisi

b.adebisi@mmu.ac.uk  
Manchester Metropolitan University  
Manchester, UK

Olamide Jogunola\*

o.jogunola@mmu.ac.uk  
Manchester Metropolitan University  
Manchester, UK

Mohammad Hammoudeh

mohammad.hammoudeh@kfupm.edu.sa  
King Fahd University of Petroleum and Minerals  
Dhahran, Saudi Arabia

## ABSTRACT

The rapid growth and pseudonymity inherent in blockchain technology such as in Bitcoin and Ethereum has marred its original intent to reduce dependant on centralised system, but created an avenue for illicit activities, including fraud, phishing, scams, etc. This undermines the reputation of blockchain network, giving rise to the need to identify these illicit activities within the blockchain network. This current work tackles this crucial problem by investigating and implementing six machine learning algorithms with a particular emphasis on striking a balance between accuracy, precision and recall. The novelty of the work lies in the utilising of the synthetic minority over-sampling technique to handle data imbalance. Thus, increasing the accuracy of the light gradient boosting machine classifier to 98.4%. The outcome of this work holds great potential for enhancing the security and credibility of blockchain ecosystems paving the way for a more secure and dependable digital future in the age of decentralised and trustless systems.

## KEYWORDS

Ethereum blockchain, machine learning, anomaly detection, blockchain security, illicit activities

### ACM Reference Format:

Chibuzo Obi-Okoli, Olamide Jogunola, Bamidele Adebisi, and Mohammad Hammoudeh. 2023. Machine Learning Algorithms to Detect Illicit Accounts on Ethereum Blockchain. In *The International Conference on Future Networks and Distributed Systems (ICFNDS '23)*, December 21–22, 2023, Dubai, United Arab Emirates. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3644713.3644838>

## 1 INTRODUCTION

Blockchain technology, initially developed as the backbone for cryptocurrencies like Bitcoin and Ethereum, has a core purpose of reducing the need for centralised intermediaries, such as banks, in financial transactions. Blockchain features, including immutability

and transparency, have enabled its broader pilot and adoption across diverse industries, including in education, healthcare, finance, energy, government, supply chain, and the Internet of Things (IoT). These industries have leveraged blockchain to enhance privacy, facilitate faster transactions, and strengthen security [6, 8]. While numerous benefits of blockchain has been assessed in previous studies [5, 8], blockchain is not immune to challenges. The prominent challenges and risks associated with blockchain include scalability and security, amongst others. Illicit activities on blockchain encompass various illegal or fraudulent transactions such as money laundering, phishing, scams, hacking, and specifically, the creation of illicit accounts. These illicit activities can severely undermine the reputation and trust associated with blockchain technology and its users [3].

Identifying and detecting these illicit activities including fake accounts can help prevent or mitigate the damages caused on the blockchain network, thereby enhancing the security and trustworthiness of the technology. The Authors in [12] discussed several methods that can be used to analyse a blockchain transaction to detect illicit activities, these includes rule-based method, graph analysis, and machine learning (ML). ML, a subdivision of artificial intelligence, has materialised as a potent instrument for diverse applications spanning multiple domains. By enabling computers to glean insights from data and render projections or judgments via explicit programming. ML algorithms leverage the rich and publicly available data on Ethereum transactions, smart contracts, and network activities to identify patterns, anomalies, and behaviours [1, 3].

In recent years, ML is explored to detect illicit activities within Ethereum blockchain network. In [11], the researchers used a random forest (RF) classifier framework to identify Ethereum entities involved in malicious activities. Through research, it's been discovered that illicit accounts often exhibit certain common characteristics, such as an unusually high number of transactions, irregular transaction amounts, involvement in known illicit activities, or connections to previously identified fraudulent accounts [3, 11]. These insights became the foundation for training their machine learning models. The authors in [1] applied temporal graph properties to detect malicious accounts on Ethereum blockchain. They applied several supervised ML including support vector machine (SVM), decision tree, RF, etc., and evaluated their model using the accuracy metric. [3] applied the Extreme gradient boosting algorithm (XGBoost), using features based on the transaction behaviour of the



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICFNDS '23, December 21–22, 2023, Dubai, United Arab Emirates  
© 2023 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0903-6/23/12  
<https://doi.org/10.1145/3644713.3644838>

accounts, such as transaction amount, transaction frequency, and transaction recency. However, [10] argued that existing methods for detecting illicit entities in Bitcoin are limited in their scope, as they typically only focus on a single type of illicit activity. They propose to address this limitation by using an ensemble of decision trees to learn discriminating features that can be used to categorise multiple groups of illicit users from licit users. They proposed a supervised learning model that combines multiple weak learners (decision trees) to form a strong learner that can improve the accuracy and reduce the variance of the predictions.

Learning algorithms have shown promising results in detecting illicit accounts on blockchain platforms using various techniques and features. However, there are still some challenges and limitations that need to be addressed, such as data quality, label availability, interpretability, scalability and generalisation [5, 7]. Besides and due to the nature of blockchain, it is vital to evaluate the performance of different ML algorithms in detecting illicit activities. For instance, the authors in [6], performed feature selection and a comparative analysis of different ML, such as K-nearest neighbours, decision tree and RF. The result showed an improvement in the F-score for the RF. Likewise, [1] utilised temporal graph properties and several ML to detect malicious accounts on blockchain, their result indicate ExtraTrees Classifier has the highest accuracy in the supervised ML cases. In [3], XGBoost was proposed and evaluated.

Selecting the ML model that performed best in these reviewed studies, the following are the specific contributions of this work:

- Evaluate and analyse six ML algorithms for the detection of illicit activities on the Ethereum blockchain network. The selected algorithms achieved the highest accuracy in the reviewed literature.
- Enhance the performance of these algorithms by performing a robust data preprocessing including the utilisation of the synthetic minority over-sampling technique (SMOTE) to handle data imbalance and the use of GridSearchCV for hyperparameter tuning.
- Discuss and evaluate the learning algorithms utilising several performance metrics, highlighting their usefulness, outcome and trade-offs.

The remaining sections are organised as follows. Section 2 presents the methodology and the experimental setup. It also discusses the ML algorithms, dataset description and preprocessing, and the evaluation metrics for evaluating the performance of the detecting capability of the ML algorithms. Section 3 discusses the result of the six ML algorithms utilising the defined performance metrics. Finally, Section 4 concludes the paper with references to potential future work.

## 2 METHODOLOGY

This section first presents the ML algorithms used in this article for illicit account detection on the Ethereum blockchain network. Followed by the dataset description, data preprocessing and the description of the evaluation metrics. The flowchart illustrating the overall workflow of this work is shown in Fig 1.

### 2.1 Machine learning algorithms

Detecting illicit accounts on Ethereum blockchain is a binary classification problem, such as a fraudulent or non-fraudulent activities. Thus, based on the literature, we considered a range of ML algorithms suitable for this type of problem.

**2.1.1 AdaBoost.** Adaptive Boosting (AdaBoost), is an ML algorithm aimed at amalgamating numerous weak classifiers to form a potent classifier. In this context, a weak classifier is a model that exhibits a marginally superior performance compared to random guessing, whereas a strong classifier showcases high accuracy and robust generalisation capabilities. The core mechanism of AdaBoost involves a sequential process of fitting a weak classifier to the training data, with the data's weights being updated based on the prior classifier's prediction errors [11]. Ultimately, the final prediction is generated by conducting a weighted aggregation of all the weak classifiers. One of AdaBoost's notable features is its compatibility with various base classifiers that support weighted samples, such as decision trees, logistic regression, and SVM. AdaBoost offers flexibility in the selection of loss functions and the determination of the number of iterations. This method boasts several advantages, including speed, simplicity, and resilience against noise and outliers. Nevertheless, AdaBoost does come with certain limitations, such as sensitivity to mislabelled data and the potential for overfitting when employing an overly complex base classifier.

**2.1.2 LightGBM classifier.** LightGBM (LGBM) is a renowned and appropriate classifier for binary classification. It is a gradient boosting framework that thrives in scenarios involving large-scale datasets. LightGBM is known for its efficiency and speed, due to its histogram-based learning approach, that minimises memory usage and accelerates training, making it ideal for handling large datasets and intricate feature spaces [2]. In addition, LGBM is designed for precise predictions, excelling at capturing complex data patterns and relationships. Its ability to construct and optimise deep trees contributes to its predictive accuracy, a critical factor in binary classification tasks where accurate class discrimination is paramount. Beside, in terms of real-world datasets that includes categorical features, posing challenges, LGBM simplifies this by automatically handling categorical features during training. This eliminates the need for manual encoding, saving time and reducing the risk of introducing errors [2]. Other benefits of LGBM is its ability to prevent regularisation and overfitting, supports parallel and distributed computing, thereby reducing the time spent on hyperparameter tuning and model selection. LGBM also permits customisation of loss function, robustness in handling noisy data and outliers. Its speed, efficiency, accuracy, automatic handling of categorical features, regularisation options, support for custom loss functions, and resilience to noisy data, making it a versatile tool for a wide range of binary datasets. When aiming for optimal performance and efficiency in binary classification, LightGBM is an impressive option worthy of selection.

**2.1.3 Random forest classifier.** RF classifier builds a set of decision trees to train its classifier. A distinct subset of the dataset is used to train each of these trees, and the algorithm randomly chooses which feature to take into account at each decision point within each tree [6]. This intentional randomness keeps the model from




Figure 1: Flowchart illustration of the workflow architecture.

overfitting the training set and makes it more capable of making predictions that generalise effectively to new data. During the classification tasks, the classifier aggregates the votes from each tree and decides which class to use as the final prediction. RF provide some level of interpretability, as a collection of decision trees, which makes it useful for comprehending how the model generates its predictions. Similarly, RF classifiers can handle huge datasets with a variety of characteristics. Because of their versatility and scalability, they are a great choice for dealing with a range of data quantities and complexities [6]. However, despite the benefits of RF, when working with large datasets, training a RF classifier can take some time. Longer training times are a result of the model's ensemble structure, which includes many decision trees. Therefore, it's crucial to take training time and computational resource limits into account when using RF. Also RF are inconsistent because the features and data subsets are chosen at random at each split. As a result, outcomes that differ slightly could be obtained if the model is trained multiple times using the same set of data. It can be more difficult to evaluate and debug models when there is randomness present.

While these algorithms represent the core of the model development, we also delved into the evaluation of other noteworthy contenders such as ExtraTrees, Gradient Boosting, and XGBoost. Each of these algorithms have set of distinctive capabilities, and the comparison aims to discern the most proficient and suitable algorithm for the detection of illicit account on the Ethereum blockchain network.

## 2.2 Dataset description and preprocessing

The Ethereum blockchain transactional data provided by Kaggle [9] contains the transaction specifics of the blockchain network. Such as, the sender and recipient addresses, transaction values, gas fees, and the total transaction history. Transaction history features provide information about Ethereum accounts' financial actions. A detailed description of the dataset and its features is provided in [3, 9]. The comprehensive features of the dataset form the foundation for Ethereum account behaviour analysis and illicit activity detection.

A data pre-processing step to deal with feature selection and missing values in the dataset is carried out. The first step in the data cleaning is to drop categorical features by reviewing each of the column of the data to determine if they are relevant for the analysis. This step is crucial in ensuring the data comprise exclusively numerical features and filtering unwanted data that may introduce noise or complexity in the model. The next step is to replace missing values with an appropriate statistical measure, the median. In this scenario, the median assumes a pivotal role as a robust indicator of central tendency for filling in the gaps within

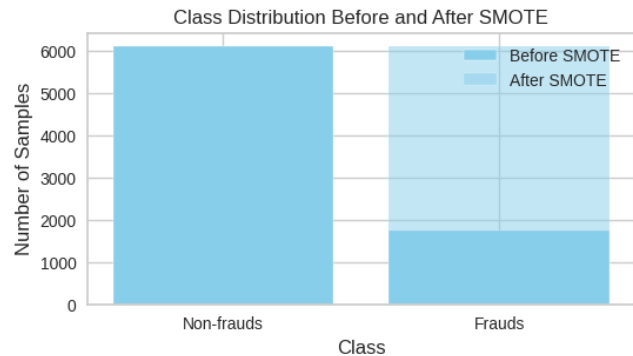


Figure 2: Class distribution before and after SMOTE.

numerical data points, thereby mitigating the potential impact of outliers or extreme values. The next step in the data preprocessing is to filter and drop features with zero-variance. This is to optimise the dataset and avert any deleterious impact on model performance due to extraneous noise.

The final step is to handle the data imbalance. An imbalanced data distribution can result in models displaying excessive bias toward the dominant class. Imbalanced datasets are characterised by a substantial underrepresentation of one class (fraudulent) compared to the other class (non-fraudulent accounts), leading to biased model outcomes. To address this challenge, the application of SMOTE was implemented. SMOTE is a resampling technique aimed at rectifying class imbalance by oversampling the minority class [11]. Instead of merely replicating existing data points, SMOTE generates synthetic samples that closely resemble the existing minority class samples. These synthetic samples contribute to balancing the class distribution within the dataset. Utilising the FLAG feature from the Kaggle Ethereum dataset, the original data instances was 7662 for non-fraudulent and 2179 for fraudulent. After data cleaning, the non-fraudulent instances reduced to 6115, and fraudulent at 1757. The application of the SMOTE technique, increased the fraudulent instances to 6116. This is to establish a balanced dataset for training the ML models and mitigate issues associated with class imbalance.

## 2.3 Experimental setup and Evaluation Metrics

The computation to train and test the developed frameworks is performed on Google Colaboratory [4] using Intel Core i7-CPU, 16 GB RAM and 64-bit operating system. The dataset was split into 80:20 ratio of training and testing data.

As the performance of ML models is largely dependant on the hyperparameter tuning, GridSearchCV was utilised to facilitate the search for the most suitable hyperparameter for each of the algorithm. GridSearchCV is a widely employed technique for systematically searching through a predefined set of hyperparameter combinations. It automates the process of hyperparameter tuning by performing an exhaustive search, evaluating the model's performance using cross-validation for each combination of hyperparameters. It returns the best combination of hyperparameters that achieves the highest score on the cross-validation, ultimately enhancing the model's ability to effectively identify fraudulent transactions.

To assess the models' performance and capabilities, five metrics were employed. These are Precision, Recall, Accuracy, Weighted F1-Score and the receiver operating characteristic (ROC) curve. Confusion matrix is the quantitative assessments used to measure the effectiveness of a classification model. The confusion matrix is a tabular representation detailing the counts of true positives, false negatives, false positives, and true negatives in the context of a binary or multiclass classification task.

Accuracy,  $A$ , is one of the most straightforward and intuitive measures to assess a model's performance. It provides a clear indication of how often the model's predictions are correct and is widely used in ML evaluation. It is the ratio of correct predictions to all predictions. It can be calculated as follows:

$$A = \frac{TP + TN}{TN + FN + TP + FP} \quad (1)$$

where  $TN$  is true negatives, which is the number of correctly predicted non-illicit accounts.  $FP$  is the false positives and depicts the number of non-illicit accounts incorrectly predicted as illicit.  $FN$  is false negatives, which is the number of illicit accounts incorrectly predicted as non-illicit;  $TP$  is true positives that predicts the number of correctly predicted illicit accounts.

Precision,  $P$ , also referred to as the positive predictive value, quantifies the ratio of correct positive predictions in comparison to all the positive predictions made.

$$P = \frac{TP}{TP + FP} \quad (2)$$

Recall,  $R$ , also recognised as sensitivity, represents the proportion of correct positive predictions out of all the actual positive instances.

$$R = \frac{TP}{TP + FN} \quad (3)$$

The F1-score is a measure of balance between precision and recall and is calculated as the harmonic mean of these two metrics. It is expressed as:

$$F1 = \frac{2 * Precision * Recall}{Precision + Recall} \quad (4)$$

Finally, the ROC, which is the measure of separability, that measures the model capability to differentiate between classes is defined as:

$$ROC = 1 - \frac{TN}{TN + FP} \quad (5)$$

### 3 RESULTS AND DISCUSSION

In this section, a comparison analysis of the selected model is performed to identify the most appropriate model to detect illicit account on the Ethereum blockchain network, using the defined evaluation metrics.

Fig. 3 shows the confusion matrix for four of the trained models. Analysing the results from the confusion matrix, the AdaBoost model, Fig. 3c, demonstrated commendable performance in correctly identifying 417 cases of fraud (True Positives). However, it also raised false alarms, flagging 644 cases (False Positives) as fraud when they were not. The False Negatives signify missed opportunities to detect fraudulent activities. The LGBM, Fig. 3b, classifier emerged as a promising candidate. With LGBM, both False Positives and False Negatives were significantly reduced, improving both recall and precision. Further model results revealed that the XGBoost, Fig. 3d, and Random Forest, Fig. 3a, classifiers exhibited effective results. They managed to reduce both False Positives and False Negatives, thereby enhancing both recall and precision. Each of these models exhibited varying degrees of effectiveness, with certain trade-offs in terms of missed fraud cases. The RF classifier, for example, missed 29 while the LGBM classifier demonstrated considerable improvement. The choice of the most suitable model depend on the specific trade-offs and priorities highlighted in Table 1.

**Table 1: Performance values of all selected ML models**

Model	Accuracy	Precision	Recall	F1-Score	ROC AUC
AdaBoost	0.670	0.393	0.988	0.562	0.634
LGBM	0.984	0.958	0.967	0.962	0.997
Random Forest	0.904	0.709	0.931	0.805	0.966
ExtraTrees	0.969	0.925	0.931	0.644	0.949
Gradient Boosting	0.768	0.480	0.979	0.644	0.949
XGB	0.976	0.931	0.960	0.945	0.994

Table 1 summarises the essential performance metrics for each of the ML algorithms. The metrics include accuracy, precision, recall, F1-score, and the ROC area under the curve (ROC AUC). These results shed light on how well each model performed in identifying illegitimate and legitimate Ethereum accounts. Particularly, the LGBMClassifier achieves the highest overall performance in detecting illicit accounts on the Ethereum blockchain. This classifier demonstrated effectiveness in identifying fraudulent accounts while maintaining a well-balanced approach to precision and recall. Specifically, the LGBMClassifier achieved an accuracy of 98.4%, indicating that it accurately classified the vast majority of transactions. Its precision score of 95.8% highlights its ability to make precise predictions, minimising false positives. Moreover, with a recall score of 96.7%, the LGBMClassifier effectively detected a significant portion of illicit accounts, minimising the number of false negatives. The impressive F1-score of 96.2% signifies a harmonious balance between precision and recall, which is particularly valuable in fraud detection scenarios where both minimising false alarms and capturing fraudulent activities are crucial. Additionally, the

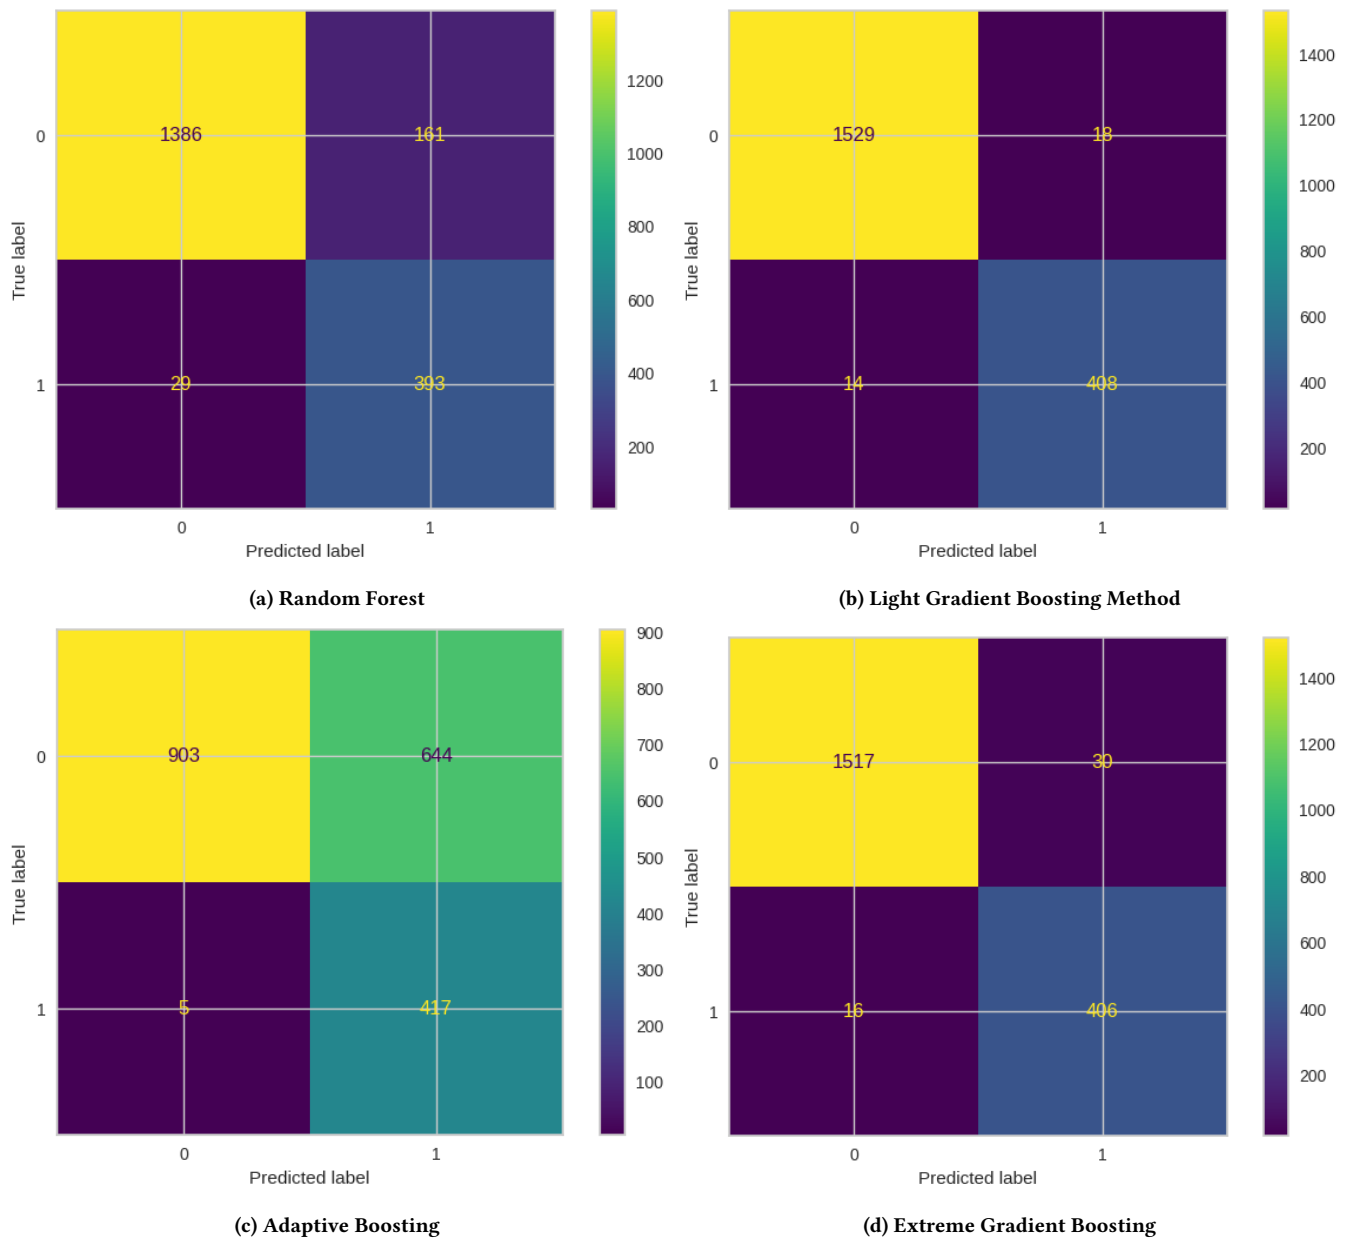
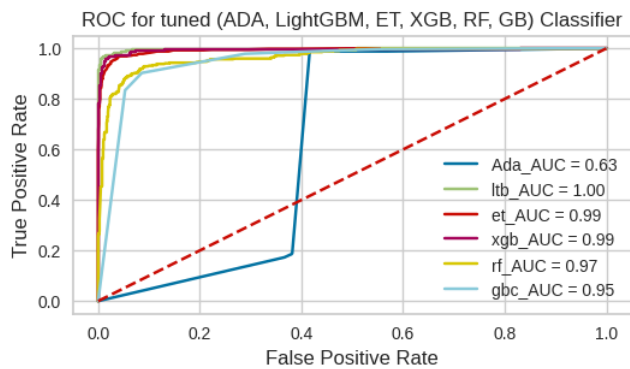


Figure 3: Confusion matrix of four different ML algorithms

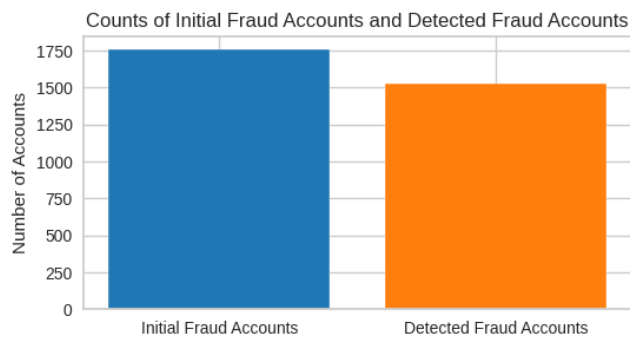
high ROC AUC score of 99.7% underscores the model’s excellent ability to distinguish between positive and negative cases. Fig. 4 illustrates the ROC curve for all the selected ML models.

**3.0.1 Light gradient boosting method.** In pursuit of detecting illicit accounts on the Ethereum blockchain, the LGBM classifier has been selected as the preferred model. To implement the model, all the previous steps involved in the data pre-processing, data splitting into training/testing sets, feature selection and scaling, and training is applied. Fig. 5 shows the total number of detected fraudulent accounts as compared to the overall fraudulent accounts.

The use of ML algorithms for the Ethereum blockchain’s illegal account detection has wider ramifications for cybersecurity. The process offers useful approaches and insights that may be used to strengthen digital security across multiple areas in addition to helping to defend blockchain networks. The proactive detection and mitigation of threats are crucial in the field of cybersecurity. Beyond the blockchain, the strategies and models used during this project can be used to protect vital infrastructure, financial institutions, and private data. Organisations can gain a proactive edge in identifying and responding to cyber threats in real time by utilising ML, potentially reducing the effect of security breaches. Additionally,



**Figure 4: Receiver Operating Characteristic (ROC) curve for the ML models.**



**Figure 5: Performance comparison of the developed frameworks.**

the combined impact of blockchain technology and ML is in line with the rising desire for decentralised and trustless systems. The necessity for strong security measures is becoming more obvious as blockchain technology continues to grow in prominence in a variety of industries and management. Enhancing security in decentralised apps and ecosystems can be modelled after the lessons discovered and approaches created in the context of identifying fraudulent accounts. In this age of rapid digital change and a constantly growing array of security risks, the partnership between ML and cybersecurity brings about an exciting shift in how we approach these challenges. It makes it possible for people and businesses to foresee and counteract the actions of those with malicious intent, ultimately paving the way for a more secure and reliable digital economy.

## 4 CONCLUSION

This work evaluated six different ML models for detecting illicit activity on the Ethereum blockchain network. The results were compared using several performance metrics. The results holds a significant promise for enhancing the security and trustworthiness of the blockchain ecosystem. Throughout the exploration, remarkable accuracy rates were achieved, with LGBM model achieving 98% accuracy. Precision and recall scores also demonstrated the ability

of these models to strike a balance between minimising false alarms and correctly identifying illicit activities. The robust F1-scores further underscored the effectiveness of the machine learning models in achieving this balance.

The future work will focus on the development of combined models that take advantage of various classifiers is one viable route. By merging the decision-making processes of several models, combined methods have the potential to further improve detection accuracy and robustness. Additionally, real-time data streams are of the most significance in the developing field of blockchain security. Blockchain transactions may be continuously monitored and analysed to produce timely alerts and responses to potential risks. These streams can greatly enhance the ML models' capacity to adjust to shifting fraudulent behaviour patterns.

## REFERENCES

- [1] Rachit Agarwal, Shikhar Barve, and Sandeep Kumar Shukla. 2021. Detecting malicious accounts in permissionless blockchains using temporal graph properties. *Applied Network Science* 6, 1 (2021), 1–30.
- [2] Weili Chen, Xiongfeng Guo, Zhiguang Chen, Zibin Zheng, and Yutong Lu. 2020. Phishing Scam Detection on Ethereum: Towards Financial Security for Blockchain Ecosystem.. In *IJCAI*, Vol. 7. 4456–4462.
- [3] Steven Farrugia, Joshua Ellul, and George Azzopardi. 2020. Detection of illicit accounts over the Ethereum blockchain. *Expert Systems with Applications* 150 (2020), 113318.
- [4] Google. 2023. Welcome to Colaboratory. Available at <https://colab.research.google.com/>, Accessed: 2023-9-16.
- [5] Gousia Habib, Sparsh Sharma, Sara Ibrahim, Imtiaz Ahmad, Shaima Qureshi, and Malik Ishfaq. 2022. Blockchain technology: benefits, challenges, applications, and integration of blockchain technology with cloud computing. *Future Internet* 14, 11 (2022), 341.
- [6] Rahmeh Fawaz Ibrahim, Aseel Mohammad Elian, and Mohammed Ababneh. 2021. Illicit account detection in the ethereum blockchain using machine learning. In *2021 Intl. Conf. Info. Tech. (ICIT)*. IEEE, 488–493.
- [7] Olamide Jogunola, Bamidele Adebisi, Khoa Van Hoang, Yakubu Tsado, Segun I Popoola, Mohammad Hammoudeh, and Raheel Nawaz. 2022. CBLSTM-AE: a hybrid deep learning framework for predicting energy consumption. *Energies* 15, 3 (2022), 810.
- [8] Olamide Jogunola, Bamidele Adebisi, Augustine Ikpehai, Segun I Popoola, Guan Gui, Haris Gačanin, and Song Ci. 2020. Consensus algorithms and deep reinforcement learning in energy market: A review. *IEEE Internet of Things Journal* 8, 6 (2020), 4211–4227.
- [9] Kaggle. 2020. Ethereum fraud detection dataset. Available at <https://www.kaggle.com/datasets/vagifa/ethereum-frauddetection-dataset>, Accessed: 2023-11-1.
- [10] Pranav Nerurkar, Yann Busnel, Romaric Ludinard, Kunjal Shah, Sunil Bhirud, and Dhiren Patel. 2020. Detecting illicit entities in bitcoin using supervised learning of ensemble decision trees. In *10th Intl. Conf. Info. Comm. and Mgt.* 25–30.
- [11] Farimah Poursafaei, Ghaith Bany Hamad, and Zeljko Zilic. 2020. Detecting malicious Ethereum entities via application of machine learning classification. In *2nd Conf. Blockchain Research & App. for Innovative Networks and Services (BRAINS)*. IEEE, 120–127.
- [12] Adam Brian Turner, Stephen McCombie, and Allon J Uhlmann. 2020. Analysis techniques for illicit bitcoin transactions. *Frontiers in Computer Science* 2 (2020), 600596.