

A SUPERVISED MACHINE LEARNING APPROACH TO DE-ANONYMIZING THE BITCOIN BLOCKCHAIN

Dr. S. V. Achuta Rao¹, M. Pavan², G. Sai Venkat², P. Shiva², M. Durga Prasad², K. Mokshith²

¹Professor, ²UG Student, ^{1,2}Department of Computer Science and Engineering (CSBS),

^{1,2}Sree Dattha Institute of Engineering and Science, Sheriguda, Ibrahimpatnam, 501510, Telangana.

Received: 09-07-2025

Accepted: 23-08-2025

Published: 30-08-2025

Abstract

Decentralized Finance (DeFi) is revolutionizing financial systems by eliminating intermediaries and enabling peer-to-peer transactions through blockchain technology. It enhances transparency, security, and accessibility, allowing users to access financial services such as lending, borrowing, and trading without relying on centralized institutions. Predictions for DeFi indicate exponential growth, with the integration of AI and machine learning driving significant advancements in fraud detection, risk assessment, and transaction analysis. Before the integration of AI, financial fraud detection primarily relied on rule-based systems, manual audits, and traditional statistical models—methods that lacked adaptability and real-time decision-making capabilities. Legacy systems such as credit scoring models and transaction monitoring frameworks struggled with scalability and required continuous human intervention. As fraudulent activities and cyber threats have grown more sophisticated, the limitations of traditional solutions have become increasingly apparent, underscoring the need for AI-driven approaches. By leveraging machine learning, transaction patterns can be analyzed with greater accuracy, enabling real-time anomaly detection and significantly reducing financial risk. The motivation behind this development is to enhance security, improve accuracy in transaction classification, and deliver scalable financial crime detection solutions. Conventional fraud detection mechanisms often fail to keep pace with evolving threats, resulting in considerable financial losses. Manual reviews are time-consuming and error-prone, while static models lack adaptability to emerging fraudulent behaviors. Machine learning enables real-time monitoring and predictive analysis, allowing financial institutions to detect suspicious activities with heightened precision. The proposed system integrates a combination of classifiers—including decision trees, logistic regression, AdaBoost, gradient boosting, k-nearest neighbors, and random forest algorithms—to improve transaction classification accuracy. AI-driven analysis enhances fraud detection by learning from historical data, reducing false positives, and enabling automated de-anonymization of transactions. This system applies advanced algorithms to identify fraudulent patterns, optimize financial security, and streamline transaction verification. By automating the process, AI-powered models offer a more robust and efficient approach to securing financial transactions, ensuring greater reliability and trust in decentralized finance.

Keywords: Decentralized Finance, DeFi, blockchain, peer-to-peer transactions, financial services, AI integration, machine learning, fraud detection, risk assessment, transaction analysis, rule-based systems, manual audits, statistical models, credit scoring, transaction monitoring, cyber threats

1. INTRODUCTION

The research focuses on enhancing financial security in decentralized finance (DeFi) by integrating machine learning for fraud detection and transaction classification. With the rise of DeFi, financial transactions have become more susceptible to fraudulent activities, money laundering, and cyber

threats. Traditional rule-based systems struggle to identify sophisticated fraudulent patterns, leading to increased financial risks. This system utilizes multiple machine learning algorithms, including decision trees, logistic regression, AdaBoost, gradient boosting, k-nearest neighbors, and random forests, to analyze transaction data and detect anomalies

with high accuracy. The approach ensures automated, real-time fraud detection by analyzing large datasets and identifying suspicious transaction behaviors. By leveraging AI, financial security is enhanced through advanced classification techniques that minimize false positives and improve detection efficiency. The system allows DeFi platforms to de-anonymize transactions and enhance regulatory compliance without compromising decentralization. The implementation provides a scalable, accurate, and efficient solution to financial risk management, ensuring trust and transparency in decentralized financial transactions. This research aims to enhance financial security in decentralized finance (DeFi) by developing an intelligent, machine learning-based fraud detection system that classifies transactions and identifies anomalies in real time. With DeFi's growing adoption, traditional rule-based fraud detection methods are increasingly ineffective against sophisticated financial crimes such as money laundering and cyber fraud. By leveraging supervised



Figure 1: Generalized hash code working machine learning algorithms—including decision trees, logistic regression, random forests, AdaBoost, gradient boosting, and k-nearest neighbors—the proposed system analyzes large volumes of transaction data, achieving high accuracy while minimizing false positives and negatives. This AI-driven approach ensures scalable, automated, and adaptive fraud detection that supports regulatory compliance, de-anonymizes suspicious transaction patterns, and strengthens the trust, transparency, and credibility of DeFi platforms without compromising decentralization. The research addresses key challenges like real-time

monitoring, behavioral pattern recognition, and scalability, making it a transformative solution for secure and efficient financial risk management in blockchain-based ecosystems.

2. LITERATURE SURVEY

Machine learning (ML) methods adapted from among deep learning algorithms have been recently applied to financial time series prediction with a number of publications in computer science journals (Greff et al. 2017; Fe-Fei et al. 2003; Zhang et al. 2018), as well as in economics and finance journals (Koutmos 2018; Kristoufek 2018). There is a gap in the existing literature, however, which is pronounced in the uncovered field of the applications of machine learning methods for time series to cryptocurrency trading data. In this work, we aim to provide a benchmark as to how efficient the modern ML algorithms can be in view of their applicability to the high-frequency trading data on the minute scale. The application of deep learning techniques faces a difficult trade-off: deep learning algorithms require a large number of data samples to learn from, implying in practice high-frequency data, such as minute-sampled trade records, whereas the training patterns over long periods are not always stationary, meaning varying patterns may be extracted from different segments of the training dataset. The applicability of deep learning to high-frequency market prediction is still an open problem. Recently, some empirical results (Mäkinen et al. 2018; Sirignano and Cont 2018; Zhang et al. 2018) with deep learning algorithms showed that there might be a universal price formulation for the deterministic part of trading behavior to some degree, which implies financial data at high frequency exhibit some stylized facts and could possess learnable patterns that are stationary over long time periods. The aforementioned references used order-driven data (limit order book) and trained recurrent neural networks with the huge number of data. In this paper, we take a

different approach: we provide a metric learning-based (Cinbis et al. 2011; Koch 2015; Vinyals et al. 2016; Xing et al. 2003) method, which we call the random sampling method (RSM). We measure the similarity between the input pattern and the training samples with the novel sampling scheme, which we describe below. Then, the label of the most similar data point becomes an output candidate for the prediction of our model.

The present approach is motivated by the highly non-stationary dynamics in digital assets as volatile as cryptocurrencies, in particular Bitcoin. State-of-the-art deep learning algorithms for time series, such as the long short-term memory (LSTM) method (Gers et al. 2000; Hochreiter and Schmidhuber 1997) require large datasets for training, and thus suffer from the fact that the causal patterns in the cryptocurrency time series may change quite substantially in the training and testing datasets, resulting therefore in insufficient prediction performance, noise fitting, and inconsistent results. For Bitcoin, recent data patterns are more relevant for trend prediction than more distant data, which practically limits the number of samples for each class. Here, we therefore adapt the metric learning method in which the algorithm finds the best recent patterns to be labeled for optimal prediction (Fe-Fei et al. 2003; Lake et al. 2014, 2011, 2015; Li et al. 2006). The works in (Graves et al. 2014; Koch 2015; Santoro et al. 2016; Vinyals et al. 2016) showed how to deploy deep learning algorithms for such purposes in various applications. Our approach was broadly inspired by recent deep learning (DL) developments in the field of image processing.

3. PROPOSED SYSTEM

The proposed system aims to enhance the prediction and analysis of decentralized finance (DeFi) transactions through the integration of machine learning, specifically Convolutional Neural Networks (CNN). Traditional financial systems have relied on

rule-based algorithms and statistical models that often struggle to accurately capture complex patterns in decentralized finance data. These methods have limitations in terms of scalability, adaptability, and accuracy in dynamic markets such as DeFi, where transactions, user behaviors, and trends evolve rapidly. This system leverages the power of CNN to address these limitations. CNNs are well-suited for extracting meaningful features from large and complex datasets, making them ideal for analyzing DeFi data, which often includes high-dimensional and time-series information. The first step in the proposed system is to gather DeFi transaction datasets, which include information about user activities, token transfers, smart contract interactions, and more. These datasets are then preprocessed to remove null values, handle categorical data through techniques like label encoding, and ensure the data is in a format that can be used by machine learning algorithms.

ML Model Building

Building the machine learning model starts with selecting an appropriate algorithm. In this case, CNN is chosen due to its ability to recognize complex patterns and perform well with high-dimensional data. The model is then trained on the preprocessed dataset. Hyperparameters, such as the number of layers, filter size, and learning rate, are optimized for best performance. The model is evaluated on the testing data, and metrics like accuracy, precision, recall, and F1-score are calculated. Based on the results, the model might undergo further training or fine-tuning to achieve optimal performance.

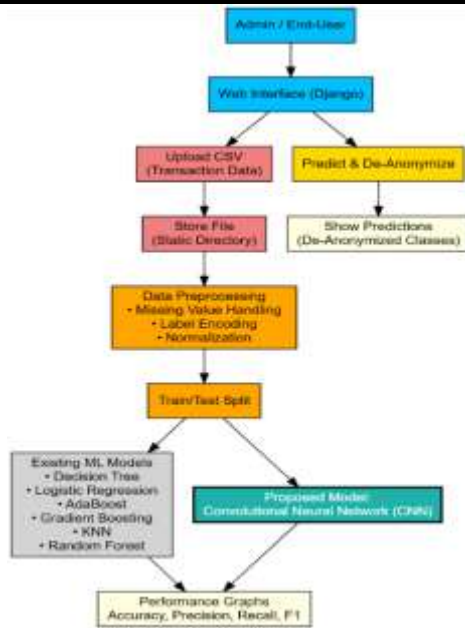


Figure 2: Block Diagram for De-anonymizing the bitcoin blockchain.

What is CNN

A Convolutional Neural Network (CNN) is a specialized deep learning architecture designed to process and analyze data with a grid-like topology, such as images. CNNs are particularly adept at identifying patterns and structures within visual data, making them essential for tasks like image classification, object detection, and facial recognition. By leveraging convolutional operations, CNNs can automatically learn hierarchical feature representations from raw input data, reducing the need for manual feature extraction.

How It Works:

CNNs operate through a series of layers that transform the input data into increasingly abstract representations: Convolutional Layers apply convolutional filters to the input data, detecting local patterns such as edges, textures, and shapes. Each filter focuses on a specific feature, and the resulting feature maps highlight the presence of these patterns across the input. Activation Layers, typically using the Rectified Linear Unit (ReLU) function, introduce non-linearity, enabling the network to learn complex patterns and relationships within the data. Pooling Layers perform downsampling operations, such as max

pooling or average pooling, to reduce the spatial dimensions of the feature maps, decreasing computational load and helping achieve spatial invariance. Fully Connected Layers, included after several convolutional and pooling layers, integrate the features learned by the previous layers to make final predictions or classifications

Architecture:

The architecture of a CNN typically consists of the following components:

- **Input Layer:** Accepts the raw data, such as an image, represented as a grid of pixel values.
- **Convolutional Layers:** Apply multiple filters to the input data to extract various features.
- **Activation Layers:** Introduce non-linearity to the network, allowing it to learn complex patterns.
- **Pooling Layers:** Reduce the spatial dimensions of the feature maps, retaining essential information while decreasing computational complexity.
- **Fully Connected Layers:** Integrate the features extracted by the convolutional and pooling layers to perform classification or regression tasks.
- **Output Layer:** Produces the final prediction or classification result.

Advantages:

CNNs offer several key advantages that make them highly effective for analyzing complex data: they enable automatic feature extraction by learning relevant patterns directly from raw input data, eliminating the need for manual feature engineering. Their hierarchical structure captures spatial hierarchies, making CNNs particularly suitable for image and video analysis. Through parameter sharing—using the same filters across different regions—they significantly reduce the number of parameters, enhancing computational efficiency. Additionally, CNNs exhibit robustness to input variations such as

translations, rotations, and scaling, thanks to their convolutional and pooling operations, making them adaptable to diverse and noisy datasets.

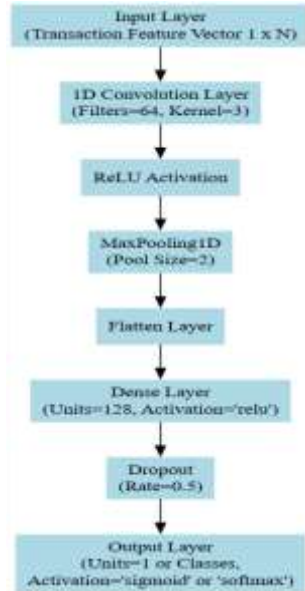


Figure 3: Working architecture of CNN model.

4. RESULTS

The below Figure shows the header and hero section of a DeFi platform's website. The design emphasizes its focus on decentralized finance and blockchain technology, with a strong association with Bitcoin. The website likely provides users with access to various DeFi services, potentially including cryptocurrency trading, lending, borrowing, or other financial tools. The imagery and navigation suggest a focus on user interaction and information dissemination.



Figure 4: Illustration of GUI interface showing home page.



Figure 5: Illustration of GUI interface showing login page.

The login form, accessible via the "Login" button on the homepage, presents a straightforward and secure interface for returning users. It consists of two primary input fields: "Username" and "Password." The "Username" field prompts users to enter the unique username they created during the signup process. The "Password" field, appropriately masked for security, requires users to input the corresponding password associated with their account. Below these fields, a prominent "login" button allows users to submit their credentials for verification. Upon successful authentication, users are granted access to their personalized accounts and the platform's mental health support services.



Figure 6: Illustration of GUI interface showing uploading dataset.

This Figure shows that users to upload datasets for analysis within the context of a DeFi platform. The application supports data processing and machine learning model training, suggesting its use for tasks like fraud detection, market analysis, or other predictive modeling in the DeFi space.



Figure 7: Illustration of GUI interface showing pre-processed dataset.

This Figure displays a table of numerical data, likely a portion of a larger dataset used in a machine learning or statistical analysis context. The table consists of rows and columns, with the top row labeled 0 through 7, serving as column headers or feature indices. Each subsequent row represents an instance or observation, with the corresponding values for each feature displayed in the cells. The data is primarily numerical, including both integers and decimal values, suggesting that it has likely been pre-processed or transformed for analysis. The varying scales and ranges of the numbers across different columns indicate that the features represent different types of measurements or characteristics. Without further context regarding the origin and purpose of the data, it's difficult to provide a specific interpretation of the values. However, the structure and format suggest that this data is likely used for training a predictive model, exploring relationships between variables, or some other form of quantitative analysis.




Figure 8: Illustration of GUI interface showing comparison of the performance of algorithms. This Figure compares the performance of six machine learning algorithms using a table and a bar chart. The CNN appears to be the best performing model based on the provided

metrics. The graph complements the table, providing a quick visual comparison of the algorithms. It's important to note that this comparison is based on a specific dataset and task. Further analysis and validation might be needed to assess the models' generalizability and suitability for real-world applications.



Figure 9: Illustration of GUI interface of uploading test dataset.

The Figure shows that users to upload data for transaction de-anonymization within a DeFi platform. The selected file, "testData," will likely be processed and potentially used for training machine learning models to identify the parties involved in anonymous cryptocurrency transactions. The "De-Anonymize Transaction Page" label clearly indicates the page's purpose.



Figure 10: Illustration of GUI interface showing predicted outcomes.

This Figure shows the predicted outcomes of a classification model for two different input data blocks. The model predicts one block as "exchange" and the other as "gambling," suggesting its potential use in identifying or categorizing activities based on their features. The specific context and meaning of the data would require further investigation.

5. CONCLUSION

In the realm of financial fraud detection, the integration of Convolutional Neural Networks

(CNNs) has significantly enhanced the accuracy and efficiency of identifying fraudulent activities. Traditional methods, such as Decision Trees, have been foundational in this field; however, they often struggle with complex, high-dimensional data and intricate patterns inherent in financial transactions. CNNs, with their advanced feature extraction capabilities, have addressed these challenges by effectively capturing spatial hierarchies and local dependencies within transaction data. The application of CNNs in fraud detection has demonstrated superior performance in various studies. For instance, a study published in the *Journal of Computer Science and Software Engineering* highlighted that CNNs could more effectively identify abnormal patterns in financial transaction data, achieving higher accuracy, recall, precision, and F1 scores compared to traditional methods. CNNs have shown promise in processing unstructured data, such as images and sequences, which are increasingly prevalent in financial transactions. This capability allows for a more comprehensive analysis of diverse data types, leading to more robust fraud detection systems. The adaptability of CNNs to various data forms and their proficiency in learning complex patterns make them a valuable tool in the ongoing battle against financial fraud.

References

1. Andrychowicz, Marcin, Misha Denil, Sergio Gomez, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando De Freitas. 2016. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*. Cambridge: The MIT Press, pp. 3981–89.
2. Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv*, arXiv:1409.0473.
3. Bariviera, Aurelio F., Maria Jose Basgall, Waldo Hasperue, and Marcelo Naiouf. 2017. Some stylized facts of the bitcoin market. *Physica A: Statistical Mechanics and its Applications* 484: 82–90.
4. Cinbis, Ramazan Gokberk, Jakob Verbeek, and Cordelia Schmid. 2011. Unsupervised metric learning for face identification in tv video. Paper presented at the 2011 IEEE International Conference on Computer Vision (ICCV), Barcelona, Spain, November 6–13; pp. 1559–66.
5. Dixon, Matthew, Diego Klabjan, and Jin Hoon Bang. 2017. Classification-based financial markets prediction using deep neural networks. *Algorithmic Finance* 6: 67–77.
6. Gers, Felix A., Jürgen Schmidhuber, and Fred A. Cummins. 2000. Learning to forget: Continual prediction with lstm. *Neural Computation* 12: 2451–71. [PubMed]
7. Gkillas, Konstantinos, and Paraskevi Katsiampa. 2018. An application of extreme value theory to cryptocurrencies. *Economics Letters* 164: 109–11.
8. Gkillas, Konstantinos, Stelios Bekiros, and Costas Siriopoulos. 2018. Extreme Correlation in Cryptocurrency Markets. Available online: <https://ssrn.com/abstract=3180934> (accessed on 14 January 2019).
9. Glorot, Xavier, Antoine Bordes, and Yoshua Bengio. 2011. Deep sparse rectifier neural networks. In Paper presented at the Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, Ft. Lauderdale, FL, USA, April 11–13; pp. 315–23.

10. Graves, Alex, Greg Wayne, and Ivo Danihelka. 2014. Neural Turing machines. *arXiv*, arXiv:1410.5401.
11. Greff, Klaus, Rupesh K. Srivastava, Jan Koutník, Bas R. Steunebrink, and Jürgen Schmidhuber. 2017. Lstm: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems* 28: 2222–32.
12. Hilliard, Nathan, Lawrence Phillips, Scott Howland, Artem Yankov, Courtney D. Corley, and Nathan O. Hodas. 2018. Few-shot learning with metric-agnostic conditional embeddings. *arXiv*, arXiv:1802.04376.
13. Hochreiter, Sepp, and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9: 1735–80.
14. Kingma, Diederik P., and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv*, arXiv:1412.6980.
15. Koch, Gregory. 2015. Siamese neural networks for one-shot image recognition. Paper presented at the 32nd International Conference on Machine Learning, Lille, France, July 6–11.
16. Koutmos, Dimitrios. 2018. Bitcoin returns and transaction activity. *Economics Letters* 167: 81–85.