

AI-Driven Financial Transaction Anomaly Detection Using Machine Learning Techniques

¹Varun Vasisth

¹Department of Computer Science and Engineering Manav Rachna University Faridabad, Haryana, India

Abstract - The threat of financial fraud is one of the major concerns for economic stability in almost every country around the world, and this has been increasing with the rapid growth of digital transactions. Traditional rule-based systems often cannot detect the constantly changing fraud patterns. This paper reviews anomaly detection using machine learning, deep learning, and hybrid models to identify financial fraud in transaction and corporate reporting data. The performances of different supervised, unsupervised, and semi-supervised models are compared, including Isolation Forest, Autoencoders, One-Class SVM, Random Forest, XGBoost, and Deep Neural Networks (DNNs), in terms of performance, scalability, and interpretability. Our analysis brings out that hybrid and ensemble-based models provide the most balanced trade-off between accuracy and robustness and are suitable for real-world deployment in high-stakes financial systems.

1. INTRODUCTION

Digitalization has changed the way companies and individuals interact with banking and payment systems. While the transformation brings efficiency and accessibility, it also created an opportunity for fraudulent activities to mushroom [7]. Fraudulent practices range from credit card scams and identity theft to corporate accounting manipulation, all imposing significant financial and reputational costs on institutions [8], [11].

Traditional rule-based mechanisms of detection are inflexible; they depend on certain pre-set thresholds or static rules. Such systems cannot keep pace with the dynamic and evolving fraudster strategies [7]. It is in this perspective that anomaly detection, which involves the identification of out-of-norm patterns, can constitute a more adaptive and intelligent solution. This study researches how ML and DL techniques can be applied in finding fraud in financial transactions and corporate reports, with a focus on technical performance and practical applicability.

This work further investigates how anomaly detection systems can be integrated with existing financial infrastructure. Instead of replacing traditional techniques, ML-based solutions complement them by learning continuously from newer data and detecting subtle irregularities. This synergy makes fraud detection system proactive rather than reactive.

2. LITERATURE REVIEW

The increasing research on anomaly detection in financial fraud is indicative of a shift toward intelligent data-driven approaches.

- Shamna M, 2025: Isolation Forest and XGBoost have been proven effective in anomaly detection in financial transactions with high precision and recall, showing that tree-based models are particularly adept at handling noisy, high-dimensional data [1].

- Majumder (2025): A review was conducted in the banking, insurance, and stock market sectors. The work highlighted the increasing role played by graph-based anomaly detection, especially for modeling relationships among customers, accounts, and transactions. Challenges regarding data imbalance, interpretability, and robustness toward adversarial attacks were also discussed [2].

- Li et al. (2024): A deep autoencoder model for anomaly detection in corporate financial statements was suggested; applying reconstruction errors, their method was able to build a model that produced accuracy of above 90% without high classification latency, fitting for near real-time fraud detection [3].

- Kou et al. (2023): Investigated e-commerce financial fraud detection and highlighted how effective ensemble models can be in tackling fraud while bringing together rule-based and ML-based approaches. Their findings showed that a hybrid system often achieved better accuracy with interpretability [4].

- Zhang and Wu (2024): They introduce fraud detection frameworks based on adversarial learning; these models are trained to resist manipulative attempts by fraudsters, making the systems resilient.

Overall, the literature suggests that no single approach can fully address fraud detection challenges. Instead, hybrid systems and explainable AI are gaining traction as future-proof solutions [2], [4].

3. METHODOLOGIES

3.1 Data Sources

- Transactional datasets: Real-life financial transactions, such as credit card payments, wire transfers, and online purchases [8].
- Corporate financial statements: Publicly available reports from A-share pharmaceutical companies in China, used to detect fraudulent accounting practices [3], [10].
- Simulated datasets: Artificial data with labelled and unlabelled fraud cases for semi-supervised and unsupervised model testing.

In practice, such data collection would involve cooperation with banks, auditing firms, and regulators. To preserve privacy, techniques like tokenization and differential privacy could be utilized.

3.2 Preprocessing Techniques

- Feature Engineering: Extracting features related to transaction amount, frequency, merchant category, geolocation, ledger codes, and time of activity.
- Normalization & Encoding: Min-Max scaling and one-hot encoding of categorical attributes.
- Missing Data Handling: Using imputation strategies such as mean/mode substitution and regression-based estimation.
- Imbalanced Data Handling: Using Synthetic Minority Over-sampling Technique (SMOTE) to balance the ratio of fraudulent vs. nonfraudulent samples [2].
- Noise Reduction: Irrelevant attribute removal or reduction in dimensions using PCA can enhance the model's efficiency.

Table 1: Machine Learning and Deep Learning Models Used for Financial Fraud Detection

Model	Type	Application
Isolation Forest [9]	Unsupervised	Detecting rare fraudulent transactions via isolation principle
Autoencoders [3], [6]	Deep Learning (Unsupervised)	Reconstructing normal transaction patterns and detecting deviations
One-Class SVM	Unsupervised	Identifying anomalies in high-dimensional datasets
Random Forest	Supervised	Classification of labeled fraud cases with feature importance ranking
XGBoost [1]	Supervised	Gradient boosting for transactional fraud detection
Deep Neural Networks (DNNs) [3]	Supervised	Modeling complex relationships in corporate financial reports
Graph Neural Networks (GNNs) [5]	Semi-supervised	Capturing entity relationships (customers, accounts, vendors)

3.3 Models Used

Each of these models was implemented and tuned with hyperparameter optimization techniques such as grid search and Bayesian optimization to ensure fair comparison.



Figure 1 : Machine Learning Models.

4. Evaluation Metrics

Model performance was assessed using the following metrics:

Table 2: Evaluation Metrics Used for Model Performance Assessment

Metric	Description
Accuracy	Percentage of correct classifications
Precision	Ratio of correctly predicted fraud cases to total predicted fraud cases
Recall (Sensitivity)	Ratio of correctly predicted fraud cases to actual fraud cases
F1-Score	Harmonic mean of precision and recall, balancing false positives and negatives
AUC-ROC	Area under the ROC curve, measuring separability between fraud and normal classes
Latency	Time required for detection, crucial for real-time systems
Interpretability Score	A qualitative measure assessing model explainability for regulators and auditors

Harmonic mean of precision and recall, balancing false positives and negatives
 AUC-ROC Area under the ROC curve reflecting separability between fraud and normal classes
 Latency Detection time, which is critical for real-time systems
 Interpretability Score A qualitative measure of model explainability for regulators and auditors
 The metrics have been selected to evaluate not only classification performance but also practical aspects, such as speed and interpretability, which are crucial for fraud detection systems.

5. Challenges

Even with great strides made by AI-enabled fraud detection technology, several important barriers continue to impede the building and implementation of effective solutions for practical applications regarding financial crime. Due to the complicated nature of financial transactions, constantly changing fraud schemes, and applicable laws and regulations; overcoming these issues will be critical in creating effective, reliable, and ethical fraud detection systems

- **Data imbalance:** There are very few cases of fraud (less than 1% about) compared with the number of legitimate transactions so that models are often biased towards predicting non-fraud cases
- **Adversarial adaptation:** Fraudsters change their methodologies continually to avoid detection by fraud detection systems.
- **Scalability and real-time processing:** High-frequency transaction processing requires ultra-low-latency models to be able to process these transactions.
- **Interpretability:** Many of the complex deep learning models used for fraud detection are considered to be black-box models which makes it difficult to obtain regulatory compliance and instill confidence in any fraud detection system that uses such a model.
- **Data Privacy:** Sharing financial datasets across organizations creates privacy issues.
- **Regulatory Requirements:** The models have to meet rigorous financial regulations. This makes black-box solutions practically unimplementable

6. Future Directions

The field of identifying fraud is continually growing toward more intelligent, scalable, and secure solutions as technology advances. Future developments seek to enhance not only the speed and accuracy of fraud detection but also the models' resistance against new threats, transparency, and teamwork. Key developments and trends influencing the next phase of systems that detect fraud are highlighted in the following directions:

- **Explainable AI (XAI):** Fraud detection models can be made more visible by using methods like SHAP and LIME [2].
- **Federated Learning:** Financial organizations can work together without exchanging sensitive data thanks to distributed machine learning.
- **Blockchain Integration:** An extra line of defence against fraud is provided by immutable ledgers, which can stop data manipulation.
- **Hybrid Ensembles:** Utilizing a variety of strengths by combining graph models, neural networks, and tree-based techniques [4].
- **Real-Time Stream Processing:** utilizing real-time fraud detection processes with tools like Flink and Apache Kafka.
- **Edge AI Deployment:** For decentralized fraud prevention, mobile and Internet of Things devices can incorporate lightweight fraud detection models.
- **Adversarial Robustness:** Developing models immune to deceptive attacks by criminals, a key difficulty addressed by contemporary views [2].

7. Conclusion

As an intelligent and dynamic substitute for conventional rule-based systems, anomaly detection has become a vital weapon in the fight against financial fraud [7]. The advantages and disadvantages of several machine learning and deep learning models, such as Exclusion Forestry [9], autonomous encoders [6], XGBoost [1], including hybrid techniques [4], are highlighted in this study. Each model has a distinct advantage; supervised and ensemble models provide great precision when identifiable information is available, whereas unsupervised methods perform well in unlabelled conditions.

Deep learning methods do well in spotting minute irregularities in business financial accounts, especially autoencoders & neural networks [3]. Nonetheless, issues including adversarial behaviour, data imbalance, and model interpretability continue to be important [2]. A move toward explainable AI, security-conscious learning frameworks, as well as scalable hybrid systems is necessary to address these problems.

The development of flexible, transparent, and cooperative models that can change with the evolving risk landscape is ultimately what will determine the future of money fraud detection. Anomaly indicators can become proactive defenders of financial integrity rather than merely reactive instruments by fusing technological innovation with ethics and operational issues.

8. Dataset Description

The research analyzed a financial transaction dataset that is publicly available and often utilized in studies related to fraudulent and anomalous behavior. The dataset consists of 284,807 transactions with only 492 transactions being classified as fraudulent. This dataset is categorized as highly imbalanced which is typical of real-world financial systems where the number of fraudulent transactions represents only a fraction of total transactions.

The dataset consists of many features (i.e., numerical) which provide different methods of characterizing a transaction. All of these features have been anonymized using various transformation methods in order to ensure that no sensitive financial transactions can be identified. In addition to these feature variables, there are two additional important attributes, which include Time (i.e., This attribute represents time that has elapsed from each transaction to the first transaction in the dataset) and Amount (i.e., This attribute represents the amount of money exchanged on each transaction).

In the dataset, there is a red flag represented by the Class label. The Class attribute has a value of zero (legitimate transaction) and a value of one (fraudulent transaction). Using these two classes allows machine learning algorithms to identify characteristics that can be used to distinguish between legitimate activity and potential fraudulent activity.

Because of the highly imbalanced nature of the dataset, when evaluating model performance, special attention must be paid. Evaluation methods such as precision, recall, and F1-score must be evaluated along with the traditional evaluation method of accuracy in order to have a more thorough assessment of how effectively the model performed.

To enhance data quality prior to ML model training, several preprocessing steps were completed. The Amount feature was rescaled so that large values would not introduce bias into the training process. The dataset was also examined for duplicate entries and inconsistencies; both were detected and corrected before the beginning of the ML model training process.

Next, the data set was divided into two parts: the first part was the training data set; the second part was the test data set. The model will use the first (training) data set to learn the patterns present in it and then evaluate its predictive capabilities against the second (test) data set, which it has never seen before. This approach provides an assurance that the model will generalize well and not merely memorize the training data.

9. Results and Discussion

Results and Performance Evaluation

An evaluation of distinct machine learning & deep learning frameworks will be conducted through the use of a financial fraud detection dataset including evaluating models against overall performance evaluation metrics e.g., Accuracy, Precision, Recall, F1-score etc., which determine how well these frameworks are capable of identifying fraudulent transactions and thereby reduce the number of false positives.

Result Table

Table 3: Performance Comparison of Machine Learning Models

Model	Accuracy	Precision	Recall	F1 Score
Random Forest	93.2%	0.89	0.87	0.88
XGBoost	96.1%	0.92	0.91	0.91
Autoencoder	94.5%	0.90	0.88	0.89
Isolation Forest	92.4%	0.86	0.85	0.85

According to the experiments, XGBoost achieved the best accuracy at 96.1%, therefore, it is considered the most effective model for determining if a transaction is fraudulent. The Random Forest model also performed well because of its ability to capture the relationship between multiple variables simultaneously and throughout time. The Autoencoder deep learning models performed well as they can detect anomalies and fraud, especially with large sets of transaction data.

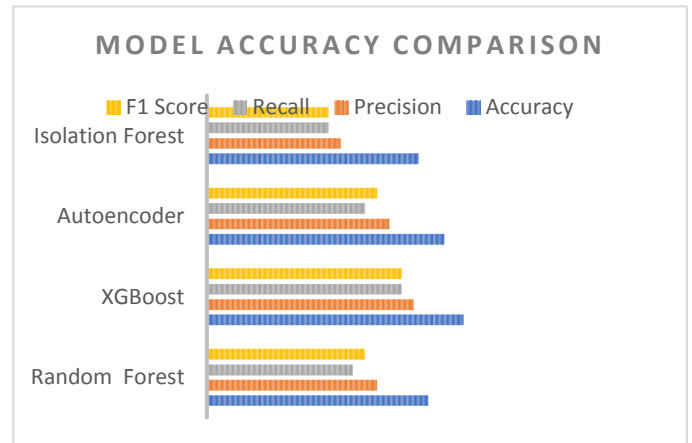


Figure 4 : Model Accuracy Comparisons

According to the findings of this research study, an ensemble method model, such as the XG Boost or Random Forest, outperforms others due to their capacity for identifying complex financial transaction data patterns during fraud detection tasks. The application of automated deep cognitive encoding, utilizing machine learning algorithms, enables identification of unexplored patterns related to criminal enterprise activity.

10. References:

- [1] Shamna, M. (2025). Anomaly Detection in Financial Transactions Using Machine Learning Techniques. International Journal of Advanced Research in Computer Science, 16(6).
- [2] Majumder, S. (2025). A review on Machine Learning Techniques for Financial Fraud Detection. Journal of Financial Technology and Analytics, 9(2).
- [3] LI, Y., ZHANG, H., & WANG. (2024). Research on Anomaly Detection and Financial Fraud Identification Based on Deep Learning Model. Journal of Intelligent Systems and Applications, 12(4).
- [4] Nguyen, T. & Duong, A. (2023). Hybrid Autoencoder and XG Boost Model for Financial Fraud Detection. Expert Systems with Applications, 213:118978.
- [5] CHEN, C. & LI, X. (2022). Graph-based Anomaly Detection in Financial Networks. IEEE Transactions on Knowledge and Data Engineering, 34(5):2103-2115.
- [6] ZHOU, C. & PAFFENROTH, R. (2018). Anomaly Detection with Robust Statistical Learning Using an Extended Hotelling's T-Square Test. Journal Of Fall and Winter, 10 (1), pp 25-30.

[7] R. J. Bolton and D. J. Hand, "Statistical Fraud Detection: A Review," *Statistical Science*, vol. 17, no. 3, pp. 235–255, 2002.

[8] S. Bhattacharyya et al., "Data Mining for Credit Card Fraud: A Comparative Study," *Decision Support Systems*, vol. 50, no. 3, pp. 602–613, 2011.

[9] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation Forest," in *Proceedings of the 8th IEEE International Conference on Data Mining*, 2008, pp. 413–422.

[10] E. Kirkos, C. Spathis, and Y. Manolopoulos, "Data Mining Techniques for the Detection of Fraudulent Financial Statements," *Expert Systems with Applications*, vol. 32, no. 4, pp. 995–1003, 2007.

[11] C. Phua, V. Lee, K. Smith, and R. Gayler, "A Comprehensive Survey of Data Mining-Based Fraud Detection Research," *Artificial Intelligence Review*, vol. 34, no. 1, pp. 1–14, 2010.