

Machine Learning-Based Identification of Financial Fraud Patterns in Digital Transaction Systems

Dr. Henry Collins

Associate Professor, Edinburgh Institute of Data Science, United Kingdom

Abstract

Digital transaction systems have transformed financial activities by facilitating electronic payments, online purchases, mobile banking operations, and automated fund transfers. These developments have increased transaction accessibility and operational efficiency while introducing additional challenges associated with fraudulent financial activities. Fraudulent transactions may exhibit complex behavioral characteristics that differ from conventional transaction patterns, making their identification difficult through fixed monitoring rules. Machine learning techniques provide computational approaches for examining transaction information, recognizing statistical irregularities, and identifying potentially fraudulent activities through predictive classification. The present paper develops a hypothetical methodological framework for examining machine learning-based identification of financial fraud patterns in digital transaction systems. The proposed framework considers transaction-level attributes, historical behavioral indicators, supervised classification, and cost-sensitive learning. Logistic regression, random forests, gradient boosting, and a cost-sensitive boosting configuration are selected as comparative predictive approaches. A fictional dataset representing digital financial transactions is introduced to demonstrate the intended analytical procedure.

The study emphasizes data preparation, minority-class representation, temporal validation, and the evaluation of classification errors. Two illustrative tables present a hypothetical transaction distribution and comparative model performance. Precision, recall, and F1-score are considered primary measures because fraudulent transactions represent a relatively small proportion of total transaction activity. The numerical examples are manually constructed and do not represent completed experiments, actual banking records, or verified detection outcomes. The methodological discussion identifies class imbalance, evolving transaction characteristics, inconsistent fraud labels, computational requirements, and data privacy as important considerations. The paper concludes that machine learning offers a relevant framework for supporting financial fraud identification, although effective implementation requires appropriate data governance, independently validated models, and continuous performance assessment. The proposed study provides a methodological foundation for future empirical research rather than evidence of an operational fraud detection system.

Keywords: Machine Learning, Financial Fraud Detection, Digital Transactions, Predictive Classification, Transaction Monitoring, Fraud Patterns, Class Imbalance, Cost-Sensitive Learning, Financial Data Analytics, Anomaly Detection.

1. Introduction

Digital transaction systems have become essential components of contemporary financial infrastructure. Electronic banking platforms, payment applications, online commercial services, and digital fund transfer systems enable individuals and organizations to conduct financial activities through interconnected technological environments. These systems process substantial volumes of transaction information involving monetary values, payment channels, transaction timing, and customer account activities. Although digital financial services improve accessibility and transaction efficiency, they also create significant requirements for transaction monitoring and financial security. Identifying potentially fraudulent activities within extensive transaction records has consequently become an important concern for financial institutions and computational researchers.

Financial fraud detection involves distinguishing potentially unauthorized or deceptive transactions from legitimate financial activities. This classification problem is complicated because fraudulent transactions may represent only a small proportion of the total records processed by a financial system. Furthermore, the characteristics of suspicious activities may overlap with legitimate customer behavior. A transaction that differs from historical spending patterns, for example, may reflect an unusual but authorized purchase rather than fraud. Consequently, effective monitoring requires careful evaluation of multiple transaction characteristics instead of relying exclusively on individual indicators or rigid decision rules.

Traditional transaction monitoring procedures frequently employ predefined conditions to identify activities requiring further examination. Such procedures can be useful when suspicious transaction characteristics are clearly established. However, fixed rules may be less effective when customer behavior changes or when relevant statistical relationships become increasingly complex. Machine learning offers an alternative analytical approach by identifying relationships among transaction features and previously established classification outcomes. Supervised algorithms can estimate fraud probabilities, while anomaly detection methods may identify unusual observations for additional review. These approaches can support structured transaction analysis without eliminating the need for appropriate human oversight.

The development of machine learning-based fraud identification systems also introduces significant methodological considerations. Financial transaction datasets often exhibit severe class imbalance, delayed fraud confirmation, incomplete observations, and changing distributions. These characteristics can influence model training and produce misleading performance estimates when conventional evaluation procedures are applied without adjustment. A model may achieve high overall accuracy by correctly classifying most legitimate transactions while identifying relatively few fraudulent activities. Consequently, precision, recall, F1-score, and the operational consequences of classification errors are particularly important for evaluating fraud detection performance.

The present paper examines machine learning-based identification of financial fraud patterns through a hypothetical methodological investigation. The proposed framework incorporates transaction feature preparation, comparative predictive modeling, cost-sensitive classification, and independent performance assessment. Four machine learning configurations are considered within a fictional digital payment environment to demonstrate how fraud identification research could be organized. The investigation emphasizes the relationship between classification quality, transaction characteristics, and operational reliability. No actual financial records have been collected or analyzed, and the numerical examples are included exclusively for academic illustration.

2. Objectives of the Study

The primary objective of this study is to examine the methodological application of machine learning techniques for identifying potential financial fraud patterns in digital transaction systems. The investigation emphasizes transaction characteristics, predictive classification, and evaluation reliability. Three specific objectives establish the conceptual direction of the proposed comparative framework.

2.1 Identification of Transaction Characteristics Associated With Fraud

The first objective is to examine how measurable transaction characteristics may contribute to the identification of potentially fraudulent financial activities. Transaction amount, payment timing, historical account activity, transaction frequency, and changes in customer behavior represent possible analytical variables. The proposed study considers how these attributes could be organized into structured predictive features. Particular attention is given to distinguishing unusual financial activity from transactions that are consistent with established customer behavior. This objective emphasizes systematic feature representation rather than assumptions based on isolated transaction characteristics.

The investigation also aims to assess the methodological usefulness of incorporating historical behavioral information into transaction classification. Individual transactions may provide limited evidence when considered independently, whereas aggregated account activity can reveal broader statistical relationships. However, historical information must be restricted to records available before the transaction being evaluated. This objective therefore considers temporal feature construction, data consistency, and appropriate observation periods. Such procedures may support more meaningful predictive evaluation while reducing the possibility of information leakage during model development.

2.2 Comparative Evaluation of Machine Learning Classifiers

The second objective focuses on establishing a comparative framework for machine learning approaches used in financial fraud identification. Logistic regression, random forests, gradient boosting, and cost-sensitive boosting are selected to represent different classification strategies. Their intended evaluation considers precision, recall, F1-score, and the relationship between missed fraudulent transactions and unnecessary alerts. The proposed comparison seeks to establish a consistent methodological procedure through which the relative suitability of these algorithms could be examined using equivalent transaction information.

A further objective involves investigating the importance of classification thresholds and operational error costs. Fraud detection systems must account for situations in which a legitimate transaction is incorrectly flagged and circumstances in which fraudulent activity is not identified. These outcomes may have different practical consequences. Therefore, model comparison should not depend exclusively on overall classification accuracy. The proposed framework emphasizes cost-aware evaluation, independent testing, and carefully documented decision criteria to support a more balanced interpretation of predictive performance.

2.3 Examination of Data Reliability and Detection Constraints

The third objective is to identify methodological challenges that may influence the reliability of machine learning-based fraud classification. Incomplete transaction information, inaccurate labels, minority-class representation, and distributional changes can affect the relationships learned by predictive algorithms. These conditions may also complicate comparisons among alternative models. The proposed investigation examines how data preparation, validation design, and suitable evaluation metrics could

address selected difficulties. Particular attention is given to recognizing limitations that may remain even after computational optimization.

The study further aims to examine the operational considerations associated with implementing predictive fraud monitoring systems. Financial transactions frequently require timely assessment, while model maintenance must accommodate changing customer behavior and newly confirmed fraud cases. Data privacy and appropriate access controls also represent important requirements. Consequently, this objective emphasizes transparent model development, responsible information processing, and ongoing evaluation. These considerations provide a foundation for understanding predictive classification as one component of a broader financial risk management process.

3. Review of Literature

3.1 Statistical Foundations of Financial Fraud Identification

Bolton and Hand (2002) examined statistical approaches to fraud detection and discussed the importance of identifying unusual activities within financial and related transaction environments. Their review established a methodological foundation for evaluating suspicious patterns through statistical analysis. Juszczak et al. (2008) subsequently investigated classification approaches for detecting fraudulent payment card transactions, including methods that distinguish one-class modeling from conventional supervised classification. These contributions demonstrated the complexity of identifying rare fraudulent observations within large transaction datasets containing substantial behavioral variation.

Whitrow et al. (2009) investigated transaction aggregation as a strategy for improving the representation of payment card activity. Their work emphasized the potential value of incorporating historical transactional information into fraud classification. Ngai et al. (2011) reviewed data mining applications across different categories of financial fraud and developed a classification framework for the existing literature. Together, these studies indicate that financial fraud identification depends on appropriate transaction representation, reliable classification methods, and an understanding of the practical differences among financial fraud problems.

3.2 Machine Learning Models and Cost-Sensitive Classification

Breiman (2001) introduced random forests as an ensemble learning approach capable of combining multiple randomized decision trees. Friedman (2001) developed gradient boosting as a method of sequential predictive optimization. Bhattacharyya et al. (2011) examined logistic regression, support vector machines, and random forests within the context of credit card fraud detection. These contributions established useful methodological foundations for comparing statistical and ensemble-based models. Their relevance lies in demonstrating how alternative learning procedures can be assessed under a common classification objective.

Elkan (2001) examined theoretical foundations of cost-sensitive learning, while Bahnsen et al. (2013) considered transaction-related financial consequences within credit card fraud classification. Dal Pozzolo et al. (2014) discussed practical difficulties involving unbalanced transaction streams and changing financial activity patterns. West and Bhattacharya (2016) provided a broader review of intelligent financial fraud detection methods. These studies indicate that predictive effectiveness should be evaluated through classification reliability and the operational consequences of errors rather than relying exclusively on conventional accuracy measurements.

3.3 Minority-Class Learning and Predictive Evaluation

Chawla et al. (2002) introduced the synthetic minority over-sampling technique as a method for addressing unequal class distributions in supervised learning. Fawcett (2006) examined receiver operating characteristic analysis, while Davis and Goadrich (2006) investigated relationships between precision-recall and ROC evaluation. He and Garcia (2009) reviewed methodological difficulties associated with learning from imbalanced datasets. These contributions established important principles for assessing classification performance when the outcomes of greatest interest occur relatively infrequently within the available observations.

Niculescu-Mizil and Caruana (2005) examined the reliability of predicted class probabilities, emphasizing the importance of calibration. Gama et al. (2014) reviewed approaches to concept drift adaptation, while Saito and Rehmsmeier (2015) demonstrated the interpretive relevance of precision-recall analysis for imbalanced classification problems. These studies provide complementary foundations for fraud detection evaluation. They suggest that suitable performance assessment requires attention to class frequencies, decision thresholds, probability quality, and the stability of predictive relationships across changing data environments.

4. Research Methodology

4.1 Research Design and Analytical Framework

The present paper proposes a hypothetical quantitative methodology for examining machine learning-based financial fraud identification. The intended research design involves supervised binary classification, with transactions categorized as legitimate or fraudulent according to established reference labels. Four predictive configurations are considered: logistic regression, random forests, gradient boosting, and cost-sensitive gradient boosting. The framework is designed to demonstrate how different classification approaches could be compared through consistent preprocessing and evaluation procedures. No actual financial transaction dataset has been processed.

The proposed analytical procedure consists of transaction feature preparation, data quality assessment, temporal partitioning, model training, parameter optimization, and independent performance testing. Each classifier would be evaluated using the same transaction records while receiving appropriately selected model configurations. The cost-sensitive approach would incorporate different consequences for classification errors during optimization or decision-making. Model selection would rely on validation data, and independent test observations would remain unavailable until final evaluation. These safeguards are intended to support fair methodological comparison.

4.2 Hypothetical Digital Transaction Dataset

For illustrative purposes, the proposed study assumes a fictional dataset containing 48,000 electronic payment transactions recorded across 180 consecutive days. The hypothetical records include transaction value, processing time, account activity indicators, payment category, and selected historical behavioral attributes. The example contains 360 transactions labeled as fraudulent and 47,640 labeled as legitimate. These figures were manually constructed to demonstrate an imbalanced classification problem. No banking information, payment records, customer identities, or operational financial systems were accessed.

The proposed data preparation procedure would involve verification of transaction labels, removal of duplicate records, examination of missing values, and consistent encoding of categorical

attributes. Numerical features would be processed according to model requirements, while historical variables would be calculated using information available before each transaction. A chronological division would separate development and evaluation periods. Any resampling or class-balancing procedure would be confined to training data, thereby protecting the independence of validation and test observations.

4.3 Model Training and Fraud Detection Evaluation

The proposed classification stage would use logistic regression as an interpretable statistical baseline and random forests to represent nonlinear relationships through ensemble decision trees. Gradient boosting would provide an additional sequential learning approach, while cost-sensitive boosting would consider the unequal consequences of classification errors. Model parameters and classification thresholds would be selected using validation information. Each approach would produce transaction-level predictions or risk scores suitable for structured comparison. No classifier training was performed for the present illustrative study.

The intended evaluation would prioritize precision, recall, and F1-score for the fraudulent class. Precision would represent the proportion of flagged transactions that are genuinely fraudulent, while recall would represent the proportion of fraudulent transactions identified. F1-score would summarize the balance between these measures. Precision-recall curves and operational cost analysis would provide additional information in an implemented investigation.

5. Analysis

5.1 Illustrative Distribution of Legitimate and Fraudulent Transactions

The distribution of legitimate and fraudulent transactions influences both predictive model development and performance interpretation. Fraud identification is commonly examined as an imbalanced classification problem because suspicious or confirmed fraudulent activities may occur less frequently than legitimate financial transactions. Table 1 presents a hypothetical dataset divided into training, validation, and testing partitions. The fictional example assigns 0.75% of observations to the fraudulent class, demonstrating the methodological implications of a relatively uncommon target outcome.

The illustrative distribution emphasizes why classification accuracy alone may provide misleading information. A hypothetical model predicting every transaction as legitimate would achieve 99.25% overall accuracy within this constructed dataset while identifying no fraudulent transactions. Consequently, classification performance must be examined using metrics sensitive to minority-class recognition.

Table 1: Hypothetical Distribution of Financial Transaction Records

Transaction Category	Training Set	Validation Set	Test Set	Total
Legitimate Transactions	33,348	7,146	7,146	47,640
Fraudulent Transactions	252	54	54	360
Total Transactions	33,600	7,200	7,200	48,000

The hypothetical distribution contains 70% training observations, 15% validation observations, and 15% independent test observations. In an actual study, transaction records would be separated chronologically to support realistic evaluation of future financial activity.

5.2 Illustrative Comparative Performance of Fraud Detection Models

Fraud detection models can be evaluated by comparing predicted transaction categories against confirmed reference labels. Precision indicates the reliability of positive fraud alerts, while recall measures the ability to identify fraudulent transactions within the evaluated dataset. Table 2 presents hypothetical performance values for four classification approaches. The cost-sensitive boosting configuration is assigned the highest illustrative recall and F1-score, demonstrating one possible analytical scenario. These values were manually constructed and do not represent executed machine learning experiments.

The hypothetical comparison demonstrates how different modeling strategies could be evaluated according to competing classification priorities. A model identifying more fraudulent transactions may also generate additional false alerts requiring review. Consequently, performance interpretation should consider the relative importance of missed fraud and unnecessary transaction investigations. A genuine experimental comparison would require independent testing, calibrated probabilities, appropriate decision thresholds, and assessment of relevant financial consequences. The fictional values presented here cannot establish the operational superiority of any classifier.

Table 2: Hypothetical Comparison of Machine Learning Fraud Classification Performance

Classification Model	Precision	Recall	F1-Score	False-Positive Cases
Logistic Regression	0.323	0.574	0.413	65
Random Forest	0.462	0.667	0.545	42
Gradient Boosting	0.557	0.722	0.629	31
Cost-Sensitive Gradient Boosting	0.537	0.815	0.647	38

6. Discussion

6.1 Transaction Behavior and Predictive Fraud Identification

Transaction behavior provides a potentially informative basis for examining unusual financial activities through machine learning. Historical spending patterns, transaction frequency, payment values, and account activity may contribute to statistical representations of legitimate financial behavior. However, deviations from historical patterns do not necessarily establish fraud because customers frequently make authorized transactions that differ from previous activity. The proposed framework therefore treats predictive outputs as indicators requiring suitable evaluation rather than definitive evidence of fraudulent conduct.

The methodological value of machine learning lies in its capacity to examine relationships among multiple transaction attributes. Ensemble models may identify nonlinear interactions that are difficult to represent through simpler statistical procedures. Nevertheless, predictive complexity does not automatically guarantee superior classification. Reliable fraud identification depends on representative

training information, meaningful reference labels, and appropriate validation. The hypothetical framework emphasizes these conditions while recognizing that operational financial security requires broader investigative and administrative procedures.

6.2 Class Imbalance and Cost-Sensitive Decision-Making

Class imbalance represents a significant consideration because legitimate transactions may substantially outnumber fraudulent activities. Under these conditions, a classifier can produce favorable aggregate accuracy while failing to identify the transactions of greatest analytical interest. The proposed evaluation therefore emphasizes precision, recall, and F1-score rather than relying solely on overall correctness. These measures provide complementary information about the reliability of fraud alerts and the proportion of fraudulent transactions successfully identified.

Cost-sensitive learning introduces an additional perspective by considering differences in the consequences of classification errors. Missing a fraudulent transaction may create financial exposure, while incorrectly flagging legitimate activity may produce unnecessary reviews and customer inconvenience. The hypothetical performance comparison illustrates how classification priorities can influence the balance between precision and recall. However, determining an appropriate operating threshold requires reliable cost information, relevant business constraints, and empirical performance evidence rather than assumed numerical improvements.

6.3 Interpretation of Comparative Classification Outcomes

The illustrative results assign different precision and recall values to the selected predictive models. The cost-sensitive configuration demonstrates higher fictional fraud recall than the conventional classifiers, although its precision is lower than that of the gradient boosting baseline. This example highlights the importance of interpreting classification measures together. Since the numerical values are hypothetical, the observed differences cannot establish statistical significance, algorithmic superiority, or actual improvements in financial fraud detection performance.

Previous methodological literature supports the need for cautious comparative interpretation. Davis and Goadrich (2006) examined precision-recall analysis, while Saito and Rehmsmeier (2015) emphasized its relevance for imbalanced datasets. Bhattacharyya et al. (2011) and Dal Pozzolo et al. (2014) investigated practical aspects of financial fraud classification. These contributions indicate that meaningful evaluation should reflect transaction distributions, operational requirements, and independent testing conditions. The present framework incorporates these considerations without claiming experimental validation.

7. Challenges

7.1 Severe Class Imbalance and Incomplete Fraud Labels

Financial transaction datasets may contain substantially fewer confirmed fraudulent records than legitimate observations. This imbalance complicates supervised learning because predictive algorithms can become disproportionately influenced by the majority class. Standard accuracy measurements may therefore conceal poor fraud recognition. Effective evaluation requires suitable class-sensitive metrics and carefully controlled learning procedures. Resampling or class weighting may be considered during training, but these methods require appropriate validation to determine whether they actually improve classification reliability.

Incomplete fraud labels introduce another significant difficulty because fraudulent activities may be identified only after investigations or customer reports. Transactions initially recorded as legitimate may subsequently receive different classifications. Such delays can affect model training and produce uncertainty regarding the reliability of evaluation labels. A future investigation should therefore document the timing and source of fraud confirmation. The proposed methodology recognizes that accurate reference labels are essential for meaningful fraud detection assessment and responsible interpretation of predictive results.

7.2 Changing Transaction Patterns and Model Stability

Digital financial behavior changes as customers adopt different payment methods, transaction channels, and purchasing practices. These developments can alter the statistical characteristics of legitimate transaction activity and influence predictive model performance. A classifier trained on earlier observations may become less reliable when applied to substantially different financial conditions. Consequently, operational fraud detection requires periodic evaluation of predictive relationships. Changes should be examined systematically rather than assuming that every temporary variation justifies immediate model retraining.

Maintaining model stability also requires appropriate management of confirmed new observations and changing class distributions. Frequent modifications may introduce inconsistency when updates are based on limited or unrepresentative data. Conversely, retaining outdated models can reduce their relevance to current transaction environments. The proposed research framework therefore emphasizes predefined monitoring procedures, documented model updates, and independent performance assessment. These safeguards may help distinguish genuine model deterioration from fluctuations arising through normal transaction variability.

7.3 Data Privacy, Interpretability, and Operational Reliability

Financial transaction information may include sensitive account activity, purchasing behavior, and other confidential customer attributes. Machine learning research involving such information requires appropriate access controls, data minimization, and responsible processing procedures. Unauthorized disclosure or unnecessary retention can create substantial privacy concerns. Therefore, a genuine fraud detection investigation should employ suitable governance arrangements and restrict the use of personally identifiable information. Predictive effectiveness cannot justify inadequate protection of confidential financial records.

Interpretability presents an additional challenge because financial institutions may need to explain why particular transactions were selected for investigation. Complex classification models can produce risk estimates through relationships that are difficult to communicate clearly. In operational environments, model performance must also be balanced against processing latency and unnecessary alerts. These considerations support the development of systems that combine predictive analysis with suitable review procedures. Automated predictions should inform responsible risk assessment rather than independently establish wrongdoing.

8. Conclusion

Machine learning provides a relevant methodological foundation for identifying potentially fraudulent patterns within digital transaction systems. The present paper developed a hypothetical comparative

framework incorporating logistic regression, random forests, gradient boosting, and cost-sensitive classification. Its principal focus was the organization of transaction features, imbalanced classification, and reliable performance evaluation. The proposed approach demonstrates how predictive algorithms could be examined under consistent analytical conditions without assuming that computational complexity necessarily produces superior fraud identification.

The illustrative dataset and comparative performance tables demonstrate possible procedures for reporting transaction distributions, precision, recall, F1-score, and classification errors. However, all numerical values were constructed exclusively for academic demonstration and do not represent actual financial observations or completed machine learning experiments. Establishing meaningful detection performance would require authentic transaction data, dependable fraud labels, reproducible model development, and independent testing. Operational conclusions would also require consideration of financial costs and the consequences of unnecessary alerts.

Overall, the study identifies class imbalance, delayed fraud confirmation, changing transaction behavior, data privacy, and model interpretability as important considerations in financial fraud analytics. Future empirical research should investigate these factors through representative datasets and carefully controlled evaluation procedures. Machine learning may support financial security and transaction monitoring when implemented within appropriate governance and investigation processes. Nevertheless, genuine claims concerning model effectiveness must be based on independently verified experimental evidence rather than hypothetical numerical comparisons.

References

1. Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90–113. <https://doi.org/10.1016/j.jnca.2016.04.007>
2. Adewumi, A. O., & Akinyelu, A. A. (2017). A survey of machine-learning and nature-inspired based credit card fraud detection techniques. *International Journal of System Assurance Engineering and Management*, 8(Suppl. 2), 937–953. <https://doi.org/10.1007/s13198-016-0551-y>
3. 2019. *Computer Science Review*, 40, Article [1] 100402. <https://doi.org/10.1016/j.cosrev.2021.100402>
4. Bahnsen, A. C., Aouada, D., Stojanovic, A., & Ottersten, B. (2016). Feature engineering strategies for credit card fraud detection. *Expert Systems with Applications*, 51, 134–142. <https://doi.org/10.1016/j.eswa.2015.12.030>
5. Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3), 602–613. <https://doi.org/10.1016/j.dss.2010.08.008>
6. Bolton, R. J., & Hand, D. J. (2002). Statistical fraud detection: A review. *Statistical Science*, 17(3), 235–255. <https://doi.org/10.1214/ss/1042727940>
7. Carcillo, F., Dal Pozzolo, A., Le Borgne, Y.-A., Caelen, O., Mazzer, Y., & Bontempi, G. (2018). SCARFF: A scalable framework for streaming credit card fraud detection with Spark. *Information Fusion*, 41, 182–194. <https://doi.org/10.1016/j.inffus.2017.09.005>
8. Chen, C.-T., Lee, C., Huang, S.-H., & Peng, W.-C. (2024). Credit card fraud detection via intelligent sampling and self-supervised learning. *ACM Transactions on Intelligent Systems and Technology*, 15(2), Article 35. <https://doi.org/10.1145/3641283>

9. Dal Pozzolo, A., Boracchi, G., Caelen, O., Alippi, C., & Bontempi, G. (2018). Credit card fraud detection: A realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8), 3784–3797. <https://doi.org/10.1109/TNNLS.2017.2736643>
10. Dastidar, K. G., Caelen, O., & Granitzer, M. (2024). Machine learning methods for credit card fraud detection: A survey. *IEEE Access*, 12, 158939–158965. <https://doi.org/10.1109/ACCESS.2024.3487298>
11. Hernandez Aros, L., Bustamante Molano, L. X., Gutierrez-Portela, F., Moreno Hernandez, J. J., & Rodríguez Barrero, M. S. (2024). Financial fraud detection through the application of machine learning techniques: A literature review. *Humanities and Social Sciences Communications*, 11, Article 1130. <https://doi.org/10.1057/s41599-024-03606-0>
12. Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P.-E., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. *Expert Systems with Applications*, 100, 234–245. <https://doi.org/10.1016/j.eswa.2018.01.037>
13. Kou, Y., Lu, C.-T., Sirwongwattana, S., & Huang, Y.-P. (2004). Survey of fraud detection techniques. In *Proceedings of the 2004 IEEE International Conference on Networking, Sensing and Control* (Vol. 2, pp. 749–754). IEEE. <https://doi.org/10.1109/ICNSC.2004.1297040>
14. Mienye, I. D., & Jere, N. (2024). Deep learning for credit card fraud detection: A review of algorithms, challenges, and solutions. *IEEE Access*, 12, 96893–96910. <https://doi.org/10.1109/ACCESS.2024.3426955>
15. Ngai, E. W. T., Hu, Y., Wong, Y. H., Chen, Y., & Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3), 559–569. <https://doi.org/10.1016/j.dss.2010.08.006>
16. Randhawa, K., Loo, C. K., Seera, M., Lim, C. P., & Nandi, A. K. (2018). Credit card fraud detection using AdaBoost and majority voting. *IEEE Access*, 6, 14277–14284. <https://doi.org/10.1109/ACCESS.2018.2806420>
17. Shahana, T., Lavanya, V., & Bhat, A. R. (2023). State of the art in financial statement fraud detection: A systematic review. *Technological Forecasting and Social Change*, 192, Article 122527. <https://doi.org/10.1016/j.techfore.2023.122527>
18. Tong, G., & Shen, J. (2023). Financial transaction fraud detector based on imbalance learning and graph neural network. *Applied Soft Computing*, 149, Article 110984. <https://doi.org/10.1016/j.asoc.2023.110984>
19. Van Vlasselaer, W., Bravo, C., Caelen, O., Eliassi-Rad, T., Akoglu, L., Snoeck, M., & Baesens, B. (2015). APATE: A novel approach for automated credit card transaction fraud detection using network-based extensions. *Decision Support Systems*, 75, 38–48. <https://doi.org/10.1016/j.dss.2015.04.013>
20. West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers & Security*, 57, 47–66. <https://doi.org/10.1016/j.cose.2015.09.005>