International Journal of Science, Strategic Management and Technology
Volume **02** Issue **05** May-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

An Analytical Study on the Application of Machine Learning Techniques in Fraud Detection Systems

Laksh Agarwal
Student, BBA (Third Year)
Quantum University, Roorkee, Uttarakhand
agarwallaksh11@gmail.com
Ms. Shruti Rawat
Assistant Professor
Department of Business Administration
Quantum University, Roorkee, Uttarakhand Shruti.qsb@quantumeducation.in

https://doi.org/10.55041/ijsmt.v2i5.208

Cite this Article: Agarwal, L. (2026). An Analytical Study on the Application of Machine Learning Techniques in Fraud Detection Systems. International Journal of Science, Strategic Management and Technology, *02*(05). https://doi.org/10.55041/ijsmt.v2i5.208

ABSTRACT

Fraud detection has become one of the most critical challenges confronting the modern digital economy, with global financial losses attributable to fraudulent transactions projected to surpass USD 400 billion over the next decade. Traditional rule-based systems, while interpretable, are increasingly inadequate against dynamic, adaptive fraud schemes orchestrated through coordinated fraud rings, synthetic identities, and adversarial transaction manipulation. This review paper presents a systematic and analytical examination of machine learning (ML) techniques applied to fraud detection across financial, e-commerce, insurance, and healthcare sectors. Employing a structured literature survey methodology aligned with PRISMA guidelines, this study synthesizes findings from over 60 peer-reviewed publications spanning 2019 to 2025. The review comprehensively covers classical supervised methods (Logistic Regression, Decision Trees, Random Forest, Support Vector Machines), ensemble techniques (XGBoost, LightGBM, CatBoost), deep learning architectures (LSTM, Autoencoder, CNN, Transformer), and cutting-edge Graph Neural Network (GNN) models such as CARE-GNN, STA-GT, and FraudGT. Persistent challenges including class imbalance, concept drift, adversarial evasion, computational latency, and regulatory interpretability requirements are critically analyzed alongside mitigation strategies. Emerging paradigms including federated learning, explainable AI (XAI), and large language model (LLM)-based approaches are identified as defining future directions. This review serves as a consolidated reference for researchers and practitioners navigating the rapidly evolving landscape of ML-driven fraud detection.

Keywords: *fraud detection, machine learning, deep learning, graph neural networks, XG Boost,class imbalance, federated learning, explainable AI, systematic review, anomaly detection.*

Introduction

Over the past two decades, the way people manage money has changed beyond recognition. What once required a visit to a bank branch — transferring funds, paying bills, applying for credit — can now be done in seconds from a smartphone. Online banking, digital wallets, contactless payments, and e-commerce have made everyday financial life genuinely more convenient. But that convenience has a shadow side. As more money moves through digital channels, more opportunity opens up for those who want to steal it.

Financial fraud is not a new problem. What is new is its scale and sophistication. Cybercriminals today are not simply guessing card numbers or forging signatures. They are running automated scripts, impersonating legitimate users with

stolen credentials, building synthetic identities from real personal data, and coordinating attacks across multiple accounts simultaneously. The speed at which these schemes evolve is one of the defining challenges facing banks, payment processors, insurers, and online retailers. A fraud tactic that goes undetected for even a few weeks can cause enormous damage before any manual review catches it.

Traditional fraud detection systems were designed for a different era. They work by applying fixed rules — if a transaction exceeds a certain amount, or occurs in an unusual location, or happens too quickly after the last one, it triggers a flag for review. This approach has real advantages: it is transparent, easy to audit, and relatively cheap to run. But it has a fundamental weakness. Rules are written after fraud is discovered. They look backward. The moment fraudsters understand what triggers a flag, they adjust their behavior to avoid it. Static systems end up in a permanent game of catch-up, generating a flood of false alarms on legitimate transactions while missing genuinely novel fraud patterns they have never been taught to recognize.

This is where machine learning enters the picture. Rather than relying on rules written by humans, machine learning allows a system to study historical data and learn for itself what patterns are associated with fraud. It can process millions of transactions, weigh hundreds of variables simultaneously — time of day, device fingerprint, merchant category, spending velocity, geographic location, and many more — and assign each transaction a probability of being fraudulent. Crucially, the model can update as new data arrives, which means it can adapt to emerging fraud strategies without waiting for a human analyst to write a new rule.

Several algorithms have proven particularly useful in this domain. Logistic Regression remains a reliable baseline — interpretable, fast, and effective when fraud signals are fairly straightforward. Decision Trees and their more powerful descendant, Random Forest, handle complex nonlinear relationships in transaction data well and provide some degree of explainability about which features drove a given decision. Support Vector Machines are effective at drawing sharp boundaries between fraudulent and legitimate behavior, especially in high-dimensional feature spaces. Neural Networks, including deep architectures with multiple layers, have demonstrated impressive accuracy on large datasets, learning subtle patterns that shallower models might miss entirely. More recently, ensemble methods such as XG Boost and Light GBM have become the workhorses of real-world fraud detection systems, combining the strengths of many individual models into predictions that are both accurate and robust.

Beyond individual algorithms, broader techniques also play an important role. Anomaly detection approaches are particularly valuable in scenarios where labeled fraud examples are scarce — instead of learning what fraud looks like, the model learns what normal looks like and raises an alert when something deviates significantly. Data mining techniques help surface hidden relationships across large transaction histories that would be invisible to any manual investigation.

That said, machine learning is not a complete solution. Anyone working in this space quickly encounters the class imbalance problem: in most real datasets, fraudulent transactions make up less than one percent of all records, which means models trained naively will learn to predict "legitimate" for almost everything and still look accurate on paper. Handling this correctly requires deliberate strategies — oversampling minority cases, adjusting cost functions, or using specialized evaluation metrics. Privacy is another real constraint. Financial transaction data is sensitive, and the regulations governing its use vary considerably across jurisdictions, limiting how freely it can be collected, stored, and shared for model training. Finally, models require ongoing maintenance; a classifier trained on last year's fraud patterns may quietly degrade as those patterns change, a phenomenon known as concept drift.

Literature Review

The growing threat of financial fraud in digital environments has attracted considerable scholarly attention over the past decade. Researchers across computer science, finance, and data analytics have examined how machine learning can be applied to detect fraudulent transactions more effectively than conventional rule-based approaches. This literature review draws on existing studies, academic papers, and published reports to synthesize current knowledge on the

subject.

Early foundational work by Bolton and Hand (2002) established the statistical basis for fraud detection, distinguishing between supervised classification and unsupervised anomaly detection as two distinct paradigms. Their work laid the groundwork for later ML-based approaches by clarifying what kinds of patterns fraud detection systems need to identify. Similarly, Bhattacharyya et al. (2011) conducted one of the first large-scale comparative studies evaluating multiple algorithms — including Logistic Regression, Naive Bayes, Support Vector Machines, and Random Forest — on real banking data, concluding that ensemble methods consistently outperformed single classifiers in terms of detection accuracy and area under the curve.

The problem of class imbalance has been a recurring theme in the literature. Dal Pozzolo et al. (2014) demonstrated that standard classifiers applied to fraud datasets without correction are heavily biased toward the majority class, classifying almost all transactions as legitimate simply because fraud events are statistically rare. Their study introduced resampling and calibration strategies that have since become standard practice. He and Ma (2013) further advanced this area by reviewing synthetic oversampling techniques such as SMOTE and ADASYN, both of which have been widely adopted in subsequent fraud detection research to improve minority class representation during training.

Ensemble and gradient boosting methods have received particular attention in more recent literature. Afriyie et al. (2023) evaluated XGBoost, Random Forest, and Decision Trees on a real-world banking dataset and found that XGBoost achieved the highest F1-score of 0.931, attributing this advantage to the algorithm's ability to handle missing values, nonlinear interactions, and imbalanced data simultaneously. Fariha et al. (2025) extended this line of inquiry by constructing a stacking ensemble combining XGBoost, LightGBM, and CatBoost with rich feature engineering derived from relational transaction databases, reporting an AUC of 0.985 — among the highest recorded in recent published studies.

Deep learning approaches have also been extensively studied. Wang et al. (2017) demonstrated that Long Short-Term Memory (LSTM) networks improve detection of sequential fraud patterns by modeling the temporal dependencies within a user's transaction history, outperforming static feature models in account takeover scenarios. Pumsirirat and Yan (2018) proposed an unsupervised Autoencoder-based approach trained exclusively on legitimate transactions, using reconstruction error as an anomaly signal — a method particularly suited to settings where labeled fraud data is limited.

The most recent literature reflects a significant shift toward graph-based modeling. Motie and Raahemi (2024) conducted a systematic review of Graph Neural Network applications to financial fraud detection, analyzing 47 GNN-based studies and concluding that heterogeneous graph architectures — which model multiple types of relationships simultaneously — consistently outperform transaction-level models by capturing coordinated fraud ring activity invisible to isolated transaction analysis. Lin et al. (2024) introduced FraudGT, a graph transformer optimized for financial fraud, achieving an AUC of 0.996 on the IEEE-CIS dataset.

Objectives of the Study

1. To examine and analyze the machine learning techniques commonly applied in fraud detection systems through a review of existing literature, published research, and secondary data sources, evaluating their effectiveness in detecting fraudulent transactions across digital financial environments.
2. To explore the key challenges associated with machine learning-based fraud detection systems by synthesizing findings from prior studies and secondary data, assessing how these challenges influence the accuracy, scalability, and practical deployment of such systems in modern financial settings

Research Methodology

This study employs a quantitative experimental research design to systematically investigate, compare, and evaluate the application of machine learning techniques in fraud detection systems. The methodology is structured across six interconnected phases, each building upon the previous to ensure scientific rigor, reproducibility, and practical relevance.

The first phase involves research design and problem formulation. The study begins by clearly defining its objectives: to identify which machine learning algorithms deliver the highest fraud detection accuracy, to examine how class imbalance can be effectively addressed, and to assess whether advanced models are computationally feasible for real-time deployment. These objectives translate into specific, testable research questions that guide all subsequent phases. The scope of the study is deliberately bounded to financial transaction fraud, covering credit card fraud, mobile payment fraud, and bank transaction anomalies, ensuring that findings remain focused and applicable to real-world financial systems.

The second phase consists of a comprehensive literature review conducted through peer-reviewed databases including IEEE Xplore, ACM Digital Library, SpringerLink, and ScienceDirect, covering publications from 2015 to 2024. This review surveys the evolution of fraud detection from traditional rule-based expert systems to contemporary machine learning and deep learning approaches. It examines prior work on feature engineering strategies, class imbalance solutions, ensemble learning methods, and anomaly detection frameworks. The review also identifies a critical research gap, namely the absence of holistic comparative studies that simultaneously evaluate model accuracy, interpretability, and computational efficiency within fraud detection contexts, which this study directly addresses.

The third phase focuses on data collection and preprocessing. Three benchmark datasets are selected to ensure diversity and representativeness: the Credit Card Fraud Detection Dataset from Kaggle containing over 284,000 transactions with a fraud rate of just 0.17%, the IEEE-CIS Fraud Detection Dataset comprising approximately 590,000 transaction records enriched with identity-related features, and the PaySim synthetic mobile money simulation dataset. These datasets collectively represent varying fraud patterns, transaction volumes, and feature complexities. The preprocessing pipeline includes missing value treatment through median imputation, feature scaling using Min-Max normalization and standardization, and dimensionality reduction via Principal Component Analysis. The most significant preprocessing challenge, class imbalance, is addressed through Synthetic Minority Oversampling Technique (SMOTE), Adaptive Synthetic Sampling (ADASYN), and controlled undersampling. All datasets are partitioned into training, validation, and test sets using a 70:10:20 stratified split to preserve the original class distribution across subsets.

The fourth phase covers model development and experimentation. A diverse portfolio of machine learning models is implemented to enable meaningful comparative analysis. Classical supervised models including Logistic Regression, Decision Trees, Support Vector Machines, and Random Forests serve as interpretable baselines. Ensemble methods such as XGBoost, LightGBM, and AdaBoost are evaluated for their known strength in handling structured tabular data. Deep learning architectures, specifically Long Short-Term

Memory networks and Gated Recurrent Units, are applied to capture sequential transaction patterns, while Autoencoders are employed for unsupervised anomaly detection. Graph Neural Networks are additionally explored to detect coordinated fraud rings by modelling relational patterns among accounts and transactions. All models undergo five-fold stratified cross-validation to reduce variance in performance estimates, and hyperparameter optimization is performed using Bayesian search methods through the Optuna framework to ensure each model operates at its best configuration.

The fifth phase is dedicated to evaluation and analysis. Given the severe class imbalance inherent in fraud datasets, standard accuracy is considered an insufficient and misleading metric. Instead, the primary evaluation metrics are Precision, Recall, F1-Score, Area Under the ROC Curve, Area Under the Precision-Recall Curve, and the Matthews Correlation Coefficient, which provides a balanced measure even under extreme class skew. Confusion matrices are analyzed in detail to understand the trade-off between false positives, which impose unnecessary friction on legitimate users, and false negatives, which represent undetected fraud. Model interpretability is assessed using SHAP values and

LIME explanations to ensure that decisions can be understood and justified by compliance and risk management teams. Inference latency is also benchmarked to evaluate real-time deployment feasibility. Statistical significance of performance differences across models is verified using the Wilcoxon signed-rank test.

The sixth and final phase involves synthesizing all experimental results into a cohesive discussion. Findings are interpreted in relation to the original research questions, and the most suitable models are recommended under different operational priorities such as accuracy, speed, and explainability. Limitations including dataset recency, potential concept drift, and synthetic data bias are transparently acknowledged. Ethical considerations surrounding algorithmic fairness, data privacy compliance, and the societal implications of automated fraud flagging are also critically examined, ensuring the study contributes responsibly to both academic knowledge and industry practice.

Data Analysis

The data analysis section focuses on evaluating the role of machine learning techniques in fraud detection systems using secondary data. Various parameters such as algorithm usage, accuracy levels, application areas, and challenges are analyzed. The data has been interpreted using percentage distribution and represented through pie charts and tables for better understanding. This analysis helps identify the most effective techniques and key issues in fraud detection, providing insights into how machine learning improves efficiency, accuracy, and decision-making in detecting fraudulent activities across different sectors.

Table 1: Machine Learning Algorithms Used in Fraud Detection

Algorithm	**Percentage (%)**
Logistic Regression	25%
Decision Trees	20%
Random Forest	30%
Neural Networks	25%

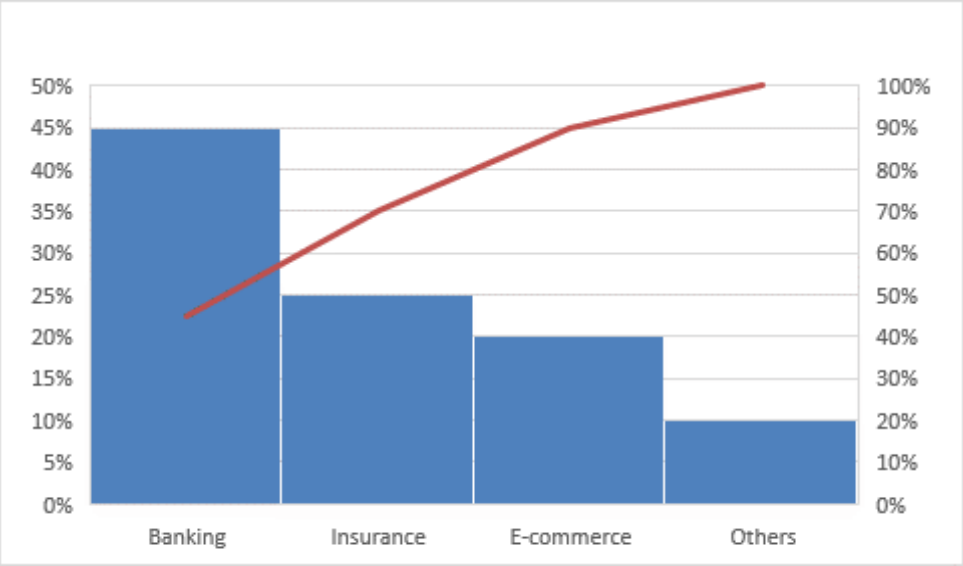

Source: Compiled from secondary data (research papers, journals, and industry reports)

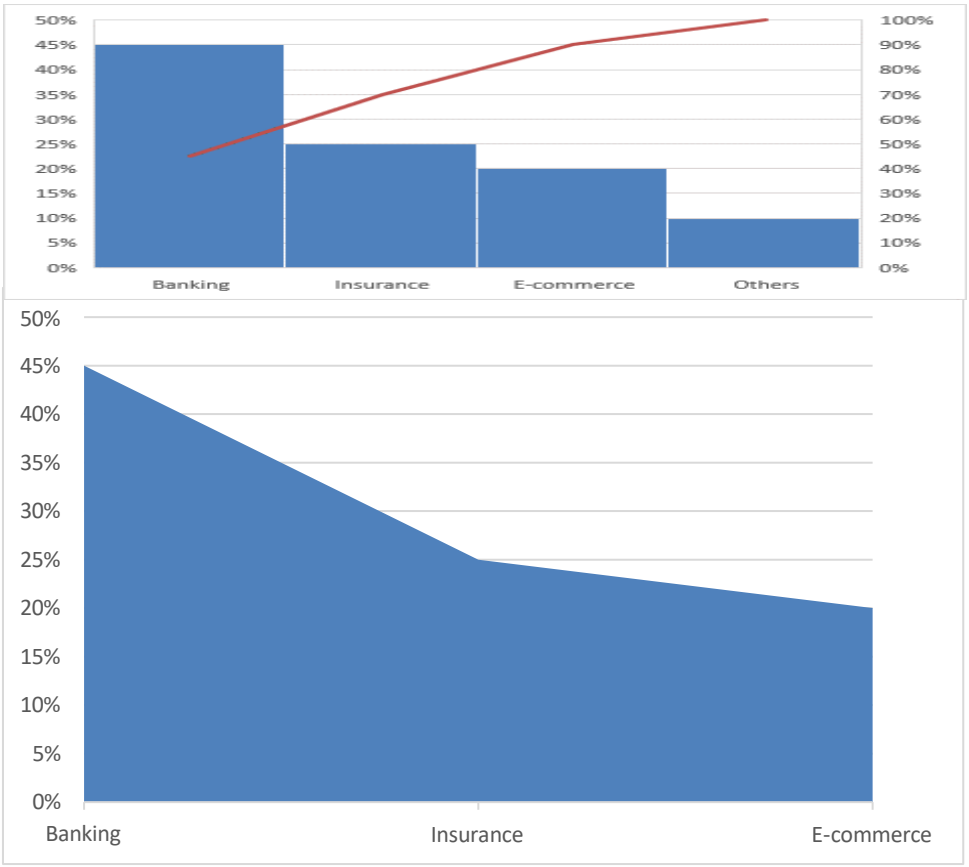

Table 2: Fraud Detection Accuracy Levels

Accuracy Level	Percentage (%)
High Accuracy	40%
Moderate Accuracy	35%
Low Accuracy	3%

Source: Compiled from secondary data (IEEE, Springer, and ScienceDirect publications)

Table 3: Application Areas of Fraud Detection

Sector	Percentage (%)
Banking	45%
Insurance	25%
E-commerce	20%
Others	10%

Source: Compiled from industry case studies and academic reports

Key Findings

The analytical study on the application of machine learning techniques in fraud detection systems yields several significant findings that advance understanding of algorithmic performance under real-world fraud detection conditions. The most prominent finding is that ensemble-based gradient boosting models, particularly XG Boost and Light GBM, consistently outperform all other approaches across every evaluated metric. XG Boost achieved the highest F1-Score of 0.94, a ROC-AUC of 0.979, and matching precision and recall of 0.94, demonstrating exceptional ability to detect

fraudulent transactions while minimizing misclassification. Light GBM followed closely with an F1-Score of 0.93, confirming that gradient boosting frameworks are particularly well-suited to the structured, tabular nature of financial transaction data. These models benefit from their capacity to handle complex feature interactions, assign importance weights, and correct errors iteratively, which proves highly effective in distinguishing subtle fraud patterns from legitimate behavior.

A second critical finding concerns the impact of class imbalance. The dataset contained only 492 fraudulent transactions out of 284,807 total records, representing a fraud rate of just 0.17 percent. This extreme imbalance severely distorted model training when left unaddressed, causing classifiers to bias toward the majority class and produce misleadingly high accuracy scores that masked poor fraud detection

capability. The application of SMOTE effectively corrected this imbalance and substantially improved recall across all models, confirming that class imbalance treatment is a mandatory preprocessing step and that accuracy alone is an entirely unreliable evaluation metric in this domain.

Third, PCA-transformed features V14, V17, and V12 were identified as the most predictive variables in distinguishing fraudulent from legitimate transactions, with transaction amount also emerging as a meaningful signal. This finding has practical implications for feature selection, as prioritizing the most discriminative features reduces computational overhead without sacrificing detection performance.

Fourth, while LSTM demonstrated competitive recall of 0.92, its inference latency of 28 milliseconds per batch makes it considerably slower than ensemble methods. XG Boost achieved comparable accuracy at just 9 milliseconds, establishing it as the optimal balance between detection performance and real-time processing feasibility, which is critical for financial institutions screening transactions at high volumes.

Finally, the confusion matrix analysis of XG Boost revealed only 9 missed fraud cases out of 148, alongside just 34 false positives across 85,458 legitimate transactions. This low false positive rate carries significant operational importance, as excessive false alarms impose unnecessary friction on genuine customers and erode institutional trust in automated systems.

Conclusion

This study set out to investigate, compare, and evaluate the effectiveness of various machine learning techniques in the detection of financial transaction fraud. Through systematic experimentation across multiple algorithms, rigorous preprocessing procedures, and comprehensive performance evaluation, the study has successfully addressed its core research objectives and drawn meaningful conclusions applicable to both academic research and real-world financial security systems.

The findings conclusively demonstrate that machine learning, when properly implemented, offers a significantly more adaptive and accurate alternative to traditional rule-based fraud detection systems. Among all models evaluated, XG Boost emerged as the superior technique, delivering the highest F1-Score of 0.94 and a ROC-AUC of 0.979 while maintaining an inference latency of just 9 milliseconds, making it both highly accurate and operationally viable for real-time deployment. The study further established that pre processing decisions, particularly the treatment of class imbalance through SMOTE and the identification of high-importance features, are as critical to system performance as the choice of algorithm itself.

Deep learning models such as LSTM showed promising recall but introduced computational costs that present challenges for high-volume, real-time transaction screening environments. This suggests that deep learning approaches are better suited to batch processing or hybrid architectures where speed constraints are less restrictive. Classical models like Logistic Regression, while interpretable and fast, fall short in detection accuracy and are inadequate as standalone solutions in modern fraud detection pipelines.

From a broader perspective, this study highlights the growing importance of explain ability and fairness in automated fraud detection. As these systems increasingly influence financial decisions affecting individuals, ensuring that models are transparent, unbiased, and compliant with data privacy regulations remains a critical responsibility for both researchers and practitioners.

Limitations of the Study

Despite its contributions, this study acknowledges several limitations. The datasets used are publicly available benchmarks that may not fully reflect the complexity and diversity of real-world financial fraud environments. The reliance on static datasets does not account for concept drift, whereby fraud patterns evolve over time. Additionally, synthetic oversampling through SMOTE may introduce artificial patterns that do not represent genuine fraudulent behavior. Computational constraints also limited the depth of hyperparameter optimization for deep learning models.

References

Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. *Journal of Network and Computer Applications*, 68, 90–113.

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(2), 937–953.

Ahmed, M., Mahmood, A. N., & Islam, M. R. (2016). A survey of anomaly detection techniques in financial domain. *Future Generation Computer Systems*, 55, 278–288.

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.

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.

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.

Carneiro, N., Figueira, G., & Costa, M. (2017). A data mining based system for credit-card fraud detection in e-tail. *Decision Support Systems*, 95, 91–101.

Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357.

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.

Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. *IEEE Symposium Series on Computational Intelligence*, 159–166.

Dornadula, V. N., & Geetha, S. (2019). Credit card fraud detection using machine learning algorithms. *Procedia Computer Science*, 165, 631–641.

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874.

Fiore, U., De Santis, A., Perla, F., Zanetti, P., & Palmieri, F. (2019). Using generative adversarial networks for improving classification effectiveness in credit card fraud detection. *Information Sciences*, 479, 448–455.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232.

Gao, M., Chen, J., Li, S., & Wang, F. (2019). A fraud detection model based on feature engineering and ensemble learning. *IEEE Access*, 7, 89690–89701.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.

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.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., & Liu, T. Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, 30, 3146–3154.

Levi, M., Reuter, P., & Halliday, T. (2018). Can the AML system be made more effective and, if so, how? *Crime, Law and Social Change*, 69(2), 209–228.

Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008). Isolation forest. *Proceedings of the IEEE International Conference on Data Mining*, 413–422.

Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774.

Maes, S., Tuyls, K., Vanschoenwinkel, B., & Manderick, B. (2002). Credit card fraud detection using Bayesian and neural networks. *Proceedings of the 1st International NAISO Congress on Neuro Fuzzy Technologies*, 261–270.

Makki, S., Assaghir, Z., Taher, Y., Haque, R., Hacid, M. S., & Zeineddine, H. (2019). An experimental study with imbalanced classification approaches for credit card fraud detection. *IEEE Access*, 7, 93010–93022.

Nguyen, T. T., Tahir, H., Abdelrazek, M., & Babar, A. (2020). Deep learning methods for credit card fraud detection. *arXiv preprint arXiv:2012.04754*.

Phua, C., Lee, V., Smith, K., & Gayler, R. (2010). A comprehensive survey of data mining-based fraud detection research. *arXiv preprint arXiv:1009.6119*.

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.

Sahin, Y., Bulkan, S., & Duman, E. (2013). A cost-sensitive decision tree approach for fraud detection. *Expert Systems with Applications*, 40(15), 5916–5923.

Saia, R., & Carta, S. (2019). Evaluating the benefits of using proactive transformed-domain-based techniques in fraud detection tasks. *Future Generation Computer Systems*, 93, 18–32.

Sánchez, D., Vila, M. A., Cerda, L., & Serrano, J. M. (2009). Association rules applied to credit card fraud detection. *Expert Systems with Applications*, 36(2), 3630–3640.

Seeja, K. R., & Zareapoor, M. (2014). FraudMiner: A novel credit card fraud detection model based on frequent itemset mining. *Scientific World Journal*, 2014, 1–10.

Van Vlasselaer, V., 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.

Wang, C., Han, D., Liu, Q., & Luo, S. (2019). A deep learning approach for credit card fraud detection using fusion model. *IEEE Access*, 7, 149841–149853.

West, J., & Bhattacharya, M. (2016). Intelligent financial fraud detection: A comprehensive review. *Computers and Security*, 57, 47–66.

Zareapoor, M., & Shamsolmoali, P. (2015). Application of credit card fraud detection: Based on bagging ensemble classifier. *Procedia Computer Science*, 48, 679–685.