Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

Real-Time Credit Card Fraud Detection Using Machine Learning: A Comprehensive Literature Review

M. Srinath[1], T. Karthik[2], Mr. Mujthaba Gulam Muqeeth[3]

[1,2] Student, Department of Computer Engineering, Methodist College of Engineering and Technology, Abids, Hyderabad, Telangana, India.
[3]Associate Professor, Department of Computer Engineering, Methodist College of Engineering and Technology, Abids, Hyderabad, Telangana, India.

Abstract

Credit card-based financial transactions form the backbone of today's digital economy, giving millions of consumers unparalleled convenience in purchases, online payments, and fund transfers. The widespread adoption of credit card-based transactions has concurrently opened avenues to potential fraud-related losses leading to significant losses among individuals, businesses, and financial organizations. Traditional fraud detection systems, which are rule- based, give preliminary protection but are not ready to cope with complex and variable fraud patterns that adapt over time. This calls for intelligent fraud detection techniques that apply machine learning to analyze the behaviours of transactions and single out fraudulent activities from the genuine ones.

It presents a credit card fraud detection system developed using a Decision Tree classifier, considering interpretability, operational efficiency, and modeling of non-linear decision boundaries. This model examines features in anonymized transactions in order to find unusual patterns that set them apart from typical customer profiles. This roadmap contains extensive data preprocessing, the removal of duplicates, normalization, and class imbalance. Experimental evaluation showed that the Decision Tree classifier yielded reliable performance on fraud transaction detection while maintaining a low false positive rate. This proves that interpretable machine learning models can be embedded into real-world financial systems to enhance security, transparency, and trust in digital transactions.

Index Terms: Credit Card Fraud Detection, Machine Learning, Decision Tree, Anomaly Detection, Financial Security, Data Preprocessing, Classification Models.

1. Introduction

With the easy availability of online gateways for payment, e-commerce websites, and digital banking services, FINANCIAL transactions have taken a complete turn. Credit cards, especially, have become a universally accepted mode of payment as they are quick and can be used anywhere in the world [1].

However, credit card fraud sees a corresponding rise with every growth in digital transactions. Fraudulent acts range from stolen card information with purchases made without the owner's consent to complex cyber- attacks designed to expose and manipulate payment systems' vulnerabilities [2]. These fraud incidents result not only in massive

financial losses but also in user privacy compromise and loss of trust in using digital banking services. Credit card fraud is particularly challenging to detect because fraudulent transactions make up only a very small portion of total transactions. This severe class imbalance causes conventional algorithms to misclassify fraud as legitimate due to overwhelmingly skewed datasets [3].

Furthermore, the patterns in fraud change dynamically and often resemble legitimate behaviors in order to avoid detection [4]. Thus, an efficient fraud detection system should incorporate the application of machine learning algorithms that learn underlying patterns, adapt evolving fraud techniques, and provide decisions interpretable by financial institutions [5].

In this work, we propose the development of a Decision Tree-based classification model in order to identify fraudulent transactions. The usage of Decision Trees is considered particularly appropriate because of the transparency of their decisions, low inference time, and ability to capture complex decision boundaries [6]. It shall be focused on an anonymized credit card dataset, where transaction features are transformed by applying dimensionality reduction techniques to protect sensitive information [7].

This study follows structured research based on a pipeline involving data preprocessing, feature evaluation, model training, and performance evaluation that shows how classical ML algorithms can deliver reliable fraud detection. The results highlight the importance of interpretable models in financial applications and will be useful for further research on hybrid and ensemble methods to enhance fraud detection systems.

BACKGROUND AND MOTIVATION

CREDIT card fraud is one of the most persistent challenges facing financial technology. Fraudsters increasingly leverage advanced technological tools and social engineering techniques to systematically bypass traditional security measures. Fraud takes many shapes, such as account theft, card cloning, phishing, online transaction manipulation, or even the creation of a synthetic identity [8]. The aftermath of this goes beyond monetary loss to customers' trust, merchant reputation, and overall financial market stability.

Historically, fraud detection relied on manually crafted rules, such as transaction amount limits, location mismatch, and velocity checks, which flag suspicious activities. These systems, though computationally inexpensive and easy to deploy, suffered from low adaptability and high rates of false positives [9]. They failed to spot subtle or emerging fraud patterns that did not conform to predefined rules [10].

That all changed when machine learning entered the fray. Learning-based systems look back at historical data to induce the patterns of transaction behavior and predict whether new transactions are fraudulent. However, fraud detection remains a difficult task because of feature anonymization, limited visibility into customer behavior, and constantly changing fraud strategies [11]. Decision Trees help alleviate part of this problem by enabling the model to concentrate on discriminative decision rules even

when the features themselves may not be directly interpretable [12].

The background of this work is improving fraud detection accuracy by using more interpretable machine learning models. While more complex models exist, Decision Trees offer a balanced solution suitable for real-time applications, regulatory compliance, and confidence from financial auditors [13].

A. MOTIVATION

The exponential increase in digital transactions for banking, online shopping, bill payments, and subscription services creates an urgent need for fraud detection. As consumers moved more towards digital payment methods, there came a corresponding development of sophisticated means by which cybercriminals exploit system vulnerabilities. Motivational factors behind this research are:

The volume of transactions is growing. Greater use of online credit card payments means greater opportunities for fraud transactions. With machine learning, it is possible to monitor large volumes of transactions in real-time [16].

Evolving Fraud Techniques: Fraudsters constantly change their behavior in ways that evade traditional rule-based systems. Such systems are appropriately complemented by learning models to cope with these always-changing patterns [17].

Need for Transparent Models: Financial institutions need interpretable systems. Decision Trees offer a clearly understandable rule-based classification structure which can easily be used in accord with regulatory standards [18].

Real-Time Fraud Detection: Detection has to occur in real time and exactly when a transaction is in process. The algorithms of Decision Trees are quite computationally efficient; because of this, they can easily be deployed in real time. Minimizing Financial Risk: The occurrence of false negatives regarding fraud detection can lead to huge monetary losses. This factor alone is a great motivator in developing highly sensitive models for fraud prediction.

These motivations accentuate the urgent need for a highly accurate, interpretable, lightweight, and adaptable fraud detection system.

2. PROBLEM STATEMENT

CREDIT card fraud detection has several challenges and is inherently complex because:

1. Severe Class Imbalance: Fraud cases represent less than 0.2% of typical transaction datasets. Models tend to favour the majority class, causing many fraudulent transactions to be misclassified as legitimate [21].

2. Dynamic Fraud Behavior: Fraudsters' strategies continuously change. The model must generalize well to unseen and evolving fraud patterns [22].

3. Anonymised and Transformed Features: Various financial datasets anonymise the features over transactions for privacy. Hence, direct interpretability and manual feature engineering over these data cannot be performed easily [23].

4. Real-Time Constraints: Fraud detection requires near-instantaneous classification in production systems. High-latency models cannot be deployed to live financial gateways [24].

5. Interpretability Requirements: Banks are regulators require clear explanations for automated decisions, especially in the case of blocking legitimate transactions. Black- box models can violate compliance policies [25].

6. Noise and Variability: Variations in transaction patterns due to seasonal behavior, customer habits, and unusual spending make fraud detection unstable without robust modeling techniques [26].

These challenges form the basis for choosing a Decision Tree model and underpin the methodological choices described in later sections.

3. LITERATURE REVIEW

CREDIT card fraud detection has been an active area of research for more than two decades, with various methodologies investigated in fighting the constantly changing nature of financial fraud. Most early studies relied on statistical models and rule-based approaches, using fixed thresholds to identify suspected transactions [27]. While simple and interpretable, these methods were not adaptive and performed poorly in dynamic fraud settings where patterns may change frequently [28]. Machine learning gradually gained prominence as a promising alternative owing to its abilities in learning complex relationships hidden in large-scale transactional data [29].

Several classical learning algorithms have been explored by different researchers for fraud detection, including Logistic Regression, Naïve Bayes, K-Nearest Neighbors, and Support Vector Machines [30]. The systems showed reasonable performance, especially when the application domain dealt with a well-balanced dataset. However, in real-world situations where fraudulent transactions are very rare, their performance becomes very poor [31]. Most traditional classifiers failed to catch the patterns of the minority class effectively while avoiding heavy reliance on oversampling and cost-sensitive techniques [32].More recent literature has also looked into sophisticated ensemble techniques such as Random Forests, Gradient Boosting Machines, and AdaBoost[33]. These models were more resilient against imbalanced data and could learn non-linear boundaries that are better representations of fraud behavior. Despite these better performances, ensemble models often come with concerns regarding interpretability, which makes them less desirable in a high- stakes financial environment where decisions must be transparent and reproducible [34].

Deep learning techniques, represented by autoencoders and neural networks, also found their applications for transactional fraud detection [35]. Specifically, autoencoders have been trained to identify abnormal patterns of transactions by reconstructing input features and analyzing reconstruction errors [36]. While these approaches showed very good results in anomaly detection, their computational cost and blackbox nature are significant drawbacks for large-scale real-time systems [37].

Some works highlight the importance of XAI in financial systems, arguing that transparent models like Decision Trees do not lose relevance even with the increasing popularity of deep learning methods [38]. With Decision Trees, fraud patterns could be effectively detected by inducing understandable rules that correspond to human decision-making logic [39]. Their strength in handling categorical and numerical

variables with fast inference time makes them suitable for live fraud detection systems [40].

In general, the literature demonstrates the evolution from rule-based detection to supervised machine learning and finally to deep learning models. However, how to balance interpretability, performance, and real-time capability remains an open challenge [41]. The paper adds to the literature by providing a Decision Tree- based approach that strikes a practical balance between model transparency and predictive accuracy.

4. OBJECTIVES

OBJECTIVES OF THE RESEARCH The
objectives of this research are formulated to address the challenges identified in current fraud detection systems. The primary goals include:

1. To design a Decision Tree–based fraud detection classifier which can classify transactions into either legitimate or fraudulent ones by inducing transparent decision rules understandable and hence validatable by financial institutions.

2. Enhance the data quality using robust preprocessing: removal of duplicates, handling missing values, normalization of numerical features, and preparation of balanced training samples to enable effective learning.

3. High sensitivity to find fraud cases with tuning/model parameter tuning for minimum false negatives, so as not to miss high-risk transactions.

4. The ability to provide real-time detection capability by developing a model that is computationally efficient and capable of providing instant predictions suitable for live transaction monitoring systems.

5. Supporting deployment in practical environments through the design of the system for easy integration with existing banking infrastructure, following regulatory compliance, and scalable transaction monitoring.

These correspond to the goals the current literature identified for research and addresscritical gaps in current fraud detection systems.

5. SYSTEM ARCHITECTURE

The proposed credit card fraud detection system is designed using a modular architecture that ensures robustness, interpretability, and performance in real time. Organized flow of data is maintained in the architecture by arranging it into sequential steps wherein each component plays its role in fraud detection. This modularized architecture allows for easy debugging, maintainability, and future enhancements without hurting the whole system.

At an extremely high level, the architecture of the fraud detection system consists of the modules:

1. Data Ingestion Module: This module is for loading raw transaction files and parsing them for preprocessing.
2. Preprocessing Module: It removes duplicates, cleans missing values, and normalizes numeric data.

3. Module for Handling Class Imbalance: resampling strategies are utilized along with class weighting so that fraud is adequately represented in the data.
4. Feature Processing Module: Feature identification, transformation, and preparation for training.
5. Model Training Module: This module will utilize pre-processed features and train a Decision Tree classifier.
6. Fraud Prediction Module: It observes and classifies new transactions into either legitimate or fraudulent.
7. Post-processing Module: Verification of prediction, threshold-based flagging and integration of alerts.
8. Evaluation and Reporting Module: This module generates performance metrics, confusion matrices, and analysis reports.

A. SYSTEM ARCHITECTURE DETAILS

1. Collecting Data and Handling Input: The dataset includes anonymized credit card transactions represented by numeric attributes preprocessed using Principal Component Analysis, PCA [7]. The system consumes such data as input in a format that is ready for processing. Anonymity reduces interpretability of features but offers privacy and allows the application to be compliant with financial regulations, such as PCI-DSS.

2. Preprocessing Pipeline: This step is very important to ensure data quality.

* Duplicate Removal: Ensures consistency in data and avoids biases in the model because of repeated fraudulent entries.

* Missing Value Handling: Confirms dataset completeness and prevents disruptions in training by using different strategies of imputation or removal.

* Normalization: Scales numerical features by Min-Max or Z-score normalization in order to avoid dominance from high magnitude attributes.

• Noise Reduction: It finds and removes the abnormalities or outliers of transaction values.

3. Class Imbalance Management: In order to handle the extreme imbalance of fraud <<1%, the system applies:

Class weighting within the Decision Tree algorithm.

oversampling techniques like SMOTE: Synthetic Minority Over-sampling Technique [42].

Stratified sampling for the creation of balanced train-test splits.

This ensures that the classifier sees enough fraud samples while training.

4. Decision Tree Model Construction: The classifier is constructed taking into consideration the following tradeoffs to balance performance and interpretability:

The main measure used for splitting is Gini impurity.

Overfitting is avoided by tuning hyperparameters to control the depth of the trees.

The minimum sample constraints are set for stable leaf nodes.

* It uses class-weight adjustments to increase sensitivity toward the minority class (fraud).

The interpretability of Decision Trees means that the bank can review the exact decision path for every flagged transaction.

5. Prediction and Classification Layer: The model, once trained, predicts fraud likelihood on incoming transactions. The transactions that are classified as suspicious are highlighted for a manual review or automated intervention such as blocking the transaction and/or requesting secondary authentication.

6. Post-processing and Thresholding: Depending on the risk appetite of the institution, the operational thresholds can be tuned to dynamically bring an optimal balance between false positives and false negatives. Again, for real-world deployment in banking systems, this is a much-neededflexibility.

6. METHODOLOGY

The research methodology adopted in this work is structured and systematic to ensure that the proposed model for fraud detection works reliably and efficiently with high interpretability. Given the sensitive nature of financial data, coupled with operational constraints related to real banking systems, the methodology will give balanced focus to aspects related to data preparation, model construction, performance evaluation, and practical deployment considerations.

A. Data Acquisition & Preprocessing

It starts by acquiring the dataset of anonymized credit card transactions from [7]. It is very necessary to perform preprocessing in order to increase the quality of the data, such as duplicate removal, treatment of missing values, and Min-Max scaling for numerical features. Since features are PCA-transformed, the domain-specific engineering is limited; hence, robust general preprocessing is even more vital.

B. Handling Class Imbalance

Because fraudulent transactions are less than 0.2% of all data, class imbalance was handled by doing a combination of class- weight adjustments within the Decision Tree algorithm and stratified sampling during data splitting. This procedure penalizes misclassifications of fraud instances more, which might allow the model to learn the patterns of the minority class much better without drastically affecting the distribution of the original data.

C. Model Development

A Decision Tree classifier is implemented by means of the Scikit-learn library for Python.The most important hyperparameters are optimized using grid search in order to prevent overfitting and ensure generalization:

max_depth: Limits the depth of the tree.

min_samples_split: Minimum number of samples required to further split an internal node.

min_samples_leaf : The minimum number of samples required to be at a leaf node.

class_weight: Here we use 'balanced' to automatically weight inversely proportional to class frequencies.

This model is trained using 70% of the preprocessed and stratified data.

D. Model Evaluation and Validation

The final classifier is tested on a separate 30% hold-out test set. Since this is an imbalanced problem, inherently, the performance is measured with a suite of metrics beyond accuracy:

Precision: The correctness of the fraud alerts.

Recall Sensitivity: It measures the model performance about detecting actual fraud cases.

F1 Score: The harmonic mean between precision and recall.

ROC-AUC Score: This determines the model's capability in distinguishing between classes at all possible thresholds.

Confusion Matrix: A detailed classification of true positives, false positives, true negatives, and false negatives.

Such multifaceted assessment can give full insight into operational strengths and weaknesses of the model.

7.RESULTS AND EVALUATION

EVALUATING the performance of a credit card fraud detection system requires metrics that truly reflect the performance of a model in identifying the rare fraudulent class. The Decision Tree model was tested on the hold- out test set, with the performance metrics summarized as shown in Table I below.

TABLE I: PERFORMANCE METRICS OF THE DECISION TREE CLASSIFIER

Metric	**Value**
Accuracy	**99.4%**
Precision	**0.82**
Recall	**0.78**
F1-Score	**0.80**
ROC-AUC	**0.94**

While the overall accuracy is high at 99.4%, this is mostly an artifact of the class imbalance. More informative metrics are precision at 0.82 and recall at 0.78. Recall at 0.78 means that the model identifies 78% of all fraudulent transactions, which is crucial for minimizing financial loss. Precision at 0.82 means that 82% of the transactions that were flagged as fraud were indeed fraudulent, helping in controlling the operational cost of false alarms. The F1- score of 0.80 and a high ROC-AUC of 0.94 indicate a strong overall discriminatory power.

Further insight is given by the confusion matrix: while the majority of the predictions lie on the diagonal (correct classifications), the off-diagonal cells show the error pattern of the model. A limited but not zero number of FNs corresponds to missed fraud cases, and thus a direct financial risk; on the other hand, a larger number of FPs corresponds to legitimate transactions wrongly flagged as fraud, thus inconveniencing customers. This trade-off between FNs and FPs can be adjusted by changing the classification threshold of the post-processing layer according to the desired risk tolerance of the financialinstitution.

8. LIMITATIONS

ALTHOUGH the fraud detection system based on a Decision Tree shows tremendous promise, a number of its limitations need to be recognized: Dataset Anonymization: The use of PCA-transformed, anonymized features limits model interpretability at a domain level and impedes meaningful, behaviour-based feature engineering [23]. Model Overfitting: Despite regularization, Decision Trees can still easily overfit to noise in the training data, which gets worse with a large number of features, possibly reducing generalization to new, evolving fraud patterns [43]. Static Training Data: The model is trained on a historical snapshot. Fraud strategies are rapidly changing; this implies that model performance may decay over time without a mechanism for continuous or periodic retraining with fresh data [22]. Operational Deployment Gap: While this work focuses on model development and offline evaluation, real-world deployment involves several other challenges including system integration latency, API design, scalability under peak load, and continuous monitoring in a production environment [24]. Inherent Trade-offs: Class weighting can diminish class imbalance, but increases false positives. Tuning a model for very high recall often comes at the cost of precision, overwhelming fraud investigation teams with alerts [44]. These limitations outline clear avenues for future work and system improvement.

9. FUTURE SCOPE

The identified limitations and the continuously changing face of both fraud and ML technology reveal several promising directions for further work: Ensemble and Hybrid Models: Integrating the Decision Tree into ensemble methods, such as Random Forest or Gradient Boosting (XGBoost, LightGBM), might provide better predictive robustness with higher accuracy and still maintain some degree of interpretability via some importance metrics of features [45]. Advanced Learning Architectures: Hybrid models that include supervised classifiers, such as Decision Trees, and unsupervised anomaly detection techniques, like Isolation Forest or Autoencoders, may lead to better detection of unseen novel fraud patterns [46]. Real-Time Streaming Pipeline: An end-to-end pipeline using stream processing frameworks (e.g., Apache Kafka, Apache Flink) would empower real-time analysis and instant

decisioning on live transaction streams [47]. XAI Integration: Enhancing the model with post-hoc explainability tools, SHAP, or LIME will enable the user to get more granular, quantitative explanations for individual predictions with an increase in trust and adherence to regulations [48]. Feature Enrichment and Federated Learning: This could be further improved by collaborating with financial institutions on contextual feature enrichment, such as merchant category, geolocation, and device ID. Collaborative model training across different institutions is possible with privacy-preserving federated learning without necessarily sharing the raw data [49].

10. CONCLUSION

THIS research covers an in-depth study on fraud detection in credit card transactions using a Decision Tree–based machine learning model. Fraud detection remains one of the critical challenges that financial institutions face due to the low rate of fraudulent activities, the dynamic nature of fraud patterns, and the need for real-time interpretable responses. The implemented Decision Tree classifier, enhanced with strategic class weighting and rigorous preprocessing, effectively balanced interpretability with computational efficiency and predictive accuracy. With a recall of 0.78 and an F1-score of 0.80, this classifier showed very strong skills in spotting fraudulent transactions while sustaining reasonable precision to limit false alarms. This study underlines the crucial importance of preprocessing steps in data analysis and the appropriate evaluation metrics for imbalanced classification problems. While promising, the model does present limitations with respect to dataset anonymization, probable overfitting, and static training that highlight further areas for refinement. Future work should involve the integration of ensemble methods, creation of real-time streaming pipelines, and application of advanced explainability techniques in order to construct a more robust, adaptive, and industry-ready solution for fraud detection.

References

1. J. West and M. Bhattacharya, "Intelligent Financial Fraud Detection: A Comprehensive Review," Computers & Security, vol. 57, pp. 47–66, 2016.
2. R. J. Bolton and D. J. Hand, "Statistical Fraud Detection: A Review," Statistical Science, vol. 17, no. 3, pp. 235–255, 2002.
3. A. Dal Pozzolo, O. Caelen, Y.-A. Le Borgne, S. Waterschoot, and G. Bontempi, "Calibrating Probability with Undersampling for Imbalanced Classification," in Proc. IEEE Symposium Series on Computational Intelligence, 2015.
4. C. Phua, V. Lee, K. Smith, and R. Gayler, "A Comprehensive Survey of Data Mining-basedFraud Detection Research," arXiv preprint arXiv:1009.6119, 2010.
5. P. K. Chan, W. Fan, A. L. Prodromidis, and S. J. Stolfo, "Distributed Data Mining in Credit Card Fraud Detection," IEEE Intelligent Systems, vol. 14, no. 2, pp. 67–74, 1999.
6. L. Breiman, J. Friedman, R. A. Olshen, and C. J. Stone, Classification and Regression Trees. Belmont, CA, USA: Wadsworth, 1984.
7. Kaggle, "Credit Card Fraud Detection Dataset," 2023. [Online]. Available: https://www.kaggle.com/mlg-ulb/creditcardfraud
8. "The Changing Face of Fraud: 2023 Report," Association of Certified Fraud Examiners (ACFE), 2023.

9. S. Bhattacharyya, S. Jha, K. Tharakunnel, and J. C. Westland, "Data mining for credit card fraud: A comparative study," Decision Support Systems, vol. 50, no. 3, pp. 602–613, 2011.
10. A. C. Bahnsen, D. Aouada, A. Stojanovic, and B. Ottersten, "Feature engineering strategies for credit card fraud detection," Expert Systems with Applications, vol. 51, pp. 134–142, 2016.Y. Sahin and E. Duman, "Detecting credit card fraud by decision trees and support vector machines," in Proc. Int. MultiConf. Engineers and Computer Scientists, 2011, vol. 1.
11. T. M. Mitchell, Machine Learning. New York, NY, USA: McGraw-Hill, 1997.
12. A. Dal Pozzolo, "Adaptive Machine Learning for Credit Card Fraud Detection," Ph.D. dissertation, Université Libre de Bruxelles, 2015.
13. N. Carneiro, G. Figueira, and M. Costa, "A data mining based system for credit-card fraud detection in e-tail," Decision Support Systems, vol. 95, pp. 91–101, 2017.
14. "Global Payments Report," Worldpay from FIS, 2023.
15. V. V. Zaslavsky and A. Strizhak, "Credit card fraud detection using self- organizing maps," Information and Security, vol. 18, pp. 48–63, 2006.
16. A. Correa Bahnsen, A. Stojanovic, D. Aouada, and B. Ottersten, "Improving credit card fraud detection with calibrated probabilities," in Proc. SIAM Int. Conf. Data Mining, 2014.
17. F. Doshi-Velez and B. Kim, "Towards A Rigorous Science of Interpretable Machine Learning," arXiv preprint arXiv:1702.08608, 2017.
18. J. H. Friedman, "Greedy function approximation: a gradient boosting machine," Annals of Statistics, vol. 29, no. 5, pp. 1189–1232, 2001.
19. Y. Liu, A. H. Gandomi, and M. Imran, "Machine learning for credit card fraud detection: A survey," IEEE Access, vol. 9,pp. 157989–158006, 2021.
20. H. He and E. A. Garcia, "Learning from imbalanced data," IEEE Trans. Knowl.
21. Data Eng., vol. 21, no. 9, pp. 1263–1284,2009.
22. I. Žliobaitė, "Learning under concept drift: an overview," arXiv preprint arXiv:1010.4784, 2010.
23. GDPR, "Regulation (EU) 2016/679 of the European Parliament and of the Council," Official Journal of the European Union, 2016.
24. M. Stonebraker, U. Çetintemel, and S. Zdonik, "The 8 requirements of real-time stream processing," ACM SIGMOD Record, vol. 34, no. 4, pp. 42–47, 2005.
25. "Principles for the Management of Credit Risk," Basel Committee on Banking Supervision, Bank for International Settlements, 2000.
26. T. Fawcett and F. Provost, "Adaptive fraud detection," Data Mining and Knowledge Discovery, vol. 1, no. 3, pp. 291–316, 1997.
27. R. Agrawal and R. Srikant, "Fast algorithms for mining association rules," in Proc. 20th Int. Conf. Very Large Data Bases (VLDB), 1994, pp. 487–499.
28. K. R. Seeja and M. Zareapoor, "FraudMiner: A novel credit card fraud detection model based on frequent itemset mining," The Scientific World Journal, vol. 2014, 2014.
29. A. K. Jain, R. P. W. Duin, and J. Mao, "Statistical pattern recognition: A review," IEEE Trans. Pattern Anal. Mach. Intell., vol. 22, no. 1, pp. 4–37, 2000.
30. E. A. Duman and M. H. Ozcelik, "Detecting credit card fraud by genetic algorithm and scatter search," Expert Systems with Applications, vol. 38, no. 10,

31. pp. 13057–13063, 2011.N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: synthetic minority over-sampling technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357,2002.
32. M. Kubat and S. Matwin, "Addressing the curse of imbalanced training sets: one- sided selection," in Proc. 14th Int. Conf. Machine Learning (ICML), 1997, pp. 179–186.
33. L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1,pp. 5–32, 2001.
34. C. Rudin, "Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead," Nature Machine Intelligence, vol. 1, no. 5, pp. 206–215, 2019.
35. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, no. 7553,pp. 436–444, 2015.
36. J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," Special Lecture on IE, vol. 2, no. 1, pp. 1–18, 2015.
37. Z. C. Lipton, "The mythos of model interpretability," Queue, vol. 16, no. 3, pp. 31–57, 2018.
38. A. B. Arrieta et al., "Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI," Information Fusion, vol. 58, pp. 82–115, 2020.
39. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. 31st Int. Conf. Neural Information Processing Systems (NeurIPS), 2017, pp. 4768–4777.
40. J. R. Quinlan, C4.5: Programs for Machine Learning. San Mateo, CA, USA: Morgan Kaufmann Publishers, 1993.
41. M. T. Ribeiro, S. Singh, and C. Guestrin, ""Why should I trust you?": Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135–1144.
42. N. V. Chawla et al., "SMOTE: Synthetic Minority Over-sampling Technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321–357,2002.
43. T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, 2nd ed. New York, NY, USA: Springer, 2009.
44. A. Correa Bahnsen, D. Aouada, A. Stojanovic, and B. Ottersten, "Feature engineering strategies for credit card fraud detection," Expert Systems with Applications, vol. 51, pp. 134–142, 2016.T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 785– 794.
45. F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in Proc. 8th IEEE Int. Conf. Data Mining, 2008, pp. 413–422.
46. P. Carbone et al., "Apache Flink: Stream and batch processing in a single engine," IEEE Data Eng. Bull., vol. 38, no. 4, 2015.
47. S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. 31st Int. Conf. Neural Information Processing Systems, 2017, pp. 4768–4777.
48. B. McMahan et al., "Communication- efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artificial Intelligence and Statistics (AISTATS), 2017, pp. 1273–1282.