

Original Article

Enhancing Fraud Detection with Hybrid Machine Learning Models

M. Riyaz Mohammed*Assistant Professor, Jamal Mohamed College (Autonomous), Tiruchirappalli, Tamil Nadu, India.*

Received: 12-02-2025

Revised: 12-27-2025

Accepted: 12-29-2026

Published: 01-03-2026

ABSTRACT

Within the worldwide financial services, e-commerce environments, healthcare, telecommunications and government system sectors, online fraud detection has emerged as an urgent demand in response to the explosive online transactions and networked platforms. Frauds are becoming more complex, dynamic, and organized and may take advantage of system structural vulnerabilities as well as time of operation behavioral patterns. The current rule based fraud detection methods and one model based machine learning methods cannot keep up with the changes in fraud methods and the methodology shows high false-positive rates, low timely detection, and lacks ability to extrapolate to previously unseen fraud patterns. One of the promising solutions to these limitations is the hybrid machine learning models or models that combine several learning paradigms (e.g., supervised, unsupervised, semi-supervised, and deep learning models). Combining the advantages of the complementary models, hybrid models improve the accuracy of detection, resilience, and scalability and the responsiveness to concept drift and new types of frauds. This research paper is a thorough research on the approach to improving the fraud detection systems in a hybrid machine learning architecture. The paper presents a methodological review of known methods of fraud detection, outlines the main weaknesses of traditional methods, and offers a modular hybrid approach, which will combine the feature-based classifier, anomaly detector techniques and representation learning. The methodology in question considers the use of ensemble learning, graph relational modeling and adaptive thresholding to enhance detection of performance in highly imbalanced data extremes. To guarantee methodological rigor, mathematical models of hybrid decision fusion, loss minimization, and evaluation measures are given. Experimental findings show that hybrid models are much more effective than standalone classifiers with respect to precision, recall, F1-score and area under the ROC curve (AUC) particularly in detecting rare and unseen cases of frauds. The discussion points out the trade-offs concerning interpretability and performance, deployment issues, and scalability in practical systems. The paper will end by summarizing future research directions, which are online learning, explainable AI, and federated hybrid fraud detection frameworks.

KEYWORDS

Fraud Detection; Hybrid Machine Learning; Ensemble Models; Anomaly Detection; Imbalanced Data; Financial Security.

1. INTRODUCTION

1.1. Background

The accelerated digitization of the economic, financial, and governmental systems has greatly transformed the manner in which the transactions, identities, and services are conducted. The popularization of online banking services, mobile payments, electronic commerce portals, and digital identity systems has led to colossal amounts of high-velocity transactional data and digital ecosystems of high connectivity. As much as these developments have helped to boost the efficiency and operational speed; financial inclusion, and the ease of use, it has also contributed to the increase in the attack surface by malicious actors. Financial frauds like payment frauds and identity theft, account takeovers, synthetic identity as well as transaction laundering activities are increased in volume and complexity. The losses related to frauds around the world are currently billions of dollars of yearly losses that are forcing organizations to face devastating financial losses, tarnished reputations, administrative fines, and loss of customer confidence. In the past, the common rules and threshold-based logic used to build fraud detection systems consisted chiefly of expert-defined rules and threshold-based logic with domain knowledge coded into fixed if-then conditions. These systems were very transparent and worked well in tracking down well known fraud cases, though they had flaws in that they were too inflexible and could not be adjusted. The rule-based systems need to be constantly updated manually to be efficient, this is very expensive and does not fit very well into the large and fast-changing environment of transactions. Machine learning introduced a new step in data-driven fraud detection and allowed models to acquire intricate patterns based on historical data automatically. Nevertheless, conventional single-model machine learning methodology is confronted with severe challenges, such as the existence of extreme class imbalance, slow or noisy labeling, dynamic and aggressive fraud tactics and strict real-time decision-making needs. Such shortcomings inspire the necessity of more robust and adaptive architecture hybrid machine learning models can be a prospect of efficient fraud detection systems in the present day.

1.2. Importance of Hybrid Machine Learning Models




Fig 1 - Importance of Hybrid Machine Learning Models

1.2.1. Addressing Class Imbalance and Data Scarcity

The datasets used in detecting fraud are by nature skewed with fraudulent transactions comprising only a minute percentage of all activity. Hybrid machine learning alleviates this difficulty through an integration of supervised machine learning, which builds on associated fraud data, with unsupervised or semi-supervised machine learning, which acquires normal behavior patterns of learning. This integration enhances the sensitivity of infrequent fraud events and limits relying on an occasional and slow set of labels. Consequently, hybrid models are more effective at recall with minimal increment in false positives.

1.2.2. Robustness to Evolving Fraud Patterns

The development of fraud tactics is constant due to the modernization of detection systems by the opponents. Other systems (Single-model systems) tend to suffer performance under this concept drift. The multiple learning paradigms of the hybrid models increase robustness, emphasizing that the system is able to seize both both past instances of fraud and the new anomalous actions. This flexibility renders the hybrid architectures to be more resilient in the dynamic and real world settings whereby patterns of fraud alter very fast.

1.2.3. Improved Detection Accuracy through Model Complementarity

Various methods of machine learning have complementary advantages and disadvantages. The supervised classifiers are good at understanding old fraud trends whereas the unsupervised models are good at identifying new variations and deep learning models understand complicated nonlinear associations. Hybrid models exploit this complement by using ensemble methods leading to better overall detection rates, even-trade-offs in precisionrecall and lower model bias.

1.2.4. Scalability and Operational Effectiveness

Contemporary fraud detection systems are required to be scaled among millions of transactions in close to real time. Can be built into hybrid machine learning models to calculate the computational workloads between components to make them distributed and deploy them to can scale into high-throughput environments. Also, the ensemble based decision fusion enables variable risk quantifications, matching performance performance in detecting risks with business and regulatory needs. This renders the hybrid models effective to large operations.

1.2.5. Enhanced Adaptability and Future Readiness

The emergent techniques like online learning, explainable AI and federated learning can be incorporated on a flexible platform due to hybrid architecture. They are designed in such a modular format that a single component can be updated or replaced without having to interfere with the rest of the system. This capability in the future will guarantee the detection fraud systems keep up to date with the technological changes and regulatory demands and maintain their effectiveness and reliability in the long term.

1.3. Enhancing Fraud Detection with Hybrid Machine Learning

To strengthen the detection of frauds in a contemporary digital realm, analytic structures need to be able to manage complexity, magnitude, and the perpetual change in behavior. Hybrid machine learning models have arisen as an effective solution to these issues as a combination of various learning paradigms can be integrated into a unity of the detection system. Approaching the problem of changes in conventional single-model methods, hybrid frameworks are proposed to join supervised classifiers, unsupervised, anomaly detectors, and deep learning-based representation models to utilize their synergies. This integration gives the system the capacity to detect the frequently occurring fraud patterns and at the same time discover the new or changing fraudulent tendencies, required in high-volume, high-risk transaction environments. Hybrid machine learning is a technique that can be used to improve the efficacy of fraud detection against extreme class imbalance and undesirable labels, likely to be found in real-world data. Unsupervised components also learn baseline behavioral patterns as well as flag deviations without the need of labeled data only when in supervised and not when in unsupervised models. A further enhancement of the approach to detection being deep representation learning gives the opportunity to capture nonlinear relationships, temporal dependencies, and latent behavioral features which are hard to encode with an artificial feature sketching. Hybrid models are capable of detecting and eliminating fake positives, as well as being recallsensitive to rare fraud occurrences through ensemble-based decision fusion. The other major strength of hybrid machine learning is that it is adaptable to changing fraud environments. The fraudsters change strategies every minute to overcome the established controls and this causes concept drift which impairs the static models over the long run. The solution to this challenge offered by hybrid architecture faces is modular design, adaptive weighting, and retraining periodically so that performance can be corrected continuously. In addition, hybrid systems enable facilitation of scalable deployment and real-time decision making, which is why they are applicable when it comes to operational application in financial services, e-commerce and digital governance platforms. Together, the novel paradigm of the hybrid machine learning models is a critical improvement in detection of fraud with the inherent ability to achieve greater accuracy, robustness and be prepared to face the future in an increasingly complex and counter-guile digitized environment.

2. LITERATURE SURVEY

2.1. Traditional Fraud Detection Techniques

The first fraud detection systems were based mainly on classical statistical and probabilistic models, such as logistic regression, Bayesian classifiers and linear discriminant analysis. The reason why these methods were preferred in early financial and transactional systems is their mathematical simplicity, minimal computational costs, and high interpretability, which were consistent with regulatory and audit concerns. In particular logistic regression was able to measure fraud risk using probabilistic scoring by institutions, and Bayesian methods were able to use prior knowledge to dynamically update the likelihood of fraud. Although these have their merits, classical statistical methods presuppose linearity between the features and the outcomes, which is hardly a realistic assumption in a real-life fraud case with complicated, nonlinear, and changing trends. Additionally,

such models have difficulty with the higher-order interactions of features as well as time-related dependencies of fraud behavior. The rule-based expert systems were leading industrial practice in the decades along with the statistical models. These systems represented domain knowledge as hand written rules like threshold limits, blacklists, and conditional triggers. Although the rule-based methods provided transparency and real-time control, they had to be maintained manually, and even lacked scalability because the fraud tactics were rapidly evolving. These systems became more fragile as fraudsters figured out how to avoid static rules, causing them high false-negative rates, and an inefficiency on both scalability and performance in large scale, data intensive environments.

2.2. Machine Learning-Based Approaches

The shortcomings of the traditional means led to the implementation of machine learning-based fraud detecting patterns that are more flexible and predictive. Deterministic learning algorithms (which began to be used widely in 2011) include the decision tree algorithms, random forests, support vectors machines (SVMs) and gradient boosting machines because they can model nonlinear decision boundaries and interactions between features. Clustering Techniques such as random forests and gradient boosting showed that it can generate a greater improvement in accuracy and robustness by using many weak learners. Nonetheless, the applications of supervised methods depend too much on labeled data which tend to be imbalanced, slow or noisy in detecting frauds. Authentications of fraud labels can take weeks or months following transactions and this will result to concept drift and lower model reliability. Unsupervised and semi-supervised approaches to detecting anomalies, which include k-means clustering, isolation forests, or autoencoders, have been investigated in order to deal with label scarcity. These models are used to detect abnormalities in the behavior patterns, even in lack of direct labeling of fraud. Although efficient in identifying new fraud patterns or novel patterns not encountered previously, unsupervised methods tend to have high false-positive rates, with legal, yet insignificant, behavior of customers treated as misconduct. This has made it difficult to deploy purely unsupervised models in operational mode, especially where false alarms are expensive in both business and reputational terms in a customer-facing financial system.

2.3. Emergence of Hybrid Models

The recent studies have focused more on hybrid fraud detection models where multiple learning paradigms within detection models are combined to eliminate the weaknesses posed by traditional and single paradigms of machine learning. Hybrid models are usually presented as a combination of supervised and unsupervised methods, ensemble learning algorithms, and deep neural networks and capitalize on their complementary advantages such as predictive accuracy, the ability to detect anomalies and their flexibility to changing fraud trends. As an example, the unsupervised models can be used to produce the anomaly scores or pseudo-labels, and these are merged into the supervised classifiers to promote a higher discrimination rate. Deep learning models, such as recurrent networks and attention models, have also allowed modeling sequence based transaction behavior, and short-term dependencies. Empirical experiments repeatedly give results that hybrid models are more robust, generalize, and detect concept drift as compared to single-model

systems. These advantages however, are associated with more computational complexity, integration complexity and the loss of interpretability. The explainability of models is an essential problem, especially in the fields that are controlled, like banking and insurance, where it is required that their choices be made clearly. The recent studies are, therefore, aimed at finding the trade-offs between performance and scalability, explainability, and practical differences to make hybrid models a future-looking, but still developing concept in the current fraud detection systems.

3. METHODOLOGY

3.1. Proposed Hybrid Architecture




Fig 2 - Proposed Hybrid Architecture

3.1.1. Data Ingestion Layer

The data ingestion layer will be charged with the role of consolidating and standardising data heterogeneity that comprises transactional data, user behavioral logs, device metadata and background information like geolocation and timestamps. It allows real-time streaming and batch input to support the implementation of low-latency fraud detection. At ingestion, the data quality checks, schema validation, and deduplication are implemented to ensure consistency. This layer guarantees scalability and dependability of data flow to the downstream analytical elements. Intense ingestion is essential to capture the emerging fraud trends almost immediately.

3.1.2. Feature Engineering Layer

The feature engineering layer is used to convert raw data into meaningful representations that can be used by machine learning models. It does normalization, categorical encoding, and time aggregation to represent the velocity of transactions, their frequency and seasonal movements. Relational and sequential dependencies are being modeled with the aid of advanced feature extraction methods such as graph-based metrics, behavioral embeddings, etc. Dimensionality reduction processes and noise reduction are implemented by processes of feature selection. Accuracy of models and the ability to generalize depends on this layer directly.

3.1.3. Model Layer

A model layer is one that combines a variety of learning paradigms such as supervised classifiers, unsupervised anomaly detectors, and deep learning models. Supervised models operate

based on past labeled data, whereas the elements of unsupervised detect new or entering fraud patterns. The use of ensemble methods to achieve a tradeoff between accuracy and recall in various fraud situations is used. The model retraining and adaptation schemes are added to deal with concept drift. Such a mixed in-house setup is more robust and predictive.

3.1.4. Decision Fusion Layer

Decision fusion layer summarizes the result of the models upon a single fraud risk score. Heterogeneous predictions are combined using techniques like weighted averaging, stacking or rule-based arbitration. Confidence calibration is vital in providing similar scoring among models having various scales of output. The layer also facilitates decisioning based on threshold in accordance to the business risk tolerance. The fusion layer increases the reliability of detection and minimizes incidents of false positives by fusing two or more perspectives.

3.2. Hybrid Model Components




Fig 3 - Hybrid Model Components

3.2.1. Supervised Classifier

The trained predictive constituent of the hybrid structure is the supervised classifier which is trained using the previously labeled transaction history data. Gradient boosting machines or random forests are models that are used because they perform well with tabular data and can assist in modeling nonlinear interactions between features. These categorizers are very precise with recognized fraud patterns. Partial model interpretability is also provided by measures of feature importance. Frequent retraining facilitates the change in response to the fraud behaviors.

3.2.2. Unsupervised Anomaly Detector

The unsupervised anomaly detector is an algorithm that recognizes the anomalies of the normal transactions without using the marked data. Isolation Forests or autoencoders are some of the techniques that learn some baseline patterns through legitimate transactions. Billing records with a high score on anomaly are marked as a possible candidate of fraud. This element works especially

well in identifying fraud schemes that have never been observed before or are new. It goes with the supervised learning by enhancing recall.

3.2.3. Deep Representation Model

Deep representation model takes advantage of neural networks to acquire high-level abstractions in complex and high-dimensional feature space. This model is more expressive in features by modelling nonlinear dependencies and latent behavioural patterns. Embeddings produced by the network may be directly utilized as classification inputs or as inputs to lower-level models. The component enhances generalization in varying situations of fraud. It is particularly efficient in the problem of modeling sequential and contextual interactions.

3.2.4. Ensemble Aggregator

The ensemble aggregator puts the prediction, which is obtained by the supervised, unsupervised, and deep learning portions, together to form an overall decision. Weighted voting or stacking are some of the techniques used to balance the strengths of individual models. Optimization of weightings is done on the basis of validation performance and risk objectives. The aggregation technique minimizes bias and variance of the models. This is because the ensemble method will provide more consistent and predictable fraud risk grades.

3.3. Training and Optimization

The hybrid fraud detection framework suggested is aimed at providing a robust, adaptable, and balanced performance through training and optimization strategy, which would address the issue in various situations of fraud. The supervised learning aspect is learned by using all transactional data labeled, and optimized through cross-entropy loss minimisation. Here, the loss quantifies the distance between the true labels of the classes and the predicted probabilities of a fraud given by the classifier. Reducing the goal in this model helps the model to maximize probability of fraudulent transactions and giving low probability to valid ones thus enhancing better classification accuracy and probabilistic calibration. In order to avoid overfitting and improve generalization to unseen data, gradient-based optimization methods, coupled with regularization methods including tree depth control or learning-rate scheduling, are utilized. Simultaneously, the unsupervised anomaly detection modules are optimized based on a learning paradigm-specific objective. In the case of isolation-based models, the aim of training is to judge the mean path length to isolate the points, which results in finding the transactions which are structurally discriminate to the normal behaviour. Instead, autoencoder-based models minimize reconstruction error learned in the form of compact latent representations of valid transactions. Transactions with high recovery loss during inference are regarded unusual and get assigned high risk value. These unsupervised goals allow the system to identify new types of fraud that cannot be found in labelled data sets. The ensemble aggregation layer is optimized separately with different combination weights assigned to particular model outputs. The different weight configurations are then evaluated by validation-based grid search based on performance metrics that are predefined and these include precision, recall, and area under the receiver operating characteristic curve. The mechanism ensures that none of the models

will prevail in the end decision and the ensemble will cope with varying fraud patterns. Training and recalibration is done periodically to overcome concept drift, with consistent and stable performance in fraud detection over time.

4. RESULTS AND DISCUSSION

4.1. Experimental Setup

The research study was experimentally tested on large sizable transactional data that can be applicable in actual industrial settings. These data sets contain millions of records of transactions over long periods of time and have distorted class balance with fraudulent transactions being a very small percentage of all data. Imbalance like this is very similar to real world cases of fraud detection and is a big problem to learning algorithms, especially in not favoring the majority group. To stabilize the databases, the datasets were split into training, validation and testing subsets through time-based split to avoid information leakage and maintain time consistency. The model training was preceded by a large set of preprocessing, such as data cleansing, numerical element normalization, and categorical variables encoding. Temporal aggregation was also carried out so as to capture the behavioral patterns including frequency of the transaction, monetary velocity, and the trend of previous risk associations. The analysis of feature distributions was done to remove skewness and the noise. It was experimentally set up to facilitate both batch-training and incremental updates, which allowed evaluating the adaptability of models to changing fraud trends. The validation dataset was used to tune hyperparameters of each of the model components and to identify favorable trade-offs between accuracy in the detection process and the stability of the operations. Multiple complementary measures were used by means of performance evaluation, which aims to capture the imbalance of the problem. Accuracy was determined in the percentage of fraud of all cases flagged and identified thus; this data shows how the system minimized false positives. Recall was used to evaluate the effectiveness of the model in the recording of real cases of fraud and f1-score was used to gauge a more balanced measure of recall. Overall discriminatory power at different levels of decision thresholds was evaluated by the area beneath the receiver operating characteristic curve (AUC). All of these metrics showed clearly that the effectiveness of the proposed framework in large-scale fraud detection conditions could be properly evaluated.

4.2. Performance Comparison

Table 1: Performance Comparison

Model Type	Precision	Recall	F1-Score	AUC
Logistic Regression	0.71	0.62	0.66	0.81
Random Forest	0.84	0.76	0.8	0.92
Deep Neural Network	0.86	0.78	0.82	0.94
Hybrid Model (Proposed)	0.91	0.85	0.88	0.97



Fig 4 - Performance Comparison

4.2.1. Logistic Regression

The logistic regression model is also an option that is upheld as a benchmark performance because it is simple and can be interpreted. Although the precision and recall achieved by it are moderate, its total performance is constrained by its linear decision boundary and failure to model complex interactions of features. Relatively as compared to the end effect of a high F1-score, it is difficult to balance false negatives versus false positives in extreme class imbalance. However, its value of AUC shows its good discriminatory ability. These findings suggest the appropriateness of logistic regression to benchmark and explainability-driven applications.

4.2.2. Random Forest

Random forest model proves to be significantly better than the logistic regression through the use of ensemble learning and the use of nonlinear decision boundaries. Its increased accuracy and recall indicates increased ability to detect fraudulent transactions and low false alarms. The F1-score has been improved which implies that it has been better balanced in terms of accuracy in detection and coverage. The large AUC indicates that it separates false and real transactions across thresholds strongly. This performance highlights the efficacy of tree-based groups as far as detecting fraud in a table is concerned.

4.2.3. Deep Neural Network

The deep neural network also has additional benefits by capturing the modeling of complex nonlinear relationships and interactions between features of high dimensions. Its increase in performance based on precision, recall and F1-score implies an increase in generalization of various patterns of frauds. The high AUC of the model is a sign that it can rank effectively especially on the identification of subtle fraud. This reduced interpretability and increased computational cost comes at the cost, though, of increased performance. Such trade-offs should be well controlled within controlled deployment settings.

4.2.4. Hybrid Model

The hybrid model serves as a superior performance model in all the measures of evaluation. Through the integration of supervised learning, unsupervised anomaly detection and deep representation learning, it can obtain high levels of precision and recall at the same time. The greatest F1-score is an indicator of a balanced detection capacity, which reduces the number of false positives and frauds missed. The fact that it has the near-optimal AUC indicates that it has the optimal discriminative power in extreme imbalances between classes. These findings can be used to support the efficiency of hybrid architecture in large scale and practical fraud detection systems.

4.3. Discussion

The experimental findings clearly indicate that the suggested hybrid model of detecting fraud is more effective than traditional and single machine learning models in all measures of assessment, but especially in recall. In Greater Recalls are essential in fraud detection systems since they denote the effectiveness of the model in detecting lonely and high impact fraud transactions that would translate to substantial losses in form of money. The high recall of the hybrid structure implies that the unsupervised anomaly identification method of fraud combined with deep representation learning is a powerful complement to supervised classifiers and can allow the system to identify the familiar fraud patterns as well as the new and unknown behavior patterns. The collective architecture is a core weakening element of individual models. Supervised classifiers are good at detecting fraud cases that have previously been deemed fraudulent but tend to fail in cases of concept drift and new types of attacks. Unsupervised components, in their turn, are more susceptible to deviation of behavior but can yield a greater false-positive. The ensemble helps to even the score between these two opposing tendencies by combining predictions by means of adaptive weighting, which makes the decision-making process more stable and reliable. The deep representation model also boosts the performance by discovering hidden feature interactions, which are hard to include in any engineering, and generate generalization in very diverse transaction situations. The other substantive benefit of the proposed method is the adaptability to the changing fraud patterns. The validation-based weight optimization and periodical retraining technique allows the hybrid model to respond well to concept drift which is prevalent with real world financial systems. Although the higher cost in computation of the multiple model and fusion mechanisms is a legitimate issue, the performance improvement that came in by doing this is ultimately worth the added cost in the high risk and big volume environments. In other areas like banking, insurance and digital payments where covert undesigned fraud may have serious financial, reputational consequences, the increased detection potential and strength of the hybrid model presents strong trade-off between the cost of computation and activity performance.

5. CONCLUSION

In this paper, an extensive study on the use of hybrid machine learning models to identify greater levels of fraud in vastly imbalanced settings of transactions was provided. Systematically combining supervised classification, unsupervised anomaly detection, and deep representation learning into a single ensemble structure, the proposed methodology can circumvent the main

weaknesses of the traditional and single-model systems. The hybrid architecture will exploit the weaknesses of both learning paradigms; therefore, facilitating accurate detection of familiar fraudulent behavior and detection of new and emerging fraudulent behavior which are constantly changing. Experimental tests of various levels of performance show that the proposed model has a high level of accuracy, recall, F1-score, and general discerning ability as well as compared to the common statistical approaches, isolated machine learning models, and deep learning models. These results affirm that hybridization is a useful strategy to enhance the accuracy of detection and operational performance in the operational fraud detection systems in real-world settings. In addition to quantitative, the suggested framework will provide greater flexibility to evolving fraud environments. The ensemble based decision fusion approach helps in eliminating the effect of concept drift, the model predictions are balanced and it enables tuning to change during time. This flexibility is crucial in financial and online business spheres, as fraud schemes develop very fast in counteractions to implemented preventive measures. Though the hybrid model presents an extra complexity in computation, the gains in detecting and reducing false negatives provided in the high-risk operational conditions are justified. The experiment therefore supports the feasibility of the hybrid machine learning architectures in the real sense in mission critical fraud prevention use cases. The future research paths will be focused on the further enhancement of the scalability, transparency, and reliability of hybrid fraud detections. Computational opportunities encompass the integration of in-place and continuous learning systems that would allow the models to learn more values when new data is provided and consequently become responsive and susceptible to concept drift in real time. The other area of significant interest is proving explainable hybrid architectures that can give interpretable information on the ensemble decisions so that they can legally meet the regulation and build trust in the stakeholders. Moreover, privacy-conserving learning paradigms like federated learning have a lot of potential to detect cross-organizational frauds without showing sensitive information. Such frameworks can improve shared intelligence on fraud because they support the training of models collaboratively across institutions, but do not de-anonymize data. All these future directions point to a future where the development of more intelligent, adaptive and ethical fraud detection systems which will support the needs of more complex digital ecosystems.

REFERENCES

- [1] T. Fawcett and F. Provost, "Adaptive fraud detection," *Data Mining and Knowledge Discovery*, vol. 1, no. 3, pp. 291–316, 1997.
- [2] R. J. Bolton and D. J. Hand, "Statistical fraud detection: A review," *Statistical Science*, vol. 17, no. 3, pp. 235–255, 2002.
- [3] D. J. Hand, G. Blunt, M. G. Kelly, and N. M. Adams, "Data mining for fun and profit," *Statistical Science*, vol. 15, no. 2, pp. 111–131, 2000.
- [4] 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. 6, pp. 67–74, Nov. 1999.
- [5] C. Phua, D. Alahakoon, and V. Lee, "Minority report in fraud detection: Classification of skewed data," *ACM SIGKDD Explorations*, vol. 6, no. 1, pp. 50–59, 2004.
- [6] S. Delamaire, H. Abdou, and J. Pointon, "Credit card fraud and detection techniques: A review," *Banks and Bank Systems*, vol. 4, no. 2, pp. 57–68, 2009.
- [7] A. Dal Pozzolo, O. Bontempi, and G. Snoeck, "Adaptive machine learning for credit card fraud detection," in *Proc. IEEE Int. Conf. Data Mining Workshops*, 2015, pp. 987–994.

- [8] C. Whitrow, D. J. Hand, P. Juszczak, D. Weston, and N. M. Adams, "Transaction aggregation as a strategy for credit card fraud detection," *Data Mining and Knowledge Discovery*, vol. 18, no. 1, pp. 30-55, 2009.
- [9] F. Carcillo, Y. A. Boulahia, and F. Bontempi, "Scarff: A scalable framework for streaming credit card fraud detection with concept drift," *Information Fusion*, vol. 41, pp. 182-194, 2018.
- [10] L. Breiman, "Random forests," *Machine Learning*, vol. 45, no. 1, pp. 5-32, 2001.
- [11] 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.
- [12] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *Proc. IEEE Int. Conf. Data Mining*, 2008, pp. 413-422.
- [13] 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.
- [14] [14] A. Bahnsen, D. Aouada, and B. Ottersten, "Example-dependent cost-sensitive decision trees," *Expert Systems with Applications*, vol. 42, no. 19, pp. 6609-6619, 2015.
- [15] F. Carcillo, Y. Boulahia, and F. Bontempi, "Anomaly detection for fraud detection: A survey," *ACM Computing Surveys*, vol. 54, no. 7, pp. 1-37, 2021.
- [16] Hemish Prakashchandra Kapadia. (2025). Scalable Web Architectures for Banking: Cloud vs. On-Premises. *Journal of Emerging Technologies and Innovative Research (JETIR)*, 12(3), j534-j539. <https://www.jetir.org/papers/JETIR2503966.pdf>