Artificial Intelligence in Fintech: A Conceptual Framework for AI-Driven Fraud Detection and Risk Management

Sunil Kumar. V[1,*], Manasa J[2], Dr. G Murali Krishna[3], Arjun Kumar K V[4], Anu Revamma Parvathi[5], Chinmayee S Dixit[6]

[1]Assistant Professor, Department of Commerce, Dayananda Sagar college of Arts, Science and Commerce, Kumaraswamy Layout, Bangalore, Karnataka, India – 560 111

[2]Assistant Professor, Department of Commerce, NMKRV College, Jayanagar 3rd Block East, Bangalore, Karnataka, India – 560 011.

[3]Assistant Professor, Department of Commerce and Management, RJS FIRST GRADE COLLEGE, KORAMANGALA, Bangalore, Karnataka, India – 560 034.

[4]Assistant Professor, Department of Commerce, Dayananda Sagar college of Arts, Science and Commerce, Kumaraswamy Layout, Bangalore, Karnataka, India – 560 111.

[5]Assistant Professor, Department of Commerce, NMKRV College, Jayanagar 3rd Block East, Bangalore, Karnataka, India – 560 011.

[6]Assistant Professor, Department of Commerce, NMKRV College, Jayanagar 3rd Block East, Bangalore, Karnataka, India – 560 011.

*Corresponding Email: sunil.vishwamanava@gmail.com

Received: 12-Dec-2025 **Revised:** 23-Feb-2026 **Accepted:** 25-June-2026

Abstract:
Artificial Intelligence (AI) and financial technology (fintech) are making an impactful stride toward revolutionizing financial fraud detection and risk management within the financial sector. Fintech and artificial intelligence (AI) are driving a significant shift in the way financial institutions detect fraud and manage risk. This paper collects from the most recent studies found in Scopus that focus on fraud detection and risk management systems in the fintech sector, from rule-based systems to machine learning (ML), deep learning (DL), ensemble and graph-based systems. The study synthesizes the results of a structured review on peer-reviewed journal and conference papers, highlighting four major technological trends: (1) Supervised learning and ensemble learning for transaction-level fraud classification, (2) Deep learning and graph networks for relational fraud and money-laundering detection, (3) Explainable AI (XAI) for regulatory transparency, and (4) Governance systems for algorithmic bias and systemic risk. The paper presents an integrated conceptual framework with five layers: Data Ingestion, Feature Engineering, Model Inference, Explain ability and Governance, and Feedback-driven Risk Monitoring, based on these strands. The diagram below shows how explain ability and regulatory compliance can be integrated into AI pipelines instead of being added on as an afterthought. Finally, implications for practitioners and regulators and future avenues of empirical study are discussed, such as federated learning with privacy protection and cross-institutional benchmarking. The results add to the ever-expanding interdisciplinary scholarship on fintech, and provide a reference framework for both research and practical risk management.

LDA is a specialized type of artificial intelligence that processes data into a graph or network structure to detect patterns and relationships among entities. LDA is a specialized form of AI designed to analyze data and identify patterns and relationships between entities, which can then be represented in a graph or network format.

Keywords: AI, Fintech, LDA, Risk Management, Machine Learning

1. Introduction

Financial technology, or fintech, is no longer a disruptive force in the traditional banking world, but a significant part of the financial infrastructure that drives the international financial system. At the heart of this shift is the integration of Artificial Intelligence (AI) into the core financial operations, such as fraud detection and risk mitigation, algorithmic trading, and customer advisory, among other areas (Belanche et al., 2019; Nguyen et al., 2023). Card-not-present fraud and complicated money-laundering schemes that exploit cross-border or decentralized payment systems are merely two examples of the sophistication and scale of financial crime that has emerged amid the proliferation of digital payment channels, peer-to-peer platforms and mobile payment applications (Rocha-Salazar et al., 2021).

With many traditional rule-based fraud detection systems employing static thresholds and manually built heuristics, it is becoming more difficult to keep up with the pace, volume and variety of transactional data being generated by today's fintech environments (Flondor et al., 2024). This deficiency has spurred a significant research effort on machine learning and deep learning algorithms that can learn fraud signatures directly from data, which are complex and non-linear (Hafez et al., 2025). In parallel, financial institutions and regulators have voiced some concerns over the opacity of many of the AI models, leading to a massive amount of research on explainable AI in the context of credit and fraud decision-making processes (Bussmann et al., 2020; De Lange et al., 2022).

Although there are many technical contributions, the literature is still scattered in disciplinary silos computer-science venues focus on the performance of algorithms, while finance and information-systems journals focus on the governance, ethics and adoption of these systems. Very few studies combine these strands in a clear and logical conceptualization of how the choices made in creating a model relate to the risk-governance needs. This paper fills this void.

The aims of this paper are then three-fold: (1) to collate the literature from Scopus indexed publications on AI applications for fintech fraud detection and risk management; (2) to extract commonalities across the papers in terms of both technological and governance; and (3) to suggest an integrated conceptual framework for data-level, model-level, and governance-level aspects of AI-enabled fraud detection and risk management systems. The rest of the paper is organized as follows. The literature review is presented in Section 2, in four thematic strands. A description of the methodology used in the review is provided in Section 3. The proposed conceptual framework is given in Section 4. Benefits, risks and challenges are addressed in Section 5. The practical and policy implications are discussed in Section 6, followed by limitations and directions for the future in Section 7, which are followed by the conclusion.

2. Literature Review

2.1 AI and the Future of Fintech Risk Management

The study of fintech has grown significantly since the mid-2010s, and in the past few years, there has been an increasing number of papers on the academic and regulatory interest in automated decision-making in the financial technology sector reviewed in the Scopus database, especially after 2018, indicating a surge in publications on this topic. In the customer-facing aspect of this transition, Belanche et al. (2019) investigated how useful and trusted robo-advisories are for their customers, the same metrics that would apply to internal risk assessment systems incorporating AI. Nguyen et al. (2023) highlighted the symbiotic nature of big data, AI, and fintech at a macro-level: big data can help power AI algorithms, while AI algorithms can help to create new products and risk management strategies in the financial sector. The COVID-19 period offers a complementary example of systemic risk from shadow banking from Bao and Huang (2021), which describes how many fintech platforms are now providing credit, many of which are AI-based, and can either ease or exacerbate systemic risk depending on the strength of the underlying credit and fraud models. Likewise, Northey et al (2022) identified that consumer confidence in AI-driven financial advice depends on consumers' perceptions of the competence and transparency of the AI service provider, which further emphasizes the importance of explain ability in fintech risk governance, as outlined in Section 2.4.

2.2 Machine Learning and Deep Learning concepts related to fraud detection.

The most extensive body of work on Scopus concerns the classification of transactions using supervised machine learning and deep learning. Hafez et al. (2025) performed a systematic review of the credit card fraud detection methods enhanced by AI, identifying and filtering a vast collection of studies published in the Scopus-indexed literature, and found that ensemble and hybrid architectures generally outperform the single-algorithm baseline, especially in scenarios where the data is imbalanced, as is common with credit card fraud datasets. To this end, Mienye and Sun (2023) showed that a deep learning ensemble with data-resampling techniques was found to significantly enhance the recall rate of the minority fraudulent-transaction classes while maintaining a similar precision rate. Following this, various other studies have further enhanced this ensemble paradigm, with Ahmed et al. (2025) using a hybrid sampling approach alongside classifiers for a balanced precision-recall trade-off, Gupta et al. (2025) introducing strong feature-selection routines to complement the stacking ensemble, and Zhao et al. (2024) applying a variant of the gradient boosting method, namely an improved LightGBM, to handle extreme class imbalance in transactional data. Tayebi and El Kafhali (2025) also fine-tuned an XGBoost classifier, achieving improvements in detection accuracy by using Bayesian hyperparameter tuning, compared to its default settings.

In addition to ensemble methods, there has been much use of deep neural architectures like auto encoders for anomaly-based fraud detection, due to their ability to learn the distribution of valid transactions and identify deviations as anomalies. Hemdan and Manjaiah (2022) used an auto encoder algorithm for anomaly detection in credit card transactions and Shabad and Kavitha (2018) and Sadgali et al. (2019) explored the neural-network classifiers at the point of sale and within transaction pipelines in general, respectively. This was furthered by the work of Parmar et al. (2021) using deep architectures trained to the structured transaction features, and Akouhar et al. (2025) by using variants of Kolmogorov–Arnold networks with dynamic oversampling to tackle the persistent imbalance challenges. Ayoub et al. (2025) presented an alternative to purely statistical resampling by proposing a granular-computing framework that is more robust against noisy transaction labels due to the granularity-aware representations. This is consistent with the literature to date, with Flondor et al. (2024) demonstrating that decision-tree algorithms are competitive in performance on the card-fraud detection task, despite their relative simplicity, and provide a level of transparency that is not offered by deep architectures.

This chapter introduces two types of approaches that rely on graph structures:

2.3 Graph-Based and Relational Approaches.

A second big area is about the relational aspects of financial fraud, which indicate that fraudulent transactions are frequently not isolated, but rather part of a coordinated effort across related accounts, merchants, or devices. Graph neural networks (GNNs) are a natural fit for this problem since they can use the information from the local neighborhood in a transaction's network. Alarfaj and Shahzadi (2024) proposed one real-time fraud prevention architecture using GNNs and auto encoders for two case studies in banking that achieved better precision-recall performance than non-relational baselines. To illustrate this, Harish et al. (2024) have performed a comparative analysis of classical methods, deep learning methods and graph-based methods, and found that graph-based methods are especially helpful in detecting groups of fraudsters that evade classical feature-based classifiers. To tackle this, Fan et al. (2023) introduced a heterogeneous GNN with attribute enhancement and structure-aware attention mechanisms, and Xie et al. (2023) designed a time-aware, attention-based gated network that explicitly considers the temporal order of transactional behavior. For real-time fraud-scoring systems, Xu et al. (2023) suggested deep boosting decision trees, which offer a more computationally economical alternative with high accuracy gain while also reducing inference latency.

2.4 Anti-Money Laundering and Transaction Monitoring

Anti-money laundering (AML) laws, which were once largely rule-driven and static alerting processes, have also been transformed by machine learning. Jullum et al. (2020) created a supervised ML model based on the historical data of alerts and investigations from a large Norwegian bank and found that their model, trained from reported and non-reported alerts, outperforms the bank's current rule-based model in prioritizing alerts for manual review. Reite et al. (2025) extended this line of research and demonstrated that a machine learning-based client-risk classification can have a meaningful impact on the efficiency of AML detection by directing investigative resources correctly to high-risk clients. Focusing on the importance of having unsupervised anomaly scores in the absence of labeled cases of money laundering and terrorism financing, Rocha-Salazar et al. (2021) introduced a neural network-based indicator to detect both types of illicit activities. On a real dataset and a synthetic AML dataset, Masood et al. (2026) compared the performance of a Hybrid Model of stacking classifiers, namely Random Forest, Support Vector Machine, and LogitBoost, with individual classifiers and observed that the hybrid model yielded better results than the classifiers, with respect to precision, recall, and F1-Score. With many AML data sets having the problem of scarce labels, several studies have tried graph and network-based representations, for example, Luna et al. (2018) did anomaly detection to find shell-company account networks, and Luo (2014) proposed suspicious-transaction detection heuristics that were precursors to later machine-learning methods; more recently, Oliveira et al. (2020) employed machine learning to detect laundering activity within the Bitcoin block chain, extending AML methods beyond the conventional banking rails into decentralized finance.

2.5 Explain ability, Fairness and Regulatory Alignment

One recurring theme through literature reviewed is between predictive performance and interpretability. Bussmann et al. (2020) proposed one of the earliest applied frameworks for explainable AI in fintech risk management, where they apply Shapley-value decomposition to explain credit-risk predictions for peer-to-peer lending, and that regulated financial services cannot responsibly use opaque, “black box” models. De Lange et al. (2022) built on this argument for bank credit assessment, showing that post-hoc explainability methods can align the accuracy and transparency expectations of bank credit models and supervisors. The results are also aligned with other information-systems literature on algorithmic bias, which Kordzadeh and Ghasemaghaei (2022) compiled and summarized in terms of causes and impacts of algorithmic bias in decision-making processes; financial services, in particular, are vulnerable to reputational and legal damages from biased algorithms affecting credit access. Gatzert and Schubert (2022) put these worries in a broader perspective of cyber-risk management in banking and insurance and demonstrated that the use of the term "AI-related operational and compliance risk" is becoming more common in textual risk disclosures, a category that has become a driver of firm value. The overall theme highlights the fact that explainability, fairness, and governance are not abstract concepts but integral aspects of the design of AI systems in financial environments with regulation.

3. Research Methodology

To align with the recent literature-driven conceptual synthesis approach, as done in other Scopus-based reviews on the use of AI in finance (e.g., Hafez et al., 2025), this study implements a literature-based conceptual synthesis approach. Instead of dealing with primary transactional data, the paper systematically combines peer-reviewed journal articles and conference papers listed mostly in Scopus-indexed journals and conferences (including IEEE Access, Journal of Big Data, Journal of Risk and Financial Management, Journal of Money Laundering Control, Neural Computing and Applications, Expert Systems with Applications, and other venues) to build a conceptual framework that is based on the findings from the available literature.

The search strategy consisted of three groups of keywords: (i) technology keywords, (ii) application keywords, and (iii) governance keywords. Similar to the practice followed in other fintech fraud detection reviews (Hafez et al., 2025), the studies were selected based on the following criteria: relevance to fintech fraud detection or risk management, being published in a peer-reviewed journal or established conference proceedings, and the most recent studies, which were published between 2018 and 2026 to capture the period of greatest activity in AI-fintech research. Studies examining purely theoretical machine-learning techniques with no financial-services application or studies involving nonfinancial services like health care or e-commerce fraud, but with a technique directly applicable to financial fraud detection, were ineligible unless they brought a technique to the table of financial fraud that was directly applicable.

Because this is a conceptual paper, it is suitable to use this qualitative synthesis approach which allows for the integration of otherwise fragmented technical and governance literatures without producing new primary empirical estimates. This restriction will be discussed again in Section 7.

4. Proposed Conceptual Framework

Based on the above four thematic strands, this paper suggests a 5-layer conceptual framework to detect and manage fraud in fintech using AI (Table 1). It provides a framework that explicitly links technical model selections to governance needs, both of which are discussed separately in the literature, to guide the design of the system as well as study of it.

Table 1. Five-Layer Conceptual Framework for AI-Driven Fraud Detection and Risk Management

Layer	Function	Representative Techniques / Sources
1. Data Ingestion & Integration	Aggregates transactional, behavioral, and network data	Multi-source transaction and graph data (Oliveira et al., 2020)
2. Feature Engineering & Representation	Converts raw data into model-ready features, sequences, and graphs	Temporal features (Xie et al., 2023); graph structures (Fan et al., 2023; Harish et al., 2024)
3. Model Inference	Generates fraud/risk scores	Ensembles (Ahmed et al., 2025; Gupta et al., 2025); gradient boosting (Zhao et al., 2024; Tayebi & El Kafhali, 2025); autoencoders (Hemdan & Manjaiah, 2022);

		GNNs (Alarfaj & Shahzadi, 2024)
4. Explainability & Governance	Produces interpretable justifications and fairness audits	SHAP/Shapley methods (Bussmann et al., 2020; De Lange et al., 2022); bias auditing (Kordzadeh & Ghasemaghaei, 2022)
5. Feedback & Continuous Monitoring	Feeds confirmed outcomes back for retraining and drift detection	Risk-based client scoring (Reite et al., 2025); alert-outcome learning (Jullum et al., 2020)

Layer 1 (Data Ingestion and Integration) gathers transactional, behavioral and third party data, such as device fingerprints and network graphs, that are fed downstream to models. The quality and representativeness of this layer has a direct impact on model performance, and given the difficulty in acquiring labels in some domains like AML, it remains a significant challenge (Oliveira et al., 2020).

Layer 2 (Feature Engineering and Representation) processes raw data into representations that are suitable for the model, such as statistical features engineered from the data, temporal sequences (Xie et al., 2023), and graph-structured representations of account and transaction networks (Fan et al., 2023; Harish et al., 2024).

Layer 3 (Model Inference) can be based on machine learning classifiers (such as ensemble classifiers) (Ahmed et al., 2025; Gupta et al., 2025), variant of gradient boosting (Tayebi & El Kafhali, 2025; Zhao et al., 2024), deep autoencoders (Hemdan & Manjaiah, 2022), or graph neural networks (Alarfaj & Shahzadi, 2024). The choice of model should explicitly consider the compromise between accuracy, processing time and interpretability, not applying the most complex model available.

Layer 4 (Explainability and Governance) applies post-hoc explainability methods to the outputs of Layer 3, including Shapley-value decomposition (Bussmann et al., 2020; De Lange et al., 2022), as well as fairness-auditing routines, based on the algorithmic-bias literature (Kordzadeh & Ghasemaghaei, 2022). This layer provides for the creation of the documentations that need to be submitted to the supervisory review and to the internal model risk governance.

This layer (Feedback and Continuous Monitoring) provides the results of the investigators' decisions and confirmed fraud/laundering results back to the data layer to be used for periodic retraining and drift detection. This completes the cycle of deployment to ongoing validation of the model, as recommended for AML client-risk classification systems (Reite et al., 2025).

This layered structure is intentionally general, and does not mandate one specific algorithm, instead it defines where explainability and governance needs to be placed, not an afterthought onto the completed model.

5. Discussion

5.1 Benefits and Opportunities

The literature reviewed highlights some common advantages of AI's implementation in the field of fintech fraud detection and risk management. First, machine learning and deep learning models have a higher classification performance than the traditional static rule-based models across common classification metrics, and especially recall on minority fraud classes (Mienye & Sun, 2023; Hafez et al., 2025). Second, unlike feature-based classifiers, which have been noted to have a blind spot with respect to coordinated fraud rings, graph-based approaches can detect coordinated fraud rings. Third, the use of explainable AI methods can help institutions balance the precision gains of more sophisticated models with the transparency demanded by regulators (Bussmann et al., 2020), which has been a barrier to the rollout of complex models in the past.

5.2 Challenges and Risks

Despite these advantages, however, there are a number of obstacles that remain. Class imbalance is a persistent technical challenge, with fraudulent and laundering transactions accounting for a minority of transactions, and there is a constant search for new methodology to re-sample and learn from the smaller number of "bad" transactions while keeping costs down (Zhao et al., 2024; Akouhar et al., 2025). The problem is particularly severe in the context of AML, where cases of confirmed laundering are scarce, and are often reported late (Oliveira et al., 2020; Rocha-Salazar et al., 2021). Explainability remains a key governance issue: While some tools for explainability are used, they only produce post-hoc approximations of a model's reasoning, and it is unclear for how long these approximations remain faithful to a model's actual decision process (De Lange et al., 2022). Another risk is algorithmic bias, where algorithms trained on past data can incorporate and reinforce access to credit or forgery detection bias by demographic group (Kordzadeh & Ghasemaghaei, 2022). Lastly, the growing deployment of interconnected AI systems in fintech platforms can lead to systemic-risk concerns where correlated model errors of institutions with similar models or data sets due to similar data quality may intensify financial shocks (Bao & Huang, 2021; Gatzert & Schubert, 2022).

6. Implications for Practice and Policy

According to the literature reviewed, fraud-detection and risk-management systems could be constructed to include explainability and monitoring features from the beginning, instead of adding them on as a post-deployment feature. Institutions implementing a graph-based or deep-learning model should use baseline models that can be explained, such as decision trees (Flondor et al., 2024), for comparative validation and regulatory reporting. Additionally, the repeated references to explainability in the literature (Bussmann et al., 2020; De Lange et al., 2022) indicate that for regulators and supervisory authorities, minimum standards for explainability and fairness auditing of AI models deployed in credit and fraud decisioning should be included in supervisory frameworks alongside only outcome-based performance measures. In the case of AML, the results from Reite et al. (2025) and Jullum et al. (2020) indicate that to enhance the effectiveness of limited investigative resources, risk-based resource allocation involving machine learning-based client and transaction risk scores can be employed without the need for a complete overhaul of the existing compliance frameworks.

7. Limitation and Future Research

There are several limitations to this paper. It is a conceptual synthesis not an empirical study, and, consequently, offers no new quantitative estimation of model performances, but rather refers to results from the underlying primary studies, with different characteristics, evaluations, and reporting quality of datasets. The literature search was systematic but was not comprehensive within every subject area in Scopus and there may be some under-representation of non-English-language research and/or emerging preprint literature not yet indexed. Future research might include cross-institutional benchmarking studies that compare different architectures on shared but private datasets; empirical research into the validity of post-hoc explainability methods for production fraud-detection models on real data; and studies of federated-learning architectures where institutions can train fraud models together while avoiding the centralization of privacy-sensitive transaction data. Moreover, more research is needed on the systemic-risk aspects of the use of AI and its interdependence in the financial sector, which, compared to its firm-level performance aspects, has been relatively under-researched.

8. Conclusion

The paper has compiled the studies published in Scopus on the use of AI in fraud detection and risk management in fintech, from the rule-based approach to ensemble learning, deep learning, and graph-based approaches, as well as studies on explainability, fairness, and regulatory aspects. On this synthesis, the paper presented the following conceptual framework that consists of five layers: data ingestion, feature engineering, model inference, explainability and governance, and continuous feedback monitoring. The document highlights that transparency and governance measures must be intentionally integrated into the lifecycle of the AI system, beyond simple predictions, to ensure effective risk management in the fintech sector. Given the ongoing evolution of financial crime and the need for interdisciplinary research to understand the potential of AI in risk management, the convergence of computer science and financial scholarship is vital for the continued development of effective and accountable risk management systems.

References

Ahmed, K. H., Axelsson, S., Li, Y., & Sagheer, A. M. (2025). A credit card fraud detection approach based on ensemble machine learning classifier with hybrid data sampling. Machine Learning with Applications, 20, 100675.

Akouhar, M., Ouhssini, M., El Fatini, M., Abarda, A., & Agherrabi, E. (2025). Dynamic oversampling-driven Kolmogorov–Arnold networks for credit card fraud detection: An ensemble approach to robust financial security. Egyptian Informatics Journal, 31, 100712.

Alarfaj, F. K., & Shahzadi, S. (2024). Enhancing fraud detection in banking with deep learning: Graph neural networks and autoencoders for real-time credit card fraud prevention. IEEE Access, 13, 20633–20646.

Ayoub, M., Abdelhamid, T., & Khalid, J. (2025). Granular computing framework for credit card fraud detection. Alexandria Engineering Journal, 121, 387–401.

Bao, Z., & Huang, D. (2021). Shadow banking in a crisis: Evidence from fintech during COVID-19. Journal of Financial and Quantitative Analysis, 56(7), 2320–2355.

Belanche, D., Casaló, L. V., & Flavián, C. (2019). Artificial intelligence in FinTech: Understanding robo-advisors adoption among customers. Industrial Management & Data Systems, 119(7), 1411–1430.

Bussmann, N., Giudici, P., Marinelli, D., & Papenbrock, J. (2020). Explainable AI in fintech risk management. Frontiers in Artificial Intelligence, 3, Article 26.

De Lange, P. E., Melsom, B., Vennerød, C. B., & Westgaard, S. (2022). Explainable AI for credit assessment in banks. Journal of Risk and Financial Management, 15(12), 556.

Fan, S., Liu, G., & Li, J. (2023). A heterogeneous graph neural network with attribute enhancement and structure-aware attention. IEEE Transactions on Computational Social Systems, 11, 829–838.

Flondor, E., Donath, L., & Neamtu, M. (2024). Automatic card fraud detection based on decision tree algorithm. Applied Artificial Intelligence, 38(1), Article 2385249.

Gatzert, N., & Schubert, M. (2022). Cyber risk management in the US banking and insurance industry: A textual and empirical analysis of determinants and value. Journal of Risk and Insurance, 89(3), 725–763.

Gupta, R. K., Hassan, A., Majhi, S. K., Parveen, N., Zamani, A. T., Anitha, R., Ojha, B., Singh, A. K., & Muduli, D. (2025). Enhanced framework for credit card fraud detection using robust feature selection and a stacking ensemble model approach. Results in Engineering, 26, 105084.

Hafez, I. Y., Hafez, A. Y., Saleh, A., Abd El-Mageed, A. A., & Abohany, A. A. (2025). A systematic review of AI-enhanced techniques in credit card fraud detection. Journal of Big Data, 12, Article 6.

Harish, S., Lakhanpal, C., & Jafari, A. H. (2024). Leveraging graph-based learning for credit card fraud detection: A comparative study of classical, deep learning and graph-based approaches. Neural Computing and Applications.

Hemdan, E. E.-D., & Manjaiah, D. H. (2022). Anomaly credit card fraud detection using deep learning. Studies in Big Data, 91, 207–217.

Jullum, M., Løland, A., Huseby, R. B., Ånonsen, G., & Lorentzen, J. (2020). Detecting money laundering transactions with machine learning. Journal of Money Laundering Control, 23(1), 173–186.

Kordzadeh, N., & Ghasemaghaei, M. (2022). Algorithmic bias: Review, synthesis, and future research directions. European Journal of Information Systems, 31(3), 388–409.

Luna, D. K., Palshikar, G. K., Apte, M., & Bhattacharya, A. (2018). Finding shell company accounts using anomaly detection. In Proceedings of the ACM India Joint International Conference on Data Science and Management of Data (pp. 167–174). ACM.

Luo, X. (2014). Suspicious transaction detection for anti-money laundering. International Journal of Security and Its Applications, 8(2), 157–166.

Masood, N. S., Sakr, R. H., & Eleneen, A. A. (2026). Money laundering detection in financial institutions using machine learning. Journal of Computer Science, 22(4), 1330–1343.

Mienye, I. D., & Sun, Y. (2023). A deep learning ensemble with data resampling for credit card fraud detection. IEEE Access, 11, 30628–30638.

Nguyen, D. K., Sermpinis, G., & Stasinakis, C. (2023). Big data, artificial intelligence and machine learning: A transformative symbiosis in favour of financial technology. European Financial Management, 29(2), 517–548.

Northey, G., Hunter, V., Mulcahy, R., Choong, K., & Mehmet, M. (2022). Man vs machine: How artificial intelligence in banking influences consumer belief in financial advice. International Journal of Bank Marketing, 40(6), 1182–1199.

Oliveira, C., Torres, J., Silva, M. I., Aparício, D., Ascensão, J. T., & Bizarro, P. (2020). Machine learning methods to detect money laundering in the Bitcoin blockchain in the presence of label scarcity. arXiv preprint arXiv:2005.14635.

Parmar, J., Patel, A., & Savsani, M. V. (2021). A novel approach for credit card fraud detection through deep learning. Lecture Notes on Data Engineering and Communications Technologies, 52, 191–200.

Reite, J. E., Karlsen, J., & Westgaard, E. G. (2025). Improving client risk classification with machine learning to increase anti-money laundering detection efficiency. Journal of Money Laundering Control, 28(1), 93–107.

Rocha-Salazar, J.-J., Segovia-Vargas, M.-J., & Camacho-Miñano, M.-M. (2021). Money laundering and terrorism financing detection using neural networks and an abnormality indicator. Expert Systems with Applications, 169, 114470.

Sadgali, I., Sael, N., & Benabbou, F. (2019). Fraud detection in credit card transaction using neural networks. In ACM International Conference Proceeding Series (p. 63). ACM.

Shabad, M. A. R., & Kavitha, M. (2018). Credit card fraud detection using neural networks at merchant side. Journal of Computational and Theoretical Nanoscience, 15(11), 3373–3375.

Tayebi, M., & El Kafhali, S. (2025). A novel approach based on XGBoost classifier and Bayesian optimization for credit card fraud detection. Cyber Security and Applications, 3, 100093.

Xie, Y., Liu, G., Yan, C., Jiang, C., & Zhou, M. (2023). Time-aware attention-based gated network for credit card fraud detection by extracting transactional behaviors. IEEE Transactions on Computational Social Systems, 10(3), 1004–1016.

Xu, B., Wang, Y., Liao, X., & Wang, K. (2023). Efficient fraud detection using deep boosting decision trees. Decision Support Systems, 175, 114037.

Zhao, X., Liu, Y., & Zhao, Q. (2024). Improved LightGBM for extremely imbalanced data and application to credit card fraud detection. IEEE Access, 12, 159316–159335.