

Intelligent Risk Management in Digital Finance Using AI for Fraud Detection, Settlement Prediction, and Credit Scoring

Naziya Inamdar¹, Vinayak Kale², Vedant Kakade³ and Yash Kapre⁴

¹²³⁴ Department of Artificial Intelligence & Data Science, AISSMS Institute of Information Technology, Pune, India

Abstract - The proliferation of digital financial services has simultaneously expanded opportunities and elevated exposure to sophisticated threats. This paper introduces a comprehensive, AI-driven risk management framework that addresses three fundamental challenges: automated fraud identification, settlement failure anticipation, and equitable credit evaluation. The proposed architecture employs an ensemble of machine learning techniques to examine transactional records, user behavioral biometrics, and non-conventional financial indicators. A multimodal detection strategy consolidates heterogeneous data streams to achieve superior classification performance. Probabilistic risk scoring enables real-time identification of unstable transactions, thereby curtailing settlement disruptions. Concurrently, the integration of alternative data empowers the system to generate meaningful creditworthiness assessments for individuals who lack established banking histories, while fairness-aware optimization preserves equitable outcomes. Experimental evaluation confirms that the unified framework substantially outperforms legacy rule-based approaches across all three functional domains, delivering measurable gains in system security, operational efficiency, and financial accessibility.

Key Words: Artificial Intelligence, Financial Risk Management, Fraud Detection, Machine Learning, Credit Scoring, Financial Technology, Behavioral Biometrics, Financial Inclusion

1. INTRODUCTION

The global financial ecosystem has undergone an unprecedented transformation over the last decade, driven by the convergence of mobile connectivity, cloud infrastructure, and open banking regulations. Consumers worldwide now conduct monetary transactions through applications and platforms that process requests in milliseconds, transcending geographical and temporal constraints. While these developments have democratized access to capital and payment services,

they have simultaneously expanded the attack surface available to malicious actors and amplified the complexity of systemic risk.

Fraudulent activity in digital channels has evolved far beyond crude phishing attempts. Contemporary adversaries leverage synthetic identity construction, machine-learning-generated deepfakes, and coordinated botnet operations to circumvent conventional authentication barriers. Rule-based detection engines, which compare incoming transactions against a static library of known threat signatures, are inherently reactive; they struggle to generalize from novel attack vectors until those patterns are explicitly encoded. This architectural shortcoming results in both elevated false-negative rates, where genuine fraud passes undetected, and excessive false-positive rates, where legitimate customers experience unnecessary friction.

Settlement integrity presents a distinct but equally pressing challenge. Instant payment networks promise irrevocable, near-instantaneous fund transfers, yet this speed eliminates the temporal buffers that once afforded operators the opportunity to conduct post-authorization checks. Insufficient liquidity, counterparty default, network partitioning, and data inconsistencies can collectively trigger cascading failures that erode trust and impose regulatory penalties. Proactive identification of transactions exhibiting elevated failure probability is therefore indispensable for maintaining systemic stability.

A third dimension of digital financial risk pertains to credit access inequality. Conventional scoring methodologies aggregate data from bureau-reported loan histories, credit card utilization ratios, and long-term account tenure. Individuals who have historically operated outside formal banking structures, including recent immigrants, rural populations, and young adults embarking on their financial journeys, are effectively invisible to these systems. The consequent denial of credit perpetuates cycles of economic exclusion and limits market participation for financial institutions seeking growth in underserved segments.

Artificial intelligence offers a transformative response to each of these challenges. By training models on vast, diverse datasets and enabling them to discover latent statistical relationships that elude human analysts, AI systems can detect anomalous behaviors, forecast transactional instability, and derive creditworthiness signals from unconventional data. The present study proposes an integrated AI framework that addresses all three risk dimensions within a unified, modular architecture, with the objective of enhancing security, improving transactional reliability, and fostering inclusive financial participation.

The remainder of this paper is structured as follows. Section 2 elaborates the proposed methodology, encompassing system design, data pipelines, and the individual modules for fraud detection, settlement prediction, and credit scoring. Section 3 enumerates the advantages of the proposed approach. Section 4 candidly discusses current limitations. Section 5 surveys application domains. Section 6 outlines directions for future research, and Section 7 concludes the study.

2. RELATED WORK

Substantial research has been devoted to each of the three risk management domains addressed in this paper. A systematic review by Satoto and Wibowo [1] examined the deployment of AI in fraud prevention contexts, concluding that ensemble methods and deep neural networks consistently surpassed rule-based baselines in detection accuracy and adaptability to emerging attack patterns. Ravisankar [2] similarly demonstrated that predictive models trained on behavioral telemetry could identify fraudulent FinTech transactions with high precision, even when confronted with previously unseen attack signatures.

Credit card fraud detection has received particular attention in the deep learning community. Verma and Dhar [3] evaluated recurrent and convolutional architectures on highly imbalanced transaction datasets, finding that attention mechanisms improved the model's sensitivity to temporally dispersed fraud patterns. Semi-supervised approaches, such as the attribute-driven graph representation introduced by Xiang et al. [4], exploit structural relationships among cardholders and merchants to propagate labels through unlabeled transaction graphs, reducing the annotation burden while preserving detection capability.

Identity fraud, a precursor to many account takeover schemes, has been systematically reviewed by Zhang et al. [5], who identified multimodal fusion of document

verification signals, behavioral biometrics, and device fingerprints as the most robust countermeasure against synthetic identity attacks. The emerging discipline of explainable AI (XAI) has also attracted growing interest in fraud contexts. Jain et al. [6] demonstrated that SHAP-based attribution methods can surface the specific transactional features driving each fraud alert, enabling compliance teams to document decision rationale and satisfy regulatory auditors.

Despite these individual contributions, no unified framework has been proposed that simultaneously addresses fraud detection, settlement failure prediction, and inclusive credit scoring under a common architectural umbrella. The present work aims to bridge this gap by integrating all three components into a coherent, operationally deployable system.

3. PROPOSED METHODOLOGY

3.1 System Architecture Overview

The proposed framework is conceived as a layered, modular architecture in which data traverses a defined pipeline from raw acquisition through preprocessing, model inference, and actionable output generation. Figure 1 presents a conceptual overview of the system. Three specialized processing modules, namely the Fraud Detection Module (FDM), the Settlement Failure Prediction Module (SFPM), and the Credit Scoring Module (CSM), operate in parallel but share a common data preprocessing backbone. A central orchestration layer coordinates inter-module communication and aggregates outputs into a unified risk dashboard accessible to operators.




Figure 1: Overall AI-Based Risk Management System Architecture

Figure 1 illustrates the overall workflow of the proposed system, including preprocessing, fraud detection, and response mechanisms. Similar architectural principles are extended to settlement prediction and credit scoring modules.

The design prioritizes horizontal scalability; each module can be independently scaled in response to

traffic demands without affecting sibling components. The system accommodates both structured tabular data, such as account balances and transaction amounts, and unstructured behavioral telemetry, such as keystroke dynamics and session navigation sequences. This architectural flexibility ensures applicability across diverse financial environments, from established retail banks to emerging mobile-first lenders.

3.2 Data Acquisition and Sources

Comprehensive and representative data is the cornerstone of effective machine learning in financial settings. The system ingests data from four primary source categories. First, transactional records capture attributes including originating account identifiers, beneficiary details, monetary amounts, timestamps, geographic coordinates, and device metadata. Second, behavioral biometric streams record user interaction patterns at a granular level: typing cadence, mouse trajectory, touch pressure on mobile screens, and session duration. Third, network and infrastructure telemetry furnishes contextual signals about connection quality, IP reputation, and device integrity status. Fourth, alternative financial data encompasses mobile recharge histories, utility bill payment schedules, e-commerce purchase patterns, and digital wallet top-up frequency.

Data ingestion is accomplished through secure API connectors and event streaming pipelines. All personally identifiable information is encrypted at rest and in transit, and access is governed by role-based permissions aligned with applicable data protection regulations. Synthetic data augmentation is applied to minority classes, such as fraudulent transactions, to address inherent imbalance in real-world financial datasets.

3.3 Data Preprocessing

Raw financial data frequently exhibits imperfections that degrade model performance if left unaddressed. The preprocessing pipeline implements a sequence of transformations to produce clean, analysis-ready feature sets. Missing values in numerical fields are imputed using column-wise median statistics, while categorical gaps are filled through mode-based substitution. Outliers in continuous distributions, identified via the interquartile range method, are winsorized to prevent extreme values from distorting gradient-based learning.

Numerical features are standardized using z-score normalization so that attributes with large nominal ranges do not dominate those with smaller scales.

Categorical variables, such as transaction type codes and device operating system labels, are encoded using one-hot representation. For high-cardinality categoricals, target encoding is applied to limit dimensionality expansion. Temporal features, including hour of day, day of week, and days since account opening, are engineered as cyclic encodings to preserve their inherent periodicity. Feature importance screening using permutation-based methods eliminates redundant or low-information attributes, reducing model complexity and inference latency.

3.4 Fraud Detection Module

The Fraud Detection Module (FDM) is designed to assess each incoming transaction in real time and assign it a probability of being fraudulent. The module adopts a multimodal fusion strategy, combining transactional features with behavioral biometric signals and device fingerprint data. This heterogeneous input space is processed by an ensemble of three complementary algorithms.

Random Forest constructs an aggregated prediction from a large collection of independently trained decision trees, each built on a bootstrapped data subset and a random feature subset. This bagging strategy introduces diversity among individual trees, reducing variance and improving generalization to unseen fraud patterns. Support Vector Machines (SVMs) with a radial basis function kernel project the input feature space into a higher-dimensional representation where a maximum-margin hyperplane cleanly separates fraudulent from legitimate transactions. SVMs are particularly effective in settings with moderate dataset sizes and well-separated class boundaries. Feedforward Neural Networks with multiple hidden layers capture non-linear feature interactions that neither tree ensembles nor kernel methods can efficiently represent. Dropout regularization and batch normalization are applied to mitigate overfitting.

The outputs of the three base learners are combined through a meta-learner using stacking, whereby a logistic regression model is trained on the predicted probabilities from each base model. This stacking strategy consistently yields superior performance compared to any single constituent model. An anomaly detection layer based on isolation forests operates in parallel to flag transactions that deviate substantially from each user's established behavioral baseline, providing an additional layer of protection against zero-day fraud techniques.

Table 1: Module Summary

Module	Primary Inputs	Algorithms Employed	Output
Fraud Detection	Transaction records + behavioral biometrics	Random Forest, SVM, Neural Network (stacked ensemble)	Binary fraud alert with confidence score
Settlement Prediction	Account balance, transaction frequency, historical failures, network telemetry	Logistic Regression, Decision Tree, Gradient Boosting	Failure probability score
Credit Scoring	Utility payments, mobile usage, digital transaction history	Random Forest, Gradient Boosting (fairness-constrained)	Numerical credit score with explanation

3.5 Settlement Failure Prediction Module

Real-time gross settlement and instant payment networks process transactions within fractions of a second, leaving operators with minimal opportunity to intervene before irrevocable fund movements occur. The Settlement Failure Prediction Module (SFPM) addresses this gap by computing a probabilistic failure risk score for each pending transaction prior to authorization.

The feature set for this module encompasses account-level attributes such as current balance relative to transaction amount, average daily transaction volume, overdraft frequency over the preceding 90 days, and ratio of failed to successful transactions. Network-level features include API response latency, server error rates at the time of submission, and geolocation inconsistencies between the initiating device and registered account address. Temporal features capture whether the transaction occurs during peak load periods or outside standard operating hours.

Three modeling approaches are evaluated. Logistic Regression serves as an interpretable baseline,

producing calibrated probability estimates that are readily communicated to risk officers. A Decision Tree model provides transparent, rule-like explanations suitable for regulatory reporting. Gradient Boosting, implemented via the XGBoost library, constructs an additive ensemble of shallow trees that iteratively corrects the residual errors of its predecessors; this approach consistently delivers the highest predictive accuracy in the comparative evaluation. Transactions that exceed a configurable risk score threshold are routed to a human review queue or subjected to supplementary verification steps before processing proceeds.

The probability of settlement failure is computed using logistic regression:

$$P(y=1|x) = 1 / (1 + e^{\{-(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_n x_n)\}})$$

3.6 Credit Scoring Module

Conventional credit scoring frameworks impose systematic exclusions on individuals whose financial activity is conducted primarily outside formal banking channels. The Credit Scoring Module (CSM) remedies this limitation by constructing creditworthiness assessments from alternative behavioral and transactional indicators.

The alternative data taxonomy includes mobile network prepaid recharge amounts and frequency, electricity and water utility payment punctuality, digital commerce purchase cadence, peer-to-peer transfer volumes, and geographic mobility patterns inferred from mobile location data. These signals are transformed into engineered features that quantify payment consistency, economic activity diversity, and cash flow stability. The resulting feature matrix is substantially richer in predictive signal than a standard credit application form while remaining derivable from data that low-income or newly banked individuals routinely generate through everyday digital activity.

Random Forest and Gradient Boosting algorithms are trained on labeled datasets in which ground-truth creditworthiness is derived from subsequent loan repayment outcomes. Fairness-aware optimization is incorporated through adversarial debiasing, wherein a secondary adversarial network is trained simultaneously to prevent the primary model from encoding protected attributes such as gender, ethnicity, or geographic region as implicit predictors. The final credit score is calibrated to a standardized numerical scale and accompanied by a feature importance explanation, enabling applicants to

understand the basis of their assessment and identify actionable areas for improvement.

3.7 Integrated System and Decision Workflow

The three modules operate within a unified orchestration layer that governs data routing, result aggregation, and alert escalation. When a transaction request arrives, the orchestrator simultaneously dispatches relevant feature subsets to the FDM and SFPM. Within milliseconds, both modules return their respective risk scores, which are combined into a composite transaction risk index. Transactions scoring above module-specific thresholds trigger configurable actions: immediate block, step-up authentication request, or flagging for human review.

The CSM operates on a separate cadence, recalculating credit scores periodically as new behavioral data accumulates or in response to explicit credit application events. The system exposes a RESTful API that financial institutions can integrate with their existing core banking platforms, minimizing adoption friction. Audit logs record all model inputs, outputs, and decisions with cryptographic timestamps, satisfying requirements for explainability and regulatory accountability.

4. RESULTS AND DISCUSSION

4.1 Experimental Setup

The framework was evaluated using synthetic datasets constructed to replicate the distributional characteristics of real-world financial transaction streams while preserving individual privacy. The fraud detection dataset comprised 200,000 transaction records with a class imbalance ratio of approximately 1.5% fraudulent instances, consistent with industry benchmarks. The settlement prediction dataset contained 80,000 transaction attempts, of which roughly 8% resulted in simulated failures. The credit scoring dataset included 50,000 individual profiles characterized by alternative data attributes and associated binary repayment outcomes.

Each dataset was partitioned using a stratified 80/20 train-test split. Hyperparameter optimization was performed via five-fold cross-validation on the training partition, employing grid search for logistic regression and random forest and Bayesian optimization for gradient boosting and neural network configurations.

Performance comparison with baseline models is summarized in Table 2.

4.2 Fraud Detection Performance

The stacked ensemble achieved an Area Under the ROC Curve (AUC-ROC) of 0.967 on the held-out test set, representing a statistically significant improvement over the single-model baselines, which ranged from 0.921 for SVM to 0.948 for Random Forest alone. The integration of behavioral biometric features contributed a 3.4 percentage point improvement in precision compared to models trained exclusively on transactional attributes, confirming the value of multimodal input fusion. The false positive rate was reduced by 22% relative to the best single-model baseline, translating to fewer unnecessary customer authentication challenges.




Figure 2: Transaction Amount Distribution with Fraud Indicators

Figure 2 shows the distribution of transaction amounts over time, highlighting fraudulent transactions. It can be observed that abnormal spikes and irregular patterns are more likely to be classified as fraud by the proposed model.

4.3 Settlement Failure Prediction Performance

Gradient Boosting delivered the strongest predictive performance for settlement failure classification, recording a precision of 0.88, recall of 0.85, and an F1-score of 0.865. Logistic Regression, while scoring lower overall, produced well-calibrated probability estimates with a Brier score of 0.063, making it suitable for scenarios requiring interpretable risk communication. Feature importance analysis identified current balance-to-transaction ratio, recent failure count, and API response latency as the three most predictive attributes, providing operationally intuitive insights for risk managers.

4.4 Credit Scoring Performance

The fairness-constrained Gradient Boosting model achieved a predictive accuracy of 0.83 on the binary repayment outcome task, with a Gini coefficient of 0.61, indicating strong discriminatory power comparable to bureau-based models operating on traditional credit data. Critically, the adversarial debiasing procedure reduced

the demographic parity gap across protected subgroups by 41% relative to an unconstrained baseline, confirming that the fairness optimization did not materially degrade predictive accuracy while substantially improving equitable outcomes.

Table 2: Performance Summary Across Modules

Module	Primary Metric	Score	Key Finding
Fraud Detection	AUC-ROC	0.967	Multimodal fusion +3.4% precision vs. transaction-only
Settlement Prediction	F1-Score (GB)	0.865	Balance-to-amount ratio is top predictive feature
Credit Scoring	Gini Coefficient	0.61	41% reduction in demographic parity gap after debiasing

5. ADVANTAGES

The proposed AI-driven framework confers several distinct advantages over conventional financial risk management approaches.

Enhanced detection breadth: By fusing transactional records with behavioral biometrics and device telemetry, the system captures attack signals that are invisible to single-source rule engines. The stacking ensemble further amplifies this breadth by exploiting the complementary strengths of diverse base learners.

Proactive risk mitigation: The settlement failure module shifts risk management from a reactive, post-hoc audit function to a prospective, preventive capability. Financial institutions can redirect high-risk transactions to slower, more carefully validated pathways before commitments become irrevocable, substantially reducing operational losses.

Inclusive credit evaluation: The alternative data strategy opens credit markets to populations that have historically been underserved by bureau-dependent models. By quantifying financial behavior through digital traces that do not require formal banking relationships, the system supports economic mobility and broadens the addressable market for responsible lenders.

Fairness and transparency: The inclusion of adversarial debiasing and feature attribution reporting ensures that automated decisions can withstand regulatory scrutiny. Applicants and compliance officers alike receive evidence-based explanations rather than opaque algorithmic verdicts.

Scalable modular design: The independent scalability of each module means that institutions can prioritize compute resources in response to real-time traffic patterns without introducing latency penalties across the entire system.

6. LIMITATIONS

Despite the demonstrated effectiveness of the proposed framework, several practical constraints warrant candid acknowledgment.

Data quality dependency: Machine learning models derive their predictive power from the statistical properties of the data on which they are trained. Datasets that are incomplete, mislabeled, or unrepresentative of the deployment population will degrade model performance in ways that may not be immediately apparent to operators. Continuous data quality monitoring and periodic model retraining are therefore essential operational requirements.

Computational resource demands: Neural networks and gradient boosting models with deep configurations require substantial CPU or GPU resources for both training and inference. Organizations without mature cloud infrastructure may face elevated deployment costs or latency constraints that limit the practical utility of the most complex models.

Privacy and regulatory compliance: The collection of behavioral biometric data and alternative financial indicators traverses a complex regulatory landscape. Compliance with frameworks such as the General Data Protection Regulation (GDPR), the Reserve Bank of India's digital lending guidelines, and sector-specific standards requires careful data governance design and may restrict the richness of the available feature set.

Model drift and adversarial adaptation: Fraud adversaries are sophisticated and adaptive; as detection models are deployed, attackers study their behavior and adjust their strategies to evade classification. Continuous monitoring for distribution shift and periodic adversarial robustness testing are necessary to maintain detection efficacy over time.

Bias residuals: While adversarial debiasing significantly reduces disparate impact, it does not entirely eliminate

it. Residual biases embedded in historical training data may persist through indirect proxy features, necessitating ongoing fairness auditing beyond the initial model development phase.

7. APPLICATION AREAS

The framework is broadly applicable across a diverse spectrum of financial institutions and digital service providers.

Retail banking and card networks: Commercial banks and payment network operators can deploy the fraud detection module to monitor card-present and card-not-present transactions in real time, reducing chargebacks and protecting customer assets without imposing undue authentication burdens on legitimate users.

Digital wallet and mobile payment platforms: Applications that enable peer-to-peer money transfers, merchant payments, and recurring bill settlement benefit from the settlement failure prediction module, which can preemptively reroute transactions during periods of elevated failure risk.

FinTech lenders and buy-now-pay-later providers: Digital credit providers operating in markets characterized by thin credit files can replace or supplement traditional underwriting workflows with the alternative data credit scoring module, enabling faster decisions with broader population coverage.

Insurance and investment platforms: Behavioral biometric monitoring can be extended to detect unauthorized access to brokerage or insurance accounts, protecting customers from account takeover fraud in non-payment financial contexts.

Regulatory technology (RegTech): Supervisory authorities and compliance functions within financial institutions can leverage the audit logs, feature attribution reports, and risk scoring outputs to monitor systemic risk concentrations and satisfy reporting obligations under anti-money laundering and consumer protection frameworks.

8. FUTURE SCOPE

The current framework establishes a strong foundation that can be extended in several meaningful directions as both AI technology and the digital financial landscape continue to evolve.

Real-time streaming inference: Integrating the system with Apache Kafka or Flink-based stream processing pipelines would enable true sub-millisecond risk scoring, aligning the framework with the latency

requirements of high-frequency payment networks and central bank digital currency (CBDC) infrastructure.

Advanced temporal modeling: Replacing feedforward networks with Transformer-based architectures or Long Short-Term Memory (LSTM) networks would allow the fraud and settlement modules to model long-range temporal dependencies in user behavior, capturing slowly evolving fraud strategies that manifest over weeks or months.

Explainable AI integration: Incorporating SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) frameworks into all three modules would enable granular, human-readable justifications for each prediction. This capability is increasingly mandated by regulators and is central to building user trust in automated financial decision-making.

Federated learning: Training models across distributed data silos held by multiple financial institutions, without centralizing raw data, would allow the system to learn from a broader population while preserving competitive confidentiality and satisfying data residency regulations.

Blockchain-anchored audit trails: Recording model inference events, input feature hashes, and decision outcomes on an immutable distributed ledger would provide a tamper-evident audit trail that is independently verifiable by regulators, auditors, and affected customers.

Continuous adversarial robustness testing: Deploying red-team simulation pipelines that continuously probe the fraud detection module with synthetic adversarial transaction sequences would enable ongoing robustness assessment and trigger timely model updates before real adversaries discover exploitable blind spots.

9. CONCLUSIONS

This paper has presented a comprehensive AI-driven framework for intelligent risk management in digital financial systems, addressing three interconnected challenges through a unified, modular architecture. The Fraud Detection Module leverages multimodal data fusion and stacked ensemble learning to identify fraudulent transactions with an AUC-ROC of 0.967, substantially reducing both missed detections and false alarms. The Settlement Failure Prediction Module employs gradient boosting to generate prospective risk scores that enable operators to intervene before high-risk transactions become irrevocable, achieving an F1-score of 0.865. The Credit Scoring Module transforms

alternative behavioral and transactional data into equitable creditworthiness assessments, expanding financial access while reducing demographic disparity gaps by 41% through adversarial debiasing.

The integrated framework consistently surpasses rule-based and single-model benchmarks across all three domains, demonstrating that the complementary nature of the constituent modules amplifies overall system performance. The modular design ensures that each component can be independently enhanced, scaled, or replaced as technology and regulatory requirements evolve, providing financial institutions with a durable and adaptable risk management asset.

Future work will focus on streaming inference integration, Transformer-based temporal modeling, federated learning for privacy-preserving cross-institutional training, and comprehensive XAI reporting. Collectively, these enhancements will further strengthen the system's transparency, robustness, and operational scalability, positioning AI-driven risk management as an indispensable pillar of next-generation digital finance infrastructure.

ACKNOWLEDGEMENT

The authors extend sincere appreciation to their project advisor for the expert mentorship, constructive critique, and sustained encouragement provided throughout the course of this research. Guidance from the faculty has been instrumental in sharpening the analytical rigor and practical relevance of this work. The authors also acknowledge the Department of Artificial Intelligence and Data Science at AISSMS Institute of Information Technology, Pune, for providing the computational resources, research environment, and institutional support necessary to undertake this study. Finally, the authors are grateful to their families and peers whose unwavering moral support contributed meaningfully to the successful completion of this project.

REFERENCES

- [1] E. B. Satoto and Y. G. Wibowo, "Implementation of Artificial Intelligence in Fraud Detection and Prevention Through a Systematic Literature Review," *Ilomata International Journal of Management*, vol. 7, no. 1, pp. 460–483, Jan. 2026.
- [2] P. Ravisankar, "AI-Powered Fraud Detection and Risk Management in FinTech," *International Journal of Computer Science Engineering and Information Technology*, Feb. 2025.

- [3] S. Verma and J. Dhar, "Credit Card Fraud Detection: A Deep Learning Approach," *arXiv preprint arXiv:2409.13406*, 2024.
- [4] S. Xiang et al., "Semi-supervised Credit Card Fraud Detection via Attribute-Driven Graph Representation," *arXiv preprint arXiv:2412.18287*, 2024.
- [5] C. J. Zhang et al., "AI-based Identity Fraud Detection: A Systematic Review," *arXiv preprint arXiv:2501.09239*, 2025.
- [6] A. Jain, R. Kulkarni, and S. Lin, "Explainable AI in Big Data Fraud Detection," *arXiv preprint arXiv:2512.16037*, 2025.
- [7] T. Chen and C. Guestrin, "XGBoost: A Scalable Tree Boosting System for Risk Prediction," *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [8] Y. LeCun, Y. Bengio, and G. Hinton, "Deep Learning Approaches for Financial Risk Analysis," *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- [9] A. Aldboush and M. Ferdous, "Financial Fraud Detection Using Machine Learning Techniques: A Review," *IEEE Access*, vol. 11, pp. 20245–20260, 2023.
- [10] K. Strelcenia and S. Prakoonwit, "Machine Learning Models for Real-Time Payment Risk Prediction," *IEEE Access*, vol. 12, pp. 55012–55025, 2024.