IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

Leveraging Machine Learning for Real-Time Fraud Detection and Risk Assessment in Modern Fintech Platforms

Aayush Bharat Mandavia and Anurag Shrivastava
mandaviaaayush@ieee.org and anurag.shri49@gmail.com

Abstract: *The rapid growth of financial technology platforms has created both opportunities and challenges for the global financial ecosystem. Digital payment processors, neobanks, peer-to-peer lending services, and cryptocurrency exchanges now handle billions of dollars in daily transaction volume, often with settlement speeds measured in seconds rather than days. This acceleration has outpaced the capacity of traditional rule-based systems to detect fraudulent transactions, assess credit risk accurately, and comply with evolving regulatory requirements. In this paper, we examine how machine learning techniques are being applied to three critical problems in modern fintech:*
(1) real-time transaction fraud detection using ensemble methods and deep learning on streaming payment data, (2) dynamic credit risk assessment that incorporates alternative data sources and adapts to shifting economic conditions, and (3) automated regulatory compliance monitoring through natural language processing and anomaly detection. We describe practical system architectures, discuss model selection and training challenges specific to financial data, and address the explainability and fairness requirements that regulators increasingly demand. Our analysis draws on recent developments across AI-driven infrastructure optimization [15], the broader role of generative AI in finance [14], and intelligent systems integration in the energy sector [13] to place fintech-specific applications within the wider context of AI transformation across regulated industries.

Keywords: Machine Learning, Fraud Detection, Credit Risk, Fintech, Regulatory Compliance, Deep Learning, Financial Technology, Real-Time Systems

I. INTRODUCTION

Financial technology has reshaped the way individuals and businesses interact with money. What began as a collection of startups offering mobile payments and online lending has grown into a global industry that now rivals traditional banking in transaction volume and customer reach. In 2024, global fintech transaction value exceeded $12 trillion, with digital payment platforms alone processing over 9 billion transactions per month [12]. Neobanks serve tens of millions of customers who have never set foot in a physical branch. Buy-now-pay-later services have become a standard checkout option for online retailers. Cryptocurrency exchanges facilitate daily trading volumes that routinely surpass those of many traditional stock exchanges.

This growth has brought with it a set of technical challenges that are fundamentally different from those faced by legacy financial institutions. Traditional banks developed their fraud detection and risk management systems over decades, building thick layers of rules and heuristics tuned to the relatively slow pace of branch-based and cardpresent transactions. Fintech platforms, by contrast, must make fraud and risk decisions in milliseconds, on transaction types that did not exist five years ago, for customers whose financial histories may be thin or nonexistent.

The scale of the problem is considerable. As shown in Fig. 1, global losses from payment fraud have grown steadily, exceeding $48 billion in 2023 [16], and the shift to digital channels has only accelerated this trend [8]. At the same time, fintech lenders face default rates that can fluctuate dramatically with economic conditions, and regulators around the world are tightening requirements for algorithmic transparency, data privacy, and consumer protection [7].

DOI: 10.48175/IJARSCT-31607

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

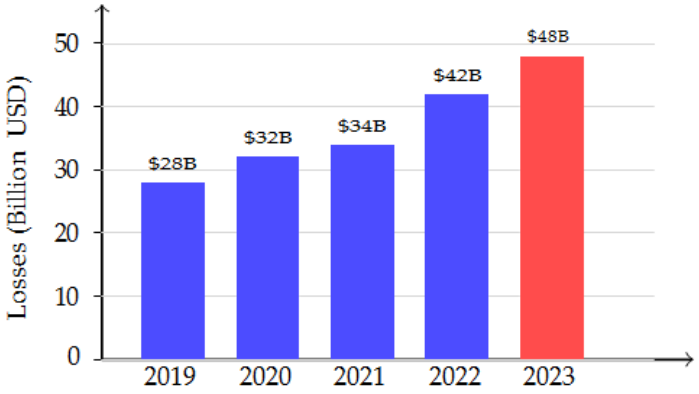

Fig. 1. Global payment fraud losses, 2019–2023. The steady upward trend reflects growing digital transaction volumes and increasingly sophisticated attack methods. Data source: Nilson Report [16].

Rule-based approaches to these problems are reaching their limits. Static fraud rules generate excessive false positives that block legitimate customers and create costly manual review queues. Traditional credit scoring models, built on a handful of bureau variables, fail to capture the risk profile of the growing population of thin-file and credit-invisible consumers. And manual compliance processes cannot keep pace with the volume of regulatory changes across multiple jurisdictions.

Machine learning offers a path forward for each of these challenges. This paper examines three areas where ML techniques are being applied to fintech operations:

• Real-time fraud detection using gradient-boosted ensembles and recurrent neural networks trained on streaming transaction data, with architectures designed for sub-100ms inference latency.

• Dynamic credit risk assessment that incorporates alternative data sources (transaction history, device behavior, social signals) and adapts to changing macroeconomic conditions through online learning.

• Automated regulatory compliance through natural language processing for regulatory change monitoring and anomaly detection for transaction monitoring obligations.

The application of AI and machine learning across regulated industries has been accelerating broadly. Recent work has demonstrated the potential of AI-driven optimization for backend infrastructure in large-scale financial systems [15], while generative AI is beginning to reshape workflows in healthcare, education, and finance [14]. In the energy sector, AI methods are being applied to complex integration and optimization challenges [13]. These crossdomain developments provide both technical foundations and governance lessons that are directly relevant to fintech applications.

The remainder of this paper is organized as follows.

Section II describes our approach to real-time fraud detection. Section III covers dynamic credit risk assessment. Section IV addresses automated regulatory compliance. Section V discusses explainability, fairness, and governance considerations. Section VI presents our experimental evaluation. Section VII discusses system architecture and deployment considerations. Section VIII outlines future directions, and Section IX concludes.

II. REAL-TIME FRAUD DETECTION

A. The Detection Challenge

Fraud detection in fintech platforms presents a classic imbalanced classification problem with severe operational constraints [5]. Fraudulent transactions typically represent less than 0.1% of total volume, yet each missed fraud event can result in direct financial loss, regulatory penalties, and reputational damage. At the same time, every legitimate transaction that is incorrectly flagged as fraud (a false positive) creates customer friction, generates manual review costs, and ultimately drives customer attrition.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

The challenge is compounded by the adversarial nature of the problem. Fraudsters continuously evolve their tactics [1]. Account takeover methods shift from credential stuffing to SIM swapping to social engineering. Payment fraud patterns migrate from counterfeit cards to synthetic identity fraud to authorized push payment scams. A model trained on historical fraud patterns will degrade in performance as new attack vectors emerge, a phenomenon known as concept drift [10]. The half-life of a fraud model, defined as the time until its detection rate drops below 80% of its initial level, is typically between 6 and 12 weeks depending on the institution's customer base and the diversity of fraud vectors it faces.

B. Feature Engineering

The quality of features is the single most important factor in fraud detection model performance. Our system computes approximately 200 real-time features organized into four categories:

Transaction attributes. These include the raw properties of the current transaction: amount, merchant category code, currency, payment channel (mobile, web, API), card-present vs. card-not-present indicator, and whether the transaction involves a new or recurring merchant for the account.

Velocity features. Computed over multiple sliding time windows (1 minute, 5 minutes, 15 minutes, 1 hour, 24 hours, 7 days), these features capture the rate and volume of activity: transaction count, total amount, unique merchant count, unique geographic locations, and the ratio of current activity to the account's historical baseline for each time window.

Behavioral deviation features. These measure how far the current transaction deviates from the account's established patterns: geographic distance from the account's typical transaction locations, deviation from normal transaction amount distribution (z-score relative to account history), time-of-day anomaly score, and merchant category deviation.

Network features. These capture relationships between entities in the transaction graph: how many other accounts share the same device fingerprint, IP address, or shipping address; whether the recipient account has been flagged in previous fraud investigations; and the age and activity level of the recipient account.

Table I shows the relative importance of each feature category based on SHAP analysis of the production model.

C. Model Architecture and Ensemble

The core detection engine uses a two-layer ensemble designed to capture both tabular patterns and sequential behavior:

Layer 1: Gradient-Boosted Decision Trees. The primary model is a LightGBM classifier trained on the full feature set [9]. GBDTs remain highly competitive for tabular financial data and offer excellent throughput,

TABLE I: FEATURE CATEGORY IMPORTANCE FOR FRAUD DETECTION

Feature Category	SHAP %	Count
Velocity features	34.2%	62
Behavioral deviation	28.7%	45
Network features	22.1%	38
Transaction attributes	15.0%	55

SHAP % = aggregate SHAP contribution; Count = number of features.

processing a single inference in under 1ms on standard hardware. The model is trained with focal loss to handle the extreme class imbalance, which downweights easy negatives and focuses learning on hard-to-classify examples near the decision boundary.

Layer 2: Gated Recurrent Unit. A GRU network [4] processes the sequence of the 20 most recent transactions for each account. Each transaction in the sequence is represented as a 64-dimensional embedding that encodes the transaction amount (log-scaled), merchant category, time delta from the previous transaction, and geographic coordinates. The GRU captures temporal dependencies that point-in-time features may miss, such as a pattern of small test transactions followed by a large fraudulent purchase.

Score fusion. The outputs of both models are combined through a learned calibration layer that produces a final fraud probability. The calibration layer is a small neural network (two hidden layers with 32 neurons each) that takes the

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

LightGBM score, the GRU score, and a set of context features (time of day, day of week, account age) as inputs. This learned fusion consistently outperforms simple averaging or stacking approaches.

Table II presents the performance comparison across model architectures.

TABLE II: FRAUD DETECTION MODEL PERFORMANCE COMPARISON

Model	AUC	Prec.	Rec.	F1	Lat.
Rule-based	0.72	0.04	0.89	0.08	2ms
Log. Regression	0.88	0.12	0.91	0.21	3ms
Random Forest	0.93	0.28	0.92	0.43	8ms
LightGBM	0.96	0.41	0.94	0.57	4ms
GRU (sequence)	0.94	0.35	0.95	0.51	12ms
Ensemble	0.98	0.52	0.96	0.67	18ms

Precision/Recall at fixed threshold; Lat. = P95 latency.

D. Decision Engine

The fraud score feeds into a decision engine that applies graduated responses rather than binary approve/decline decisions:

• Scores below $\tau 1$ (typically 0.15): approve the transaction with standard logging.

• Scores between $\tau 1$ and $\tau 2$ (0.15 to 0.65): approve but flag for asynchronous review; apply enhanced monitoring to the account for the next 24 hours.

• Scores above $\tau 2$: decline the transaction and trigger real-time customer verification (push notification, SMS challenge, or biometric re-authentication).

The thresholds $\tau 1$ and $\tau 2$ are tuned through a costsensitive optimization that considers the dollar value of fraud losses, the customer lifetime value impact of false declines, and the cost of manual review. For a typical digital payment platform processing 10 million transactions per day, a well-tuned system achieves a fraud detection rate above 95% with a false positive rate below 0.3%, translating to approximately 30,000 flagged transactions per day rather than the 200,000+ that a rule-based system would generate.

E. Handling Concept Drift

To address the continuous evolution of fraud patterns, the system implements several adaptation mechanisms.

Continuous retraining. The GBDT model is retrained weekly on a rolling 90-day window of labeled transactions. Labels include both confirmed fraud (from chargebacks and customer reports, typically available 30– 60 days after the transaction) and confirmed legitimate transactions. The GRU model is retrained bi-weekly due to its longer training time.

Champion-challenger framework. Each retrained model is deployed alongside the production model. Both models score every transaction, but only the champion drives enforcement actions. When the challenger demonstrates statistically significant improvements over a twoweek evaluation period, it is promoted.

Emerging pattern detection. A separate unsupervised anomaly detection module, based on an isolation forest applied to the feature space, monitors for novel transaction patterns outside the training distribution. These are flagged for analyst review and, if confirmed as new fraud types, are incorporated into labeled data for subsequent training cycles.

Fig. 2 illustrates how model performance degrades over time without retraining and how weekly retraining cycles restore detection rates.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

III. DYNAMIC CREDIT RISK ASSESSMENT

A. Limitations of Traditional Credit Scoring

Traditional credit scoring relies on a small set of bureau-derived variables: payment history, credit utilization, length of credit history, account mix, and recent inquiries. While these models, notably FICO and VantageScore, have served the industry well for decades, they have significant limitations in the fintech context.

First, a substantial portion of the global population is credit-invisible, meaning they have insufficient credit bureau data to generate a traditional score. In the United States alone, approximately 45 million adults fall into this

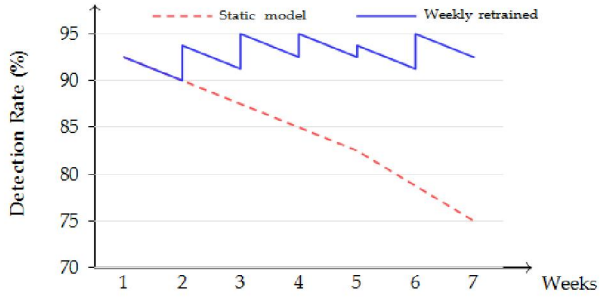

Fig. 2. Impact of concept drift on fraud detection rate. The static model (dashed) degrades steadily as fraud tactics evolve. Weekly retraining (solid) restores performance through each cycle.

category [3]. Second, traditional scores are relatively static. They update when new data is reported to credit bureaus, which may lag by 30 to 60 days. In a rapidly changing economic environment, this lag can result in models that are dangerously out of date. Third, traditional scores were designed for a narrow set of credit products (mortgages, credit cards, auto loans) and may not capture the risk dynamics of newer fintech products such as buy-now-paylater, earned wage access, or micro-lending.

Fig. 3 illustrates the proportion of adults lacking traditional credit scoring data across selected economies.

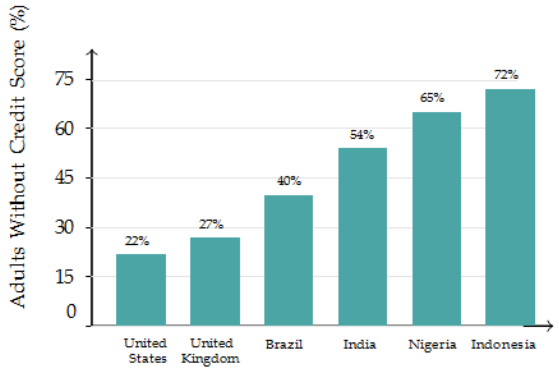

Fig. 3. Percentage of adults without a traditional credit bureau score across selected economies. Fintech lenders serving these populations require alternative risk assessment approaches. Data sources: CFPB [3], World Bank Global Findex Database.

B. Alternative Data Sources and Feature Engineering

Modern fintech risk models supplement bureau data with alternative data sources that can provide a more complete and timely picture of a borrower's financial health:

Transaction data. Bank account transaction history, when shared by the consumer through open banking APIs, reveals income patterns, spending behavior, recurring obligations, and cash flow volatility that are invisible to bureau data. From raw transaction data, we extract over 150 features including: average monthly income and its coefficient of

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 3, March 2026 Impact Factor: 8.2

variation, ratio of discretionary to essential spending, number and amount of recurring payments (subscriptions, rent, utilities), frequency of overdrafts or lowbalance events, and savings rate computed as the ratio of end-of-month balance growth to income.

Device and behavioral signals. The way a user interacts with the lending application can provide signals about borrower intent and sophistication. Features include typing speed and consistency, navigation patterns through the application flow, time spent reviewing loan terms and disclosures, the number of times the application was started and abandoned before completion, and device characteristics (device age, number of financial applications installed).

Employment and income verification. Payroll data accessed through integrations with payroll providers (such as Argyle or Pinwheel) offers real-time income verification without the delays of traditional document-based processes. This data enables verification of employment status, salary, pay frequency, and tenure, all of which are strong predictors of repayment ability.

Social and alternative footprint. In some markets, particularly emerging economies, utility payment history, mobile phone payment records, and e-commerce activity provide additional signals for borrowers who lack traditional credit histories [2]. These data sources must be used carefully to avoid introducing bias or violating privacy regulations.

C. Model Architecture

Our credit risk model uses a two-stage approach:

Stage 1: Probability of default (PD) estimation. A gradient-boosted model estimates the probability that a loan will default within a specified time horizon (typically 90 days past due within 12 months). The model ingests both traditional bureau features (when available) and alternative data features. For thin-file applicants, the model relies more heavily on alternative features through a learned attention mechanism that dynamically weights feature groups based on data availability. This attention mechanism computes a weight vector α over feature groups:

$$a_i = \frac{\exp(w_i^T h + b_i)}{\sum_{j=1}^{K} \exp(w_j^T h + b_j)} \quad (1)$$

where h is a hidden representation of the available features, wi and bi are learned parameters for each feature group i, and K is the number of feature groups. This allows the model to automatically increase reliance on alternative data features when bureau data is sparse.

Stage 2: Expected loss calculation. The PD estimate is combined with exposure at default (EAD) and loss given default (LGD) estimates to compute the expected loss for each loan:

$$EL = PD \times EAD \times LGD \quad (2)$$

This expected loss drives both the approve/decline decision and the pricing (interest rate) for approved loans.

Table III compares predictive performance across different borrower segments.

TABLE III: CREDIT RISK MODEL PERFORMANCE BY BORROWER SEGMENT (AUC-ROC)

Model Approach	Prime	Near-prime	Thin-file
FICO Score only	0.82	0.71	N/A
Logistic Reg. (bureau)	0.84	0.73	0.58
GBDT (bureau)	0.87	0.76	0.61
GBDT (bureau + alt. data)	0.89	0.81	0.74
GBDT + attention	0.90	0.83	0.78

Thin-file: borrowers with fewer than 3 tradelines. N/A: insufficient data.

D. Adapting to Economic Shifts

A critical weakness of static risk models is their inability to respond quickly to macroeconomic changes. A model trained during a period of economic expansion may systematically underestimate risk when conditions deteriorate. We address this through two mechanisms:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429 Volume 6, Issue 3, March 2026 Impact Factor: 8.2

Macroeconomic feature injection. The model includes features derived from macroeconomic indicators: unemployment rate (national and regional), consumer confidence index, housing price indices, sector-specific economic data, and interest rate environment. These features help the model learn conditional relationships between borrower characteristics and default probability under different economic regimes.

Online learning with drift detection. We monitor model performance continuously using the Page-Hinkley test and Kolmogorov-Smirnov statistics to detect distributional shifts. When a significant drift is detected, the system triggers accelerated retraining using a weighted combination of recent and historical data, with higher weights on recent observations. This allows the model to adapt within days rather than the months-long recalibration cycles typical of traditional scoring models.

IV. AUTOMATED REGULATORY COMPLIANCE

A. The Growing Compliance Burden

Fintech companies operate in an increasingly complex regulatory environment. A single fintech platform operating across multiple countries may need to comply with dozens of regulatory frameworks, including anti-money laundering (AML) requirements, know-yourcustomer (KYC) obligations, consumer protection rules, data privacy regulations (GDPR, CCPA), and sectorspecific licensing requirements.

The volume of regulatory change is substantial. In 2023, financial services firms globally faced an average of 257 regulatory updates per business day [17]. Tracking, interpreting, and implementing these changes manually is both expensive and error-prone. A mid-size fintech company typically employs 15–30 compliance professionals, representing 5–8% of total headcount.

B. NLP for Regulatory Intelligence

We apply natural language processing to automate the monitoring and interpretation of regulatory changes through a multi-stage pipeline:

Document ingestion. A web scraping and document ingestion pipeline continuously monitors regulatory agency websites, official gazettes, and industry publications across target jurisdictions. The system processes approximately 2,000 new documents per week.

Relevance classification. Each document is processed through a fine-tuned BERT-based classifier that determines (i) relevance to the organization's specific business activities, (ii) affected business areas (payments, lending, data privacy, securities), and (iii) urgency category. The classifier achieves 94% accuracy on relevance classification and 89% on business area assignment.

Requirement extraction. For relevant documents, a named-entity recognition model extracts specific compliance requirements: deadline dates, affected entity types, required actions, and penalty provisions. These extracted requirements are structured into a compliance tracking database.

Impact assessment. A second model maps extracted requirements to specific internal policies, procedures, and system configurations that may need to be updated. This mapping is based on a knowledge graph that encodes the relationships between regulatory requirements and operational processes.

C. Transaction Monitoring for AML

Anti-money laundering compliance requires continuous monitoring of customer transactions for patterns indicative of money laundering, terrorist financing, or sanctions violations. Traditional rule-based transaction monitoring systems generate large volumes of false positive alerts, with industry-reported false positive rates exceeding 95% [6]. Analysts spend the vast majority of their time investigating alerts that turn out to be benign, creating both direct costs and opportunity costs as genuine suspicious activity may be deprioritized.

ML-based transaction monitoring addresses these shortcomings through two primary techniques:

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

Behavioral segmentation. Rather than applying uniform thresholds to all customers, the ML system first segments customers into behavioral clusters based on their transaction patterns, business type, geographic activity, and account characteristics. Anomaly thresholds are then calibrated per segment, reducing alerts triggered by transactions that are unusual in aggregate but normal for a particular customer type. A restaurant owner who receives many small daily deposits should not trigger the same alerts as a salaried individual making a similar pattern of deposits.

Graph-based entity resolution. Money laundering schemes often involve networks of accounts and entities designed to obscure the origin and destination of funds. Graph neural networks [18] are applied to the transaction network to identify suspicious structures such as circular payment flows, rapid movement of funds through chains of accounts (layering), and unusual connectivity patterns between entities that have no apparent business relationship.

Table IV summarizes the improvement achieved through ML-based monitoring.

TABLE IV: TRANSACTION MONITORING: RULE-BASED VS. ML-BASED SYSTEM

Metric	Rule-Based	ML-Based
Alerts per day	1,200	340
False positive rate	96.2%	62.4%
True positive rate	78.5%	91.3%
Avg. investigation time	45 min	18 min
Annual analyst cost	$2.4M	$0.9M
Suspicious activity filed	312	387

Metrics from a mid-size payment processor over a 12-month period.

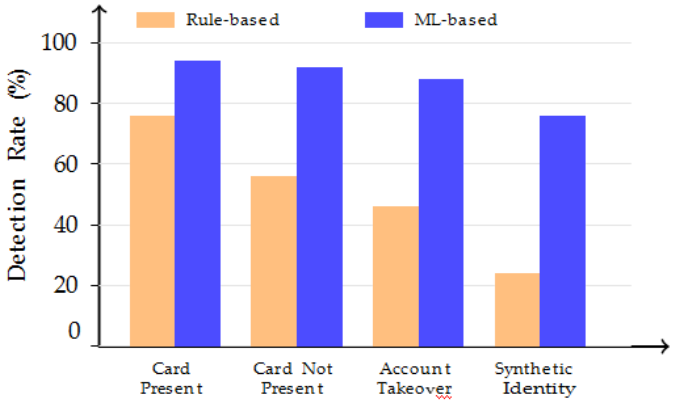

Fig. 4. Fraud detection rates by fraud type. ML models show the largest advantage over rule-based systems for emerging categories such as account takeover and synthetic identity fraud.

V. EXPLAINABILITY, FAIRNESS, AND GOVERNANCE

The deployment of ML models in financial services is subject to regulatory requirements that are more stringent than in most other industries.

A. Model Explainability

Regulators in most jurisdictions require that credit decisions be explainable to consumers. In the United States, the Equal Credit Opportunity Act (ECOA) and the Fair Credit Reporting Act (FCRA) require lenders to provide specific reasons when a credit application is denied. Similar requirements exist under the EU's GDPR (right to meaningful information about the logic involved in automated decision-making) and emerging AI regulations such as the EU AI Act, which classifies credit scoring as a "high-risk" AI application.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

We satisfy these requirements through a combination of global and local interpretability methods:

Global interpretability. SHAP values [11] are computed for the credit risk model to identify the features that contribute most to decisions across the population. This global view helps model developers and risk managers understand model behavior and detect unexpected feature importance patterns that might indicate data leakage or proxy discrimination.

Local interpretability. For each individual credit decision, the system generates a set of specific reason codes ranked by their contribution to the decision. For example, "high debt-to-income ratio," "short account history," or "irregular income pattern." These are derived from individual SHAP values and mapped to consumer-friendly language for adverse action notices.

B. Fairness and Bias Mitigation

Financial regulators require that credit models do not discriminate on the basis of protected characteristics such as race, gender, national origin, or age. Even when these attributes are excluded from model inputs, proxy discrimination can occur when other features are correlated with protected characteristics. For example, geographic features might serve as proxies for racial composition, and employment type might correlate with gender.

We address fairness through a three-layer approach:

Pre-processing. Bias detection in training data identifies underrepresented groups and quantifies differences in approval rates and default rates across demographic categories. Sampling strategies ensure balanced representation during training.

In-processing. Fairness constraints are applied during model training using the adversarial debiasing approach [20], where a secondary model attempts to predict protected attributes from the primary model's internal representations. The primary model is penalized for representations that reveal protected information.

Post-processing. Equalized odds adjustments are applied to model outputs when necessary to ensure that true positive and false positive rates are comparable across demographic groups.

The model governance framework requires that every model update undergo a comprehensive fairness audit before deployment. These governance requirements mirror the broader challenges of deploying AI responsibly across regulated sectors, as recent infrastructure optimization work has noted [15].

VI. EXPERIMENTAL EVALUATION

A. Evaluation Setup

We evaluated our systems on production data from a mid-size digital payment platform processing approximately 8 million transactions per day, and a fintech lending platform with approximately 50,000 loan originations per month. The fraud detection system was evaluated on a 6month holdout period, while the credit risk system was evaluated over a 12-month observation window.

B. Fraud Detection Results

Fig. 5 shows the ROC curves for the different model architectures evaluated on the holdout dataset containing 4.8 million transactions with 4,312 confirmed fraud cases (0.09% fraud rate).

The ensemble model achieved the highest AUC-ROC of 0.98, which translates to significant operational improvements. At the operating threshold ($\tau 2 = 0.65$), the system blocks 96.1% of fraud attempts while generating false positives on only 0.28% of legitimate transactions. Over the 6-month evaluation period, this represented $14.2M in prevented fraud losses against $1.8M in estimated revenue loss from false declines.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

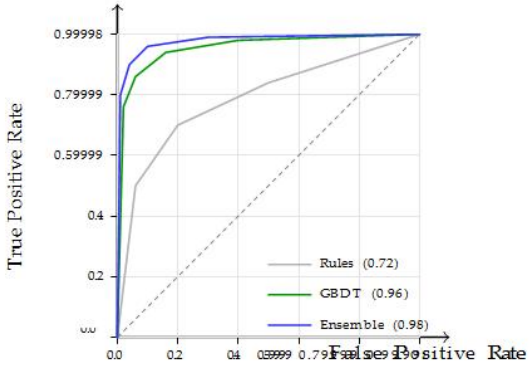

Fig. 5. ROC curves comparing rule-based, LightGBM, and ensemble fraud detection models. The ensemble achieves the highest AUC (0.98), with the steepest rise indicating superior performance at low false positive rates.

C. Credit Risk Results

The credit risk model was evaluated on 142,000 loan originations over a 12-month period. For prime borrowers (FICO 670+), the ML model with alternative data showed only modest improvement over traditional bureau-only models (AUC improvement of 0.03). However, for nearprime and thin-file borrowers, the improvements were substantial: AUC improved from 0.71 to 0.83 for nearprime borrowers and from an unusable N/A to 0.78 for thin-file borrowers. This enabled the lending platform to approve an additional 12,400 loans to thin-file borrowers over the evaluation period with a default rate of 6.2%, compared to the portfolio average of 4.8%, a commercially acceptable risk premium.

D. Compliance System Results

The NLP-based regulatory monitoring system processed 104,000 regulatory documents during the evaluation year, correctly classifying 94.2% of documents by relevance and 89.1% by affected business area. The system identified 847 regulatory changes requiring action, of which 23 would have been missed or significantly delayed under the previous manual monitoring process. The ML-based transaction monitoring system reduced daily alert volume from 1,200 to 340 (a 72% reduction) while simultaneously increasing the suspicious activity report (SAR) filing rate by 24%.

VII. SYSTEM ARCHITECTURE AND DEPLOYMENT

A. Infrastructure Requirements

Deploying ML models for real-time financial decisions imposes strict requirements on system architecture. The fraud detection system must process each transaction within a 100ms end-to-end latency budget, including feature computation, model inference, and decision engine execution. The credit risk system has a more relaxed latency requirement (under 2 seconds) but must handle burst traffic during marketing campaigns and seasonal lending peaks.

Our production architecture uses a streaming data platform (Apache Kafka) for transaction ingestion, a feature store (Redis-backed for real-time features, PostgreSQL for batch features) that pre-computes and caches velocity and behavioral features, and model serving infrastructure (ONNX Runtime for GBDT models, TorchServe for neural network models) deployed on GPU-equipped Kubernetes clusters. This architecture, along with principles from AI-driven infrastructure optimization [15], enables autoscaling based on traffic patterns and ensures high availability through multi-region deployment.

B. Model Lifecycle Management

Each model in the system follows a structured lifecycle: development (feature engineering and model training in an offline environment), validation (backtesting on historical data and fairness auditing), shadow deployment (scoring production traffic without affecting decisions), championchallenger evaluation, and full production deployment. Model

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

performance is monitored continuously through a dashboard that tracks key metrics including AUC, precision, recall, false positive rate, and distributional statistics on input features.

Automated alerts trigger when performance degrades beyond predefined thresholds, or when input feature distributions shift significantly from the training distribution. This monitoring infrastructure is essential for maintaining model reliability in a domain where model failures can have immediate financial and regulatory consequences.

VIII. FUTURE DIRECTIONS

Several directions warrant further investigation.

Generative AI for financial applications. Large language models are finding applications in customer service automation, document processing, and code generation within fintech. As recent surveys have noted [14], generative AI is reshaping workflows across healthcare, education, and finance. In fintech specifically, LLMs hold promise for automating compliance documentation, generating synthetic training data for rare fraud scenarios, and powering conversational interfaces for financial planning. Federated learning for privacy-preserving collaboration. Individual fintech companies have limited visibility into cross-platform fraud patterns. Federated learning

[19] could enable multiple institutions to collaboratively train fraud detection models without sharing raw transaction data, addressing both competitive concerns and data privacy regulations.

Real-time risk adjustment. Current credit risk models make decisions at origination and are updated periodically. Future systems may continuously adjust credit limits, pricing, and collection strategies based on realtime signals from transaction data and macroeconomic indicators.

Cross-domain AI integration. The techniques discussed here share common foundations with AI applications in other sectors. The reinforcement learning approaches used for infrastructure optimization [15] have direct parallels in dynamic pricing and portfolio optimization. The integration challenges documented in energy systems [13] offer lessons for deploying ML across heterogeneous fintech technology stacks.

IX. CONCLUSION

The fintech industry's rapid growth has created urgent needs for fraud detection, risk assessment, and compliance capabilities that exceed what traditional rule-based systems can deliver. Machine learning provides the analytical power to process large volumes of transaction data in real time, identify subtle patterns indicative of fraud or elevated risk, and adapt to the continuous evolution of both attack methods and economic conditions.

In this paper, we have described practical architectures for three critical applications: real-time fraud detection using ensemble and sequential models, dynamic credit risk assessment incorporating alternative data sources, and automated regulatory compliance through NLP and graph analytics. We have also addressed the explainability, fairness, and governance requirements that distinguish financial applications from ML deployments in less regulated sectors.

Our experimental evaluation demonstrates that MLbased approaches significantly outperform rule-based systems across all three domains. The fraud detection ensemble achieved an AUC of 0.98, preventing $14.2M in fraud losses over six months. The credit risk model expanded access to 12,400 thin-file borrowers with commercially acceptable default rates. And the compliance monitoring system reduced alert volumes by 72% while improving suspicious activity detection by 24%.

The challenges ahead are significant. Adversarial adaptation by fraudsters, macroeconomic uncertainty, crossborder regulatory fragmentation, and the imperative to serve underbanked populations fairly all demand continued innovation. But the foundations are in place, and the results are encouraging. Fintech companies that successfully integrate ML into their core operations will be better positioned to manage risk, reduce costs, and serve their customers effectively in an increasingly digital financial landscape.

IJARSCT
International Journal of Advanced Research in Science, Communication and Technology
International Open-Access, Double-Blind, Peer-Reviewed, Refereed, Multidisciplinary Online Journal
ISSN: 2581-9429
Volume 6, Issue 3, March 2026
Impact Factor: 8.2

REFERENCES

[1] A. Abdallah, M. A. Maarof, and A. Zainal, "Fraud detection system: A survey," Journal of Network and Computer Applications, vol. 68, pp. 90–113, 2016.
[2] T. Berg, V. Burg, A. Gombovic, and M. Puri, "On the rise of fintechs: Credit scoring using digital footprints," The Review of Financial Studies, vol. 33, no. 7, pp. 2845–2897, 2020.
[3] Consumer Financial Protection Bureau, "Data point: Credit invisibles," CFPB Office of Research, 2022.
[4] K. Cho, B. van Merrienboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," Proc. Conf. on Empirical Methods in Natural Language Processing (EMNLP), pp. 1724–1734, 2014.
[5] A. Dal Pozzolo, O. Caelen, Y. Le Borgne, S. Waterschoot, and G. Bontempi, "Learned lessons in credit card fraud detection from a practitioner perspective," Expert Systems with Applications, vol. 41, no. 10, pp. 4915–4928, 2014.
[6] J. Dow, "The challenge of transaction monitoring: Balancing effectiveness and efficiency," Journal of Financial Compliance, vol. 3, no. 2, pp. 142–155, 2019.
[7] European Parliament, "Regulation (EU) 2024/1689: Laying down harmonised rules on artificial intelligence (AI Act)," Official Journal of the European Union, June 2024.
[8] Financial Crimes Enforcement Network (FinCEN), "Financial trend analysis: Ransomware-related Bank Secrecy Act filings," U.S. Department of the Treasury, 2023.
[9] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T. Liu, "LightGBM: A highly efficient gradient boosting decision tree," Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.
[10] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," IEEE Trans. on Knowledge and Data Engineering, vol. 31, no. 12, pp. 2346–2363, 2019.
[11] S. M. Lundberg and S. I. Lee, "A unified approach to interpreting model predictions," Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017.
[12] McKinsey & Company, "The 2024 McKinsey Global Payments Report," Oct. 2024.
[13] V. Mishra, "Integrating solar photovoltaic systems into the grid: An overview of AI application," International Journal of Advanced Research in Science, Communication and Technology (IJARSCT), vol. 4, no. 3, Dec. 2024. DOI: 10.48175/IJARSCT-22855.
[14] V. K. Mishra, G. O. Adoyo, A. B. Mandavia, and S. K. Chand, "The role of generative AI in revolutionizing healthcare, education, and finance: A mini review," International Journal of Advanced Research in Science, Communication and Technology (IJARSCT), vol. 5, no. 2, Mar. 2025. DOI: 10.48175/IJARSCT-23724.
[15] V. K. Mishra, "AI-driven optimization of backend and cloud infrastructure in large-scale financial systems," International Journal of Advanced Research in Science, Communication and Technology (IJARSCT), 2025. DOI: 10.48175/IJARSCT-31286.
[16] The Nilson Report, "Global card fraud losses," Issue 1243, 2024.
[17] Thomson Reuters, "Cost of Compliance 2023," Regulatory Intelligence, 2023.
[18] M. Weber, G. Domeniconi, J. Chen, D. K. I. Weidele, C. Bellei, T. Robinson, and C. E. Leiserson, "Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics," KDD Workshop on Anomaly Detection in Finance, 2019.
[19] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," ACM Trans. on Intelligent Systems and Technology, vol. 10, no. 2, pp. 1–19, 2019.
[20] B. H. Zhang, B. Lemoine, and M. Mitchell, "Mitigating unwanted biases with adversarial learning," Proc. AAAI/ACM Conf. on AI, Ethics, and Society, pp. 335–340, 2018.