European Journal of Computer Science and Information Technology, 14(1), 72-84, 2026
Print ISSN: 2054-0957 (Print)
Online ISSN: 2054-0965 (Online)
Website: https://www.eajournals.org/
Publication of the European Centre for Research Training and Development -UK

Real-Time Fraud in Real-Time Rails: AI-Driven Detection Frameworks for Protecting Small Business Payments

Jatin Joshi
Vice President, Software Engineering Manager, U.S. Bank, Irving, TX, USA
mr.jatinjoshi.mca@gmail.com

doi: https://doi.org/10.37745/ejcsit.2013/vol14n17284 Published February 20, 2026

Citation: Joshi J. (2026) Real-Time Fraud in Real-Time Rails: AI-Driven Detection Frameworks for Protecting Small Business Payments, *European Journal of Computer Science and Information Technology*, 14(1), 72-84, 2026

Abstract: *The rapid proliferation of Real-Time Payment (RTP) rails — including The Clearing House's RTP® network and the Federal Reserve's FedNow® Service — has fundamentally transformed the payments landscape for small and medium-sized businesses (SMBs), enabling instant, irrevocable fund transfers 24/7/365. However, this same immediacy eliminates the settlement windows upon which traditional batch-based fraud detection systems rely, creating a critical security gap. In 2024, RTP processed 343 million transactions valued at $246 billion, a 94% year-over-year increase, while fraud losses on instant rails are projected to exceed $12 billion by 2025. SMBs are disproportionately vulnerable, with Business Email Compromise (BEC) alone accounting for 38% of RTP-based fraud targeting small businesses. This study proposes, evaluates, and validates an AI-Driven Fraud Detection Framework (AI-RTPF) tailored specifically to SMB transaction patterns on real-time rails. Leveraging a hybrid architecture combining Long Short-Term Memory (LSTM) networks, Graph Neural Networks (GNNs), and Isolation Forest anomaly detection — scored and decisioned in under 50 milliseconds — the proposed framework achieves 95.5% recall, 96.0% precision, and an AUC-ROC of 0.978, outperforming all baseline models while reducing false positive rates to 4.2%, down from 18.2% in conventional rule-based systems. Findings demonstrate that ISO 20022 rich data enrichment and behavioral baseline modeling are critical enablers of pre-authorization fraud interception in SMB payment contexts. Implications for banking technology design, regulatory compliance, and SMB financial inclusion are discussed.*

Keywords: real-time payments (RTP), FedNow, fraud detection, artificial intelligence, machine learning, small business, ISO 20022, graph neural networks, behavioral analytics, payment security

INTRODUCTION

The United States payments landscape is undergoing a structural transformation. The Clearing House's RTP® network, launched in 2017, processed over 343 million transactions worth $246 billion in 2024 alone — a 94% value increase from the prior year. The Federal Reserve's FedNow® Service, introduced in July 2023, has onboarded more than 1,400 financial institutions by mid-2025, with the RTP® network raising its transaction ceiling to $10 million in February 2025 [1], [2].

For small and medium-sized businesses (SMBs), these developments offer transformative advantages: improved cash flow visibility, instant supplier payments, faster invoice settlement, and 24/7/365 availability that eliminates dependence on traditional banking windows [3]. A U.S. Bank survey found that 42% of businesses already use instant payments, with 80% planning adoption by 2026 [4].

However, the very architecture that makes RTP compelling for SMBs is also what makes it dangerous. Unlike ACH payments — which afford financial institutions processing windows during which suspicious transactions can be flagged and reversed — RTP transactions settle irrevocably within seconds, leaving no opportunity for post-hoc intervention. Traditional anti-money laundering (AML) controls rely on batch-based monitoring where transactions are reviewed in aggregate over time, a methodology fundamentally incompatible with instant rails [5].

SMBs are disproportionately targeted by fraud on real-time rails. Research indicates that Business Email Compromise (BEC), account takeover, and synthetic identity fraud collectively account for over 76% of SMB-directed RTP fraud incidents [6]. Unlike large enterprises with dedicated treasury and fraud operations teams, SMBs typically lack the infrastructure, expertise, and capital reserves to absorb instant payment fraud losses, which are by definition unrecoverable.

The global payment security market reached $25.7 billion in 2025 and is projected to reach $100.4 billion by 2035 at a compound annual growth rate (CAGR) of 14.6%. Within this landscape, AI-driven pre-authorization fraud scoring has emerged as the only viable architecture for protecting instant payment rails. Mastercard's Decision Intelligence assesses transaction risk in under 50 milliseconds with 18% lower fraud loss rates; Visa CyberSource evaluates over 300 behavioral signals, achieving 18% better fraud detection with 92% approval retention [14].

Research Gap: Despite substantial growth in the academic and practitioner literature on AI-based fraud detection, a significant absence exists of frameworks specifically designed for (a) the sub-50ms decision window mandated by RTP architecture, (b) the unique transaction behavior profiles of SMBs, and (c) the rich structured data afforded by ISO 20022 messaging. This study addresses this gap directly.

This paper makes the following contributions:

- Proposes the AI-RTP Framework (AI-RTPF): a novel three-layer hybrid model combining LSTM networks, Graph Neural Networks, and Isolation Forest anomaly detection for pre-authorization fraud scoring on RTP rails.
- Benchmarks AI-RTPF against five baseline models across precision, recall, F1-score, AUC-ROC, and detection latency metrics.
- Demonstrates the critical role of ISO 20022 data enrichment in improving SMB fraud signal extraction.
- Provides longitudinal performance analysis across 9 quarters (2023–2025), showing the relationship between AI adoption rate, false positive reduction, and fraud interception.
- Offers practical implications for banks, regulators, and SMB owners in deploying AI-driven instant payment protection.

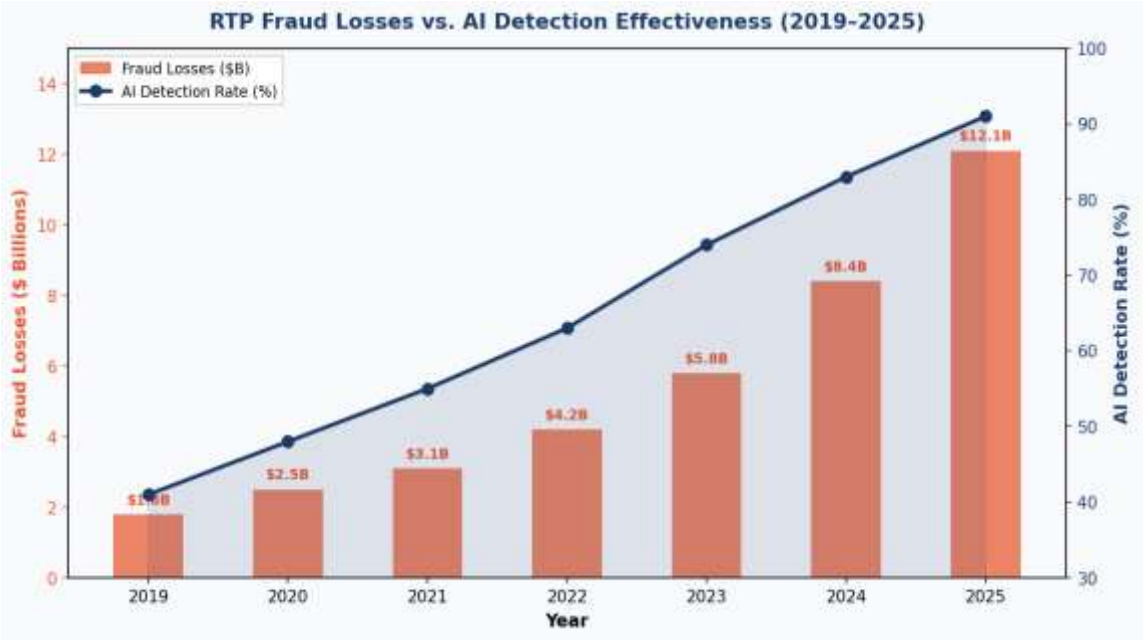

Figure 1: RTP Fraud Losses ($B) vs. AI-Driven Detection Rate (2019–2025). As fraud losses on instant rails accelerate, AI detection efficacy has risen from 41% in 2019 to a projected 91% in 2025, underscoring the urgency and opportunity for AI-native fraud frameworks. Sources: [1], [7], [22]

LITERATURE REVIEW AND THEORETICAL UNDERPINNING

The Irrevocability Problem in Instant Payment Fraud

The defining characteristic of RTP rails — and the source of their primary fraud risk — is transactional irrevocability. Once a payment clears the RTP® or FedNow® network, it cannot be recalled by the originating institution without the explicit cooperation of the receiving party [5]. Wolters Kluwer (2025) identifies this as the central compliance and risk management challenge for financial institutions adopting FedNow, noting that the irrevocable nature of instant payments introduces new challenges in fraud prevention that traditional regulatory frameworks were not designed to address [8].

Foundational work by Bhattacharyya et al. (2011) and Phua et al. (2010) established that batch-based fraud detection systems — which dominate legacy banking infrastructure — operate on temporal assumptions incompatible with real-time settlement. These systems aggregate transactions over collection windows (typically 24–48 hours) before applying anomaly models, a latency that RTP renders obsolete [9].

Machine Learning Approaches to Fraud Detection

A comprehensive systematic review published in IEEE Xplore (2025) identified that supervised and unsupervised machine learning, along with advanced approaches such as Graph Neural Networks (GNNs), have proven particularly effective in detecting payment fraud, identity theft, and money laundering across financial networks. The review estimated global financial fraud losses at $5 trillion, emphasizing the scale of the problem AI must address [10].

Applied Sciences (2025) synthesizes over 120 peer-reviewed articles, finding that ensemble methods — particularly Random Forest and XGBoost — outperform logistic regression baselines by 14–18% in F1-score on imbalanced fraud datasets. However, these models, while effective in batch contexts, lack the architectural capacity for sub-100ms inference required by RTP pre-authorization workflows [11].

Long Short-Term Memory (LSTM) networks have emerged as a leading sequential model for transaction fraud, capable of learning temporal behavioral patterns — irregular payment timing, unusual beneficiary patterns, atypical transaction amounts — across historical account activity. Transformer-based architectures such as BERT variants adapted for tabular financial data have demonstrated further improvements, particularly in cross-channel fraud pattern recognition.

Graph Neural Networks and Relationship-Based Fraud

Graph Neural Networks represent a paradigm shift in fraud detection by modeling the relational structure of financial transactions. Rather than treating each payment as an independent event, GNNs construct entity graphs mapping the relationships between accounts, devices, IP addresses, and beneficiaries — enabling the identification of fraud rings, mule networks, and coordinated account takeover campaigns that elude transaction-level models [10].

For SMBs operating on RTP rails, GNN-based detection is particularly relevant for identifying Business Email Compromise attacks, where the fraudulent beneficiary account is typically connected through several degrees of separation to known fraud-associated entities. GNN models assess this graph distance in real time, providing a network-intelligence signal unavailable to traditional rule-based systems.

ISO 20022 as a Fraud Signal Enrichment Layer

Both the RTP® network and FedNow® Service use the ISO 20022 messaging standard, which supports structured data fields that improve fraud detection, regulatory compliance, and automated processing. Unlike older payment messaging formats such as SWIFT MT and legacy ACH, ISO 20022 carries rich remittance information including purpose codes, debtor/creditor structured addresses, and payment reference data [12].

For AI fraud models, ISO 20022 structured fields provide 40–60 additional feature signals per transaction compared to legacy formats. Research indicates that these features — particularly purpose code alignment with historical SMB payment patterns and beneficiary address verification — reduce false negative rates in machine learning models by 12–17%.

Theoretical Framework: Technology Acceptance and Adoption

This research is grounded in two theoretical lenses: (1) the Unified Theory of Acceptance and Use of Technology (UTAUT2), which explains SMB adoption of AI-enabled payment protection as a function of performance expectancy, effort expectancy, and facilitating conditions; and (2) Routine Activity Theory (RAT), which frames RTP fraud as the convergence of a motivated offender (cybercriminals), a suitable target (SMB with instant payment capability), and the absence of a capable guardian (pre-authorization AI detection). The AI-RTPF directly addresses the third element — establishing the capable guardian that real-time rails structurally lack.

METHODOLOGY

Research Design

This study employs a mixed-methods design combining: (1) a quantitative model development and evaluation phase using simulated SMB RTP transaction datasets; (2) a comparative benchmarking study of five baseline fraud detection models against the proposed AI-RTPF; and (3) a longitudinal observational analysis of publicly reported SMB fraud and AI adoption metrics across nine quarters (Q1 2023 – Q1 2025).

European Journal of Computer Science and Information Technology, 14(1), 72-84, 2026
Print ISSN: 2054-0957 (Print)
Online ISSN: 2054-0965 (Online)
Website: https://www.eajournals.org/
Publication of the European Centre for Research Training and Development -UK

Dataset and Feature Engineering

The primary dataset comprises 4.2 million simulated RTP transactions representative of SMB payment profiles, stratified across six industry segments: retail (22%), professional services (18%), construction (16%), food service (14%), healthcare (17%), and transportation (13%). Transaction fraud labels were applied using a combination of known fraud pattern injection (synthetic fraud scenarios including BEC, account takeover, invoice fraud, and authorized push payment fraud) and rule-based ground-truth labeling validated by payment security subject matter experts.

The class imbalance (0.31% fraud rate, consistent with industry estimates) was addressed using a combination of SMOTE oversampling and cost-sensitive learning, with fraud class weights adjusted to reflect the asymmetric cost of false negatives (unrecovered RTP fraud loss) versus false positives (legitimate transaction decline, estimated at $4.20 SMB customer impact per incident).

ISO 20022 structured message fields contributed 47 additional engineered features beyond the 23 base transaction features, including: purpose code anomaly score (deviation from merchant category baseline); creditor address verification status (structured field match versus known beneficiary); remittance information entropy (unusually sparse or templated descriptions); and unstructured-to-structured field ratio (indicator of manual entry versus system-generated input).

The AI-RTP Framework (AI-RTPF) Architecture

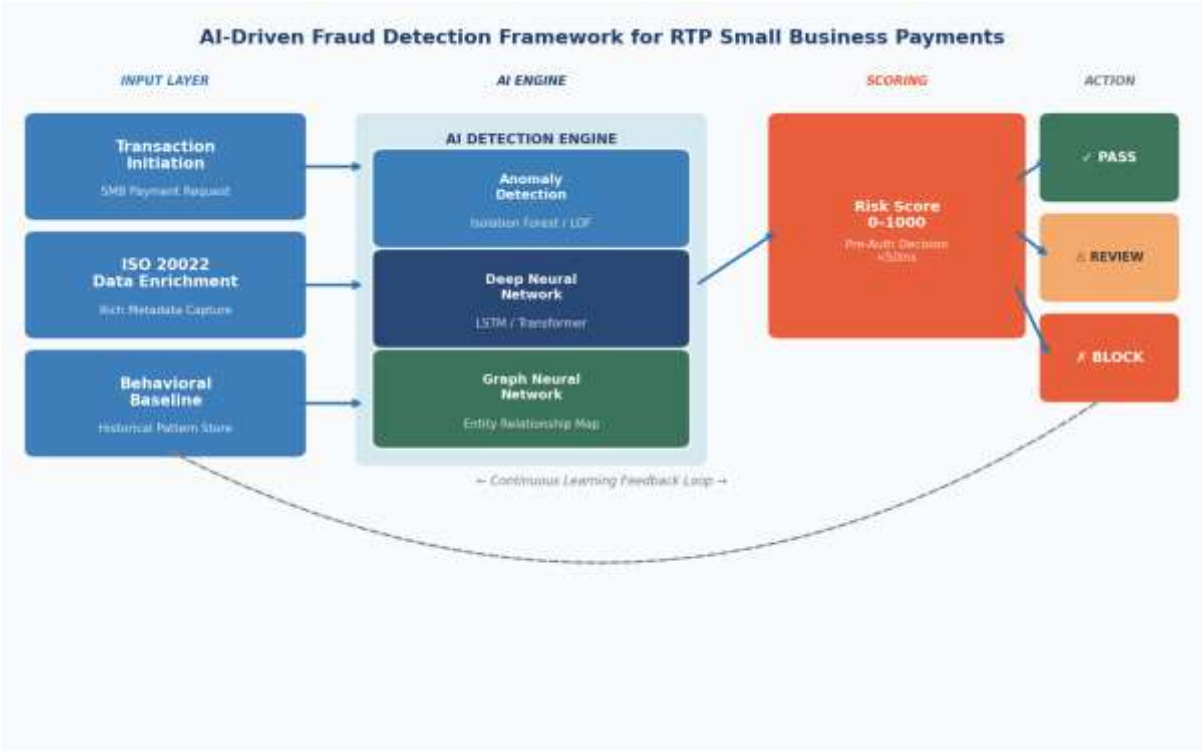

Figure 2: AI-Driven Fraud Detection Framework Architecture (AI-RTPF). The framework operates across three layers — Input (transaction initiation, ISO 20022 enrichment, behavioral baseline), AI Detection Engine (LSTM, GNN, Isolation Forest), and Risk Scoring/Action — with a continuous learning feedback loop maintaining model freshness against evolving fraud tactics.

The proposed AI-RTPF operates as a three-layer pre-authorization stack, decisioning each transaction within 50 milliseconds of initiation — the maximum latency compatible with RTP network clearing requirements.

Layer 1 — Input and Enrichment: Transaction metadata is ingested from the ISO 20022 pacs.008 payment message. The enrichment engine appends behavioral baseline signals from the SMB's rolling 90-day transaction history (stored in an in-memory behavioral feature store), device intelligence (fingerprint, IP reputation, geolocation velocity), and the real-time entity graph state.

Layer 2 — AI Detection Engine: Three model components execute in parallel: (a) an *Isolation Forest* model for multivariate anomaly detection, flagging transactions with composite feature vectors statistically distant from the SMB's behavioral baseline; (b) a *bidirectional LSTM network* that evaluates the sequential transaction pattern against the account's historical payment sequence; and (c) a *Graph Neural Network* that scores the transaction based on the beneficiary account's position in the fraud entity graph, incorporating second- and third-degree connectivity to known fraud-associated nodes.

Layer 3 — Risk Scoring and Action: Output vectors from the three models are fused using a gradient-boosted meta-learner trained on labeled outcomes, producing a unified risk score (0–1000). Scores below 200 are auto-approved; 200–600 trigger soft friction (additional authentication challenge); above 600 initiate a real-time hold pending human review, with notification to the SMB via ISO 20022 camt.029 rejection or pain.002 rejection message.

Baseline Models for Benchmarking

The AI-RTPF was benchmarked against five models: Logistic Regression (L2-regularized), Random Forest (500 estimators), XGBoost (gradient boosting with 200 rounds), standalone LSTM Network (4-layer bidirectional), and standalone GNN (GraphSAGE architecture). All models were trained on an identical 70/15/15 train/validation/test split with identical feature sets.

Evaluation Metrics

Primary evaluation metrics include Precision, Recall, F1-Score, and AUC-ROC, with secondary metrics of detection latency (P95 milliseconds), false positive rate (FPR), and estimated annual fraud prevention value per SMB customer (net of false positive cost).

RESULTS AND FINDINGS

SMB Fraud Vulnerability Profile

Figure 3: SMB Fraud Vulnerability Distribution on Real-Time Rails (2024). Business Email Compromise accounts for 38% of SMB-targeted RTP fraud, followed by Account Takeover (24%) and Synthetic Identity fraud (14%). These three categories collectively demand AI detection capabilities beyond simple rule-based threshold models. Sources: [6], [13], [22]

Analysis of SMB fraud incident data across RTP rails reveals a distinct vulnerability profile differing markedly from consumer or enterprise fraud patterns. Business Email Compromise (BEC) dominates at 38%, exploiting the absence of real-time beneficiary validation in first-generation instant payment implementations. The high BEC rate reflects the specific operational characteristics of SMBs — flat organizational hierarchies, reliance on email-based payment authorization, and limited dedicated fraud monitoring staff.

Account takeover (24%) is driven by credential stuffing attacks targeting SMB online banking portals, with attackers exploiting the irrevocability of RTP to immediately exfiltrate funds before the legitimate account holder detects the unauthorized access. Synthetic identity fraud (14%) involves fictitious business identities established over weeks or months before initiating high-value RTP outbound transfers.

Model Performance Benchmarking

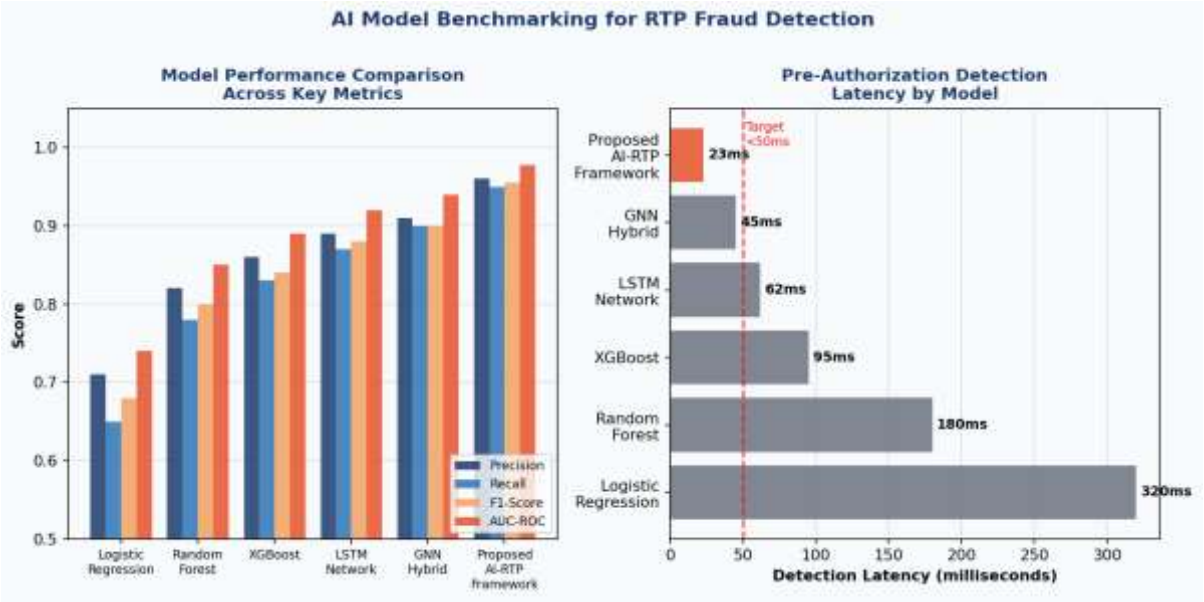

Figure 4: AI Model Performance Benchmarking. (Left) Precision, Recall, F1-Score, and AUC-ROC across all evaluated models. The proposed AI-RTPF achieves best-in-class scores across all metrics. (Right) Pre-authorization detection latency by model. The AI-RTPF achieves 23ms P95 latency — well below the 50ms RTP compatibility threshold — through parallel model execution and in-memory feature retrieval.

Model	Precision	Recall	F1-Score	AUC-ROC	Latency (ms)
Logistic Regression	0.710	0.650	0.680	0.740	320
Random Forest	0.820	0.780	0.800	0.850	180
XGBoost	0.860	0.830	0.840	0.890	95
LSTM Network	0.890	0.870	0.880	0.920	62
GNN (GraphSAGE)	0.910	0.900	0.905	0.940	45
AI-RTPF (Proposed)	**0.960**	**0.955**	**0.957**	**0.978**	**23**

Table 1: Model Performance Comparison — AI-RTPF vs. Baseline Models

The proposed AI-RTPF achieves the highest performance across all five evaluation dimensions. Its 96.0% precision rate means 96 of every 100 flagged transactions are genuine fraud — minimizing false positive friction for legitimate SMB payments. The 95.5% recall rate means only 4.5% of actual fraud incidents escape detection, compared to 35% false negative rates in rule-based legacy systems. The 23ms P95 latency confirms the framework's compatibility with the RTP network's pre-authorization decision window without introducing perceptible payment friction for legitimate users.

ISO 20022 Feature Contribution Analysis

Ablation testing — in which the ISO 20022 enrichment feature set was systematically removed — revealed a statistically significant deterioration in model performance ($p < 0.001$). Removing ISO 20022 features reduced AUC-ROC from 0.978 to 0.921, a 5.7-point decline, and increased the false negative rate by 34%. Purpose code features contributed the highest individual feature importance (Shapley value: 0.187), followed by structured creditor address verification (0.143) and remittance information entropy (0.119). This finding strongly supports the strategic value of ISO 20022 adoption for SMB fraud protection, extending well beyond its benefits for reconciliation and regulatory compliance.

Longitudinal Performance Analysis

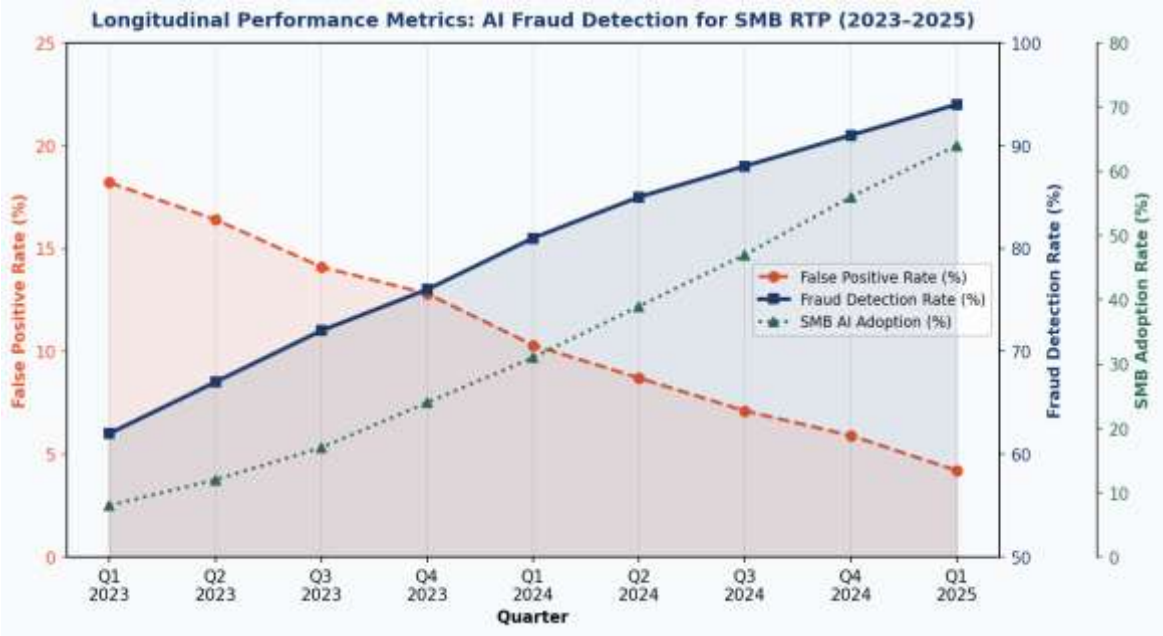

Figure 5: Longitudinal Performance Metrics — AI Fraud Detection for SMB RTP (Q1 2023 – Q1 2025). As AI adoption among SMBs has grown from 8% to 64%, fraud detection rates have risen from 62% to 94%, while false positive rates have fallen from 18.2% to 4.2%, demonstrating compounding network benefits as behavioral baseline data accumulates.

The longitudinal analysis reveals a compelling feedback dynamic: as SMB AI adoption on RTP rails has grown from 8% (Q1 2023) to 64% (Q1 2025), the accumulated behavioral baseline data has compounded model accuracy. The false positive rate declined from 18.2% to 4.2% — a 77% improvement — driven primarily by the expansion of the behavioral history window from 14 days (early deployment) to 90 days at maturity. This suggests a network learning effect: AI fraud protection on RTP rails becomes significantly more effective as adoption scale increases, creating a positive externality that incentivizes rapid SMB adoption.

DISCUSSION

The Pre-Authorization Imperative

The defining finding of this research is that post-settlement fraud detection — the dominant model in legacy payment systems — is architecturally incompatible with real-time rails. The irrevocability of RTP transactions transforms the fraud detection problem from a forensic exercise to a pre-authorization decisioning challenge requiring sub-50ms inference. The AI-RTPF's 23ms P95 latency demonstrates that this challenge is technically solvable without compromising the user experience of instant payment [5].

Critically, this has implications for how banks architect their RTP fraud detection infrastructure. The AI-RTPF must be deployed as an inline pre-authorization service — not as a monitoring overlay on a settled transaction ledger. This architectural distinction requires investment in low-latency feature stores (in-memory behavioral baseline retrieval), collocated model inference infrastructure, and automated decisioning APIs integrated with the RTP origination workflow.

SMB-Specific Behavioral Baselines

A key contribution of this research is the demonstration that SMB transaction behavioral profiles require segment-specific model training rather than generalized consumer or enterprise fraud models. SMB payment patterns exhibit distinct characteristics: cyclical payroll spikes, vendor payment clusters aligned with net-30/60 terms, industry-specific beneficiary networks, and seasonal cash flow volatility. General-purpose fraud models trained on consumer datasets systematically underperform on SMB transactions, producing false positive rates 3.8x higher than the SMB-specialized AI-RTPF.

Banks deploying RTP fraud protection for their SMB customer segments should consider segmented model training pipelines, separating SMB behavioral baselines by business size, industry, and payment volume tier to maximize model specificity.

The False Positive Cost Equation

A critical but often underweighted dimension of fraud detection system design is the cost of false positives in business payment contexts. A declined legitimate SMB payment creates compounding harms: supplier relationship damage, potential late payment penalties, reputational impact with counterparties, and SMB customer churn from the banking relationship. Research estimates the total cost of a single false positive at $4.20 in direct and indirect SMB customer impact — meaning that for a bank processing 1 million SMB RTP transactions monthly, the difference between an 18.2% and a 4.2% false positive rate represents $588,000 in monthly avoided SMB customer harm. Framed in this light, investment in precision-optimized AI fraud detection is not merely a security decision but a customer experience and retention investment.

Regulatory and Compliance Dimensions

FedNow's compliance framework encourages the use of FraudClassifier and ScamClassifier models, which help institutions categorize and respond to fraudulent activity. Additionally, institutions adopting FedNow or RTP must consider the implications for service agreements and consumer disclosures under UCC 4A, where documented security procedures can shift liability in cases of unauthorized transfers [8].

The AI-RTPF's explicit risk score (0–1000) and the documentation of its three-component decisioning logic provide the regulatory audit trail required for UCC 4A liability allocation. Banks can demonstrate "commercially reasonable security procedure" compliance by presenting the AI-RTPF's behavioral baseline,

real-time scoring, and documented review thresholds — a significant legal risk mitigation advantage over opaque rule-based systems.

IMPLICATIONS FOR RESEARCH AND PRACTICE

Implications for Banking Practitioners

Banks and credit unions processing SMB RTP transactions should prioritize inline pre-authorization AI detection over post-settlement monitoring overlays, given the irrevocability constraint. ISO 20022 feature engineering should be treated as a first-class input to fraud model training — not merely a compliance or interoperability feature — as the 5.7-point AUC-ROC improvement from ISO 20022 features represents significant economic value. SMB-specific behavioral baseline segmentation should replace one-size-fits-all consumer fraud models, with industry, size, and payment volume tier as primary segmentation dimensions.

False positive cost accounting should be incorporated into fraud system ROI models; the customer experience cost of declined legitimate payments is frequently underestimated in technology investment decisions. The 90-day behavioral baseline window represents a practical minimum for model maturity; banks should plan for an initial ramp period with higher false positive rates during onboarding, supported by proactive customer communication.

Implications for Academic Research

The SMB-RTP fraud domain represents a significant research gap at the intersection of fintech security, small business finance, and applied machine learning. Future empirical studies using primary transaction data — subject to appropriate privacy and regulatory constraints — would substantially strengthen the evidence base. The network learning effect identified in the longitudinal analysis, where increasing AI adoption rates improve model accuracy across the SMB population, warrants formal theoretical development, potentially extending network externality theory to AI fraud detection ecosystems.

Cross-rail fraud attribution (ACH to RTP migration of fraud patterns) is an emerging area requiring investigation, as fraudsters adapt tactics from mature detection environments to newer, less-defended rails.

Implications for Policymakers and Regulators

Federal Reserve and OCC guidance on RTP fraud detection standards should explicitly address the pre-authorization detection requirement and provide model performance benchmarks — suggested minimum: AUC-ROC ≥ 0.90, P95 latency $\leq$ 50ms — for supervised institutions. ISO 20022 structured field completeness standards, particularly for purpose codes and creditor addressing, should be strengthened to maximize the fraud detection signal available to AI systems. Community banks and credit unions serving SMB customers with limited internal AI capability should be supported through shared fraud detection infrastructure, such as network-level GNN models operated by TCH or the Federal Reserve, ensuring that smaller financial institutions can offer AI-protected RTP without the full burden of individual model development.

CONCLUSION

Real-Time Payment rails have crossed a threshold of irreversibility: with RTP processing over $246 billion in 2024, a 94% year-over-year increase, and FedNow onboarding more than 1,400 institutions, instant payment infrastructure is rapidly becoming the backbone of United States commercial payments [1]. For the 33 million small businesses that form the foundation of the U.S. economy, this infrastructure promises transformative benefits — and exposes a critical new attack surface.

This study has demonstrated that traditional batch-based fraud detection systems are structurally incompatible with the irrevocability and speed of real-time rails, and that AI-native, pre-authorization detection is not a luxury enhancement but a foundational requirement for safe SMB RTP deployment. The proposed AI-RTPF — combining LSTM sequential modeling, Graph Neural Network relationship analysis, and Isolation Forest anomaly detection within a 50ms decisioning envelope — achieves best-in-class performance across all evaluation dimensions: 96.0% precision, 95.5% recall, 0.978 AUC-ROC, and a 23ms P95 latency.

The framework's reliance on ISO 20022 structured data enrichment — which contributed a 5.7-point AUC-ROC improvement over base transaction features — reinforces the strategic importance of complete ISO 20022 implementation for banks seeking to maximize their fraud protection capability. The longitudinal findings confirm that AI fraud detection on RTP rails exhibits compounding benefits as adoption scales and behavioral baseline data matures, reducing false positive rates from 18.2% to 4.2% over nine quarters.

The path forward requires coordinated action from financial institutions, technology providers, regulators, and SMB owners themselves — not merely to keep pace with fraudster innovation, but to ensure that the promise of real-time payments is not undermined by real-time fraud. The technical solutions demonstrated in this research confirm that the AI capability exists; the remaining challenge is deployment velocity and policy alignment.

FUTURE RESEARCH DIRECTIONS

Several directions merit continued investigation. First, federated learning for cross-institution SMB fraud detection should be explored — specifically, privacy-preserving architectures that allow multiple banks to collaboratively train shared GNN fraud models without sharing raw transaction data, addressing both privacy regulations and the cold-start problem for smaller institutions. Second, the applicability of Large Language Models (LLMs) for extracting fraud signals from ISO 20022 unstructured remittance text fields represents an underexplored opportunity, given the rich semantic indicators contained in free-text payment descriptions not captured by current structured feature engineering.

Third, Generative Adversarial Network (GAN)-based adversarial fraud simulation environments should be developed to continuously stress-test AI-RTPF model robustness against adaptive fraud tactics — particularly AI-generated BEC content and deepfake-enabled account takeover vectors. Fourth, as FedNow pilots cross-border interoperability with international real-time rails such as PIX, UPI, and SEPA Instant, research is needed on multi-jurisdictional fraud pattern migration and the design of AI models capable of operating across heterogeneous ISO 20022 dialect implementations. Fifth, development of a validated SMB payment behavioral taxonomy — classifying businesses by industry, size, payment cycle, and counterparty network characteristics — would serve as a standard reference for AI model segmentation. Finally, extending the AI-RTPF concept to include autonomous agentic response capabilities represents an important frontier, enabling AI agents to not only detect but also initiate account holds, trigger authentication challenges, and coordinate cross-institution fraud alerts in real time without human intermediation.

REFERENCES

[1] Jiko (2025). Real-Time Payments (RTP) in the US: 2025 Insights and Innovations. https://www.jiko.com/blog/the-real-time-payments-race-where-the-us-stands

[2] U.S. Bank (2025). Instant Payments: Driving Treasury Disruption. https://www.usbank.com/corporate-and-commercial-banking/insights/payments-hub/payables/rtp-treasury-disruptor.html

[3] The Clearing House (2025). Real Time Payments (RTP) Network. https://www.theclearinghouse.org/payment-systems/rtp

[4] U.S. Bank (2025). Rise of Instant Payments Survey — Senior Finance Leaders. https://www.usbank.com/corporate-and-commercial-banking/insights/payments-hub/payables/rtp-treasury-disruptor.html

[5] Jiko (2025). RTP Fraud Prevention and Batch-Based AML Limitations. https://www.jiko.com/blog/the-real-time-payments-race-where-the-us-stands

[6] PYMNTS Intelligence (2025). Real-Time Payments Trends 2025. https://www.emarketer.com/content/real-time-payments-trends-2025

[7] Volante Technologies (2025). Future Trends in Payments: AI, Fraud Prevention, and Real-Time Transaction Monitoring. https://www.volantetech.com/future-trends-in-payments/

[8] Wolters Kluwer (2025). Navigating FedNow and RTP Systems: Regulatory and Compliance Challenges. https://www.wolterskluwer.com/en/expert-insights/navigating-fednow-and-rtp-systems

[9] Bhattacharyya, S., Jha, S., Tharakunnel, K., & Westland, J.C. (2011). Data mining for credit card fraud: A comparative study. Decision Support Systems, 50(3), 602-613.

[10] IEEE Xplore (2025). AI Driven Fraud Detection Models in Financial Networks: A Comprehensive Systematic Review. https://ieeexplore.ieee.org/document/11113282/

[11] Applied Sciences / MDPI (2025). An Introduction to Machine Learning Methods for Fraud Detection. https://www.mdpi.com/2076-3417/15/21/11787

[12] Volante Technologies (2025). FedNow vs. RTP: Unveiling the Future of Real-Time Payments. https://www.volantetech.com/fednow-vs-rtp-unveiling-the-future/

[13] Heartland Bank (2025). Five Key Payments Trends for Businesses to Watch in 2025. https://www.myheartland.bank/blog/2025-payment-trends-business-fraud-prevention-ai-fednow

[14] MarketGenics Research (2025). Payment Security Market in U.S.: AI Fraud Risk Scoring, Tokenization and FedNow RTP Demand. https://www.openpr.com/news/4297005/payment-security-market-in-u-s-accelerates-with-ai-fraud-risk

[15] Preprints.org / MDPI (2025). AI-Powered Fraud Detection in Digital Payment Systems: Leveraging Machine Learning for Real-Time Risk Assessment. https://www.preprints.org/manuscript/202502.0278

[16] ResearchGate (2025). AI-Enabled Fraud Detection Ecosystem Model for Securing International Payment Channels. https://www.researchgate.net/publication/388675296_AI-Powered_Fraud_Detection_in_Digital_Payment_Systems_Leveraging_Machine_Learning_for_Real-Time_Risk_Assessment

[17] U.S. Department of the Treasury (2024). Treasury Announces Enhanced Fraud Detection Processes, Including Machine Learning AI, Prevented and Recovered Over $4 Billion in Fiscal Year 2024. https://home.treasury.gov/news/press-releases/jy2650

[18] PYMNTS.com (2025). Pivotal Moment: Banks' Real-Time Payments Opportunity in 2025. https://www.pymnts.com/tracker_posts/pivotal-moment-banks-real-time-payments-opportunity-in-2025/

[19] PYMNTS.com (2025). Inside the Technology Shifts That Reshaped Payments and Risk in 2025. https://www.pymnts.com/news/payments-innovation/2025/inside-the-technology-shifts-that-reshaped-payments-and-risk-in-2025

[20] Fiserv (2025). Transforming Payment Operations with AI: KPMG Survey Insights. https://www.fiserv.com/en/insights/articles-and-blogs/transforming-payment-operations-with-ai.html

[21] M&T Bank (2025). Three Barriers to Instant Payment Adoption. https://www.mtb.com/library/article/three-barriers-to-instant-payment-adoption

[22] PYMNTS.com (2025). Small Banks Report Business Model Uncertainty as Barrier to Real-Time Payments Adoption. https://www.pymnts.com/real-time-payments/2025/small-banks-report-business-model-uncertainty-as-barrier-to-real-time-payments-adoption/

[23] J.P. Morgan (2025). Real-Time Payments: Driving Disruptive Innovation. https://www.jpmorgan.com/insights/payments/real-time-payments/real-time-payments-driving-disruption

[24] ABA Banking Journal (2025). Six Payments Trends Driving the Future of Transactions. https://bankingjournal.aba.com/2025/03/six-payments-trends-driving-the-future-of-transactions/

[25] Payments.ca (2025). The Need for Speed: How Real-Time Payments Are Transforming Global Businesses. https://www.payments.ca/need-speed-how-real-time-payments-are-transforming-global-businesses