Journal of Science, Technology and Social Transformation Volume 2, Issue 02, 2026

Artificial Intelligence and Automation in Modern Payment Gateways:A Comprehensive Review

Pramod Kumar

Independent Researcher

Abstract

The rapid integration of artificial intelligence (AI) and automation technologies into financial payment infrastructures has fundamentally transformed the architecture, security posture, and operational efficiency of modern payment gateways. This comprehensive review systematically examines the current state of AI adoption across the global payment processing landscape, encompassing machine learning-based fraud detection, intelligent transaction routing, natural language processing for customer dispute resolution, robotic process automation in compliance workflows, and emerging biometric authentication mechanisms. Drawing upon peer-reviewed literature, industry white papers, and empirical deployment data published between 2018 and 2026, this article identifies key technological trajectories, evaluates comparative performance metrics among leading gateway providers—including Stripe, Adyen, CyberSource, PayPal, and Razorpay—and critically assesses persistent challenges related to algorithmic bias, explainability, adversarial robustness, and regulatory compliance under frameworks such as PCI DSS v4.0 and GDPR. The review further explores the convergence of AI with blockchain-based settlement, open banking APIs, and large language model (LLM) applications in payment orchestration. Findings indicate that AI-driven gateways achieve fraud detection accuracy exceeding 99.2%, false-positive reductions of up to 67%, and authorisation rate improvements of 3–8 percentage points over rule-based systems. The article concludes with a forward-looking agenda identifying federated learning, quantum-resistant cryptography, and real-time explainable AI as the most consequential research frontiers in this domain.

Keywords: Artificial intelligence; payment gateways; fraud detection; machine learning; transaction routing; robotic process automation; fintech; deep learning; PCI DSS; open banking DOI: 10.64235/ff0t5h25

Introduction

The global digital payments ecosystem processed an estimated USD 14.8 trillion in transaction value during 2025, with compound annual growth rates exceeding 12% across emerging economies and 7% in mature markets [1].

Payment gateways serve as the critical nexus between merchants, acquiring banks, card networks, and issuing institutions. Historically, these systems prioritised throughput and reliability, with security enforced through deterministic rule engines—velocity checks, block lists, and BIN-range restrictions. While effective against known attack vectors, such approaches proved brittle against novel fraud patterns, generated high false-positive rates that degraded customer experience, and demanded substantial manual intervention from risk analysts.

The emergence of machine learning (ML) in the mid-2000s, followed by deep learning breakthroughs after 2012, provided financial institutions with tools capable of ingesting millions of transaction attributes simultaneously, detecting non-linear relationships invisible to human analysts, and adapting continuously to evolving threat landscapes.

Despite rapid industry adoption, rigorous academic synthesis of AI applications within payment gateways remains fragmented. Existing reviews tend to focus narrowly on fraud detection [2,3], omitting the broader automation stack encompassing compliance, customer service, treasury operations, and infrastructure orchestration. Furthermore, critical concerns regarding algorithmic fairness, model transparency, and regulatory conformance have received insufficient attention in the practitioner literature.

This review addresses these gaps by providing a comprehensive, multi-dimensional analysis of AI and automation technologies across the payment gateway lifecycle. Our objectives are fourfold: (i) to map the current

Corresponding Author: Pramod Kumar Independent Researcher Email id:Pramod.kumar987789@gmail.com

How to cite this article: Kumar P. (2026). Artificial Intelligence and Automation in Modern Payment Gateways:A Comprehensive Review *Journal of Science, Technology and Social Transformation* 2(2), 22-28.

Source of support: Nil

Conflict of interest: None

technological landscape of AI deployment in payment systems; (ii) to compare the AI capabilities of leading global gateway providers; (iii) to critically evaluate persistent technical and regulatory challenges; and (iv) to articulate a research agenda for the near-term future. The remainder of this article is structured as follows: Section 2 presents the methodology; Section 3 reviews core AI technologies and their applications; Section 4 examines major gateway implementations; Section 5 addresses challenges and limitations; Section 6 surveys emerging trends; and Section 7 concludes with recommendations.

Methodology

This review follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) framework, adapted for technology-domain reviews. Literature was retrieved from IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, Google Scholar, and SSRN using Boolean search strings combining terms from three semantic clusters: (a) "payment gateway" OR "payment processor" OR "merchant acquirer"; (b) "artificial intelligence" OR "machine learning" OR "deep learning" OR "natural language processing" OR "robotic process automation"; and (c) "fraud detection" OR "transaction routing" OR "authentication" OR "compliance automation".

Search results were restricted to publications between January 2018 and March 2026, yielding 1,247 initial records. After de-duplication (n = 183), title and abstract screening (n = 612 excluded), and full-text eligibility assessment against inclusion criteria—empirical studies, peer-reviewed reviews, or high-quality grey literature with verifiable data—214 sources were retained for synthesis. An additional 38 industry reports from Nilson Report, McKinsey Global Payments, Juniper Research, and gateway vendors were incorporated to supplement peer-reviewed findings with deployment metrics.

Core AI Technologies in Payment Gateways

Machine Learning for Fraud Detection

Fraud detection represents the most mature and well-documented application of ML in payment gateways. The fundamental challenge is one of extreme class imbalance: fraudulent transactions typically constitute 0.01%–0.5% of transaction volume, yet their financial impact disproportionately exceeds their frequency [4]. Early ML approaches applied logistic regression and decision tree ensembles, achieving modest improvements over rule-based baselines. The watershed development was the application of gradient boosting machines—particularly XGBoost and LightGBM—to high-dimensional transaction feature spaces, enabling real-time scoring at sub-50 millisecond latencies with area-under-curve (AUC) values consistently above 0.97 [5].

Feature engineering constitutes the critical determinant of model performance. Contemporary fraud models ingest 400–1,200 features per transaction, spanning four principal categories: (i) transaction attributes (amount, currency, merchant category code, time-of-day); (ii) device and network signals (IP geolocation, device fingerprint, browser entropy, VPN detection); (iii) behavioural sequences (typing cadence, mouse movement velocity, session navigation patterns); and (iv) network graph features (shared account attributes, IP co-occurrence, device reuse across flagged accounts) [6].

Deep learning models, particularly recurrent neural networks (RNN) and their variants—Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU)—have demonstrated superior performance in capturing temporal patterns within transaction sequences. A landmark study by Jurgovsky et al. [7] demonstrated that LSTM models outperformed random forests by 19% on precision-recall AUC for credit card fraud when sequential transaction context was incorporated. [8].

Graph Neural Networks for Network-Level Fraud

A significant limitation of transaction-level ML models is their inability to detect coordinated fraud rings—organised groups that distribute fraudulent activity across multiple accounts, devices, and merchants to evade per-account velocity thresholds. Graph Neural Networks (GNNs) address this by modelling the payment ecosystem as a heterogeneous graph wherein nodes represent entities (accounts, devices, merchants, IPs) and edges represent interactions [9].

Implementation at scale presents substantial engineering challenges, as payment graphs can contain billions of nodes and hundreds of billions of edges. Recent deployments at Alipay and PayPal have demonstrated that real-time GNN inference is achievable through graph sampling, feature caching, and approximate neighbour aggregation strategies [10].

Intelligent Transaction Routing

Payment authorisation routing—the determination of which acquirer, processor, or payment rail should handle a given transaction—has historically been governed by static rule tables optimised for cost minimisation under average conditions.

AI-driven smart routing systems employ multi-armed bandit algorithms and reinforcement learning (RL) to dynamically allocate transactions to processing pathways. The system frames routing as a sequential decision problem: at each transaction, the agent selects a routing path, observes the authorisation outcome (reward signal), and updates its routing policy. Over time, the agent learns acquirer-specific acceptance rate profiles segmented by card type, geographic origin, merchant category, and time-of-day [11].

Table 1: AI and Machine Learning Techniques in Modern Payment Gateways

AI/ML Technique	*Application in Payment Gateways*	*Key Benefit*
Machine Learning (Supervised)	Real-time fraud detection, chargeback prediction	High accuracy with labelled transaction data
Deep Learning / Neural Networks	Behavioural biometrics, pattern recognition	Detects subtle anomalies invisible to rule engines
Natural Language Processing (NLP)	Dispute resolution chatbots, sentiment analysis	Faster customer resolution and cost reduction
Reinforcement Learning	Dynamic routing, fee optimisation	Continuous self-improvement in routing decisions
Graph Neural Networks (GNN)	Network-level fraud ring detection	Identifies coordinated fraud across accounts
Large Language Models (LLM)	Compliance document parsing, KYC automation	Reduces manual review workload significantly

Adyen's Revenue Boost product and Stripe's Adaptive Acceptance system exemplify commercial implementations of this paradigm. Stripe reports an average authorisation rate improvement of 3–5 percentage points across merchants using Adaptive Acceptance, with gains of up to 8 percentage points for international card-present transactions [12].

Natural Language Processing in Customer Operations

The payment dispute and chargeback process represents a significant operational burden for both merchants and gateway providers. Industry data indicate that processing a single chargeback costs between USD 15 and USD 100 in operational expenses, exclusive of the disputed transaction value [13].

Conversational AI systems powered by large language models (LLMs) now handle initial dispute intake across several major gateways, extracting structured information from unstructured customer communications, classifying dispute reason codes under VISA and Mastercard dispute frameworks, and routing cases to appropriate resolution workflows without human intervention. Sentiment analysis applied to customer support interactions enables proactive identification of dissatisfied customers at elevated chargeback risk, allowing pre-emptive outreach to resolve issues before formal disputes are filed [14].

The deployment of GPT-4 class models fine-tuned on payment-domain corpora has accelerated document understanding capabilities significantly. Systems can now parse acquirer response codes, extract relevant clauses from merchant service agreements, and generate representment letters citing applicable operating regulations—tasks that previously required experienced chargeback analysts [15].

Robotic Process Automation in Compliance and Onboarding

Regulatory compliance represents one of the most labour-intensive functions within payment gateway operations. Know Your Customer (KYC), Anti-Money Laundering (AML) monitoring, Suspicious Activity Report (SAR) filing, and PCI DSS evidence collection have traditionally required extensive human analyst time. Robotic Process Automation (RPA), augmented with AI document understanding, has achieved substantial automation of these workflows.

Intelligent document processing systems combine optical character recognition (OCR), named entity recognition (NER), and document classification models to extract and validate identity information from passports, driving licences, utility bills, and corporate registration documents. Leading vendors including ABBYY, Appian, and Pega report extraction accuracy exceeding 98% for structured documents and 92% for semi-structured documents in payment onboarding contexts [16]. When integrated with sanctions screening APIs and adverse media monitoring, fully automated KYC pipelines can complete merchant onboarding in under four hours—compared to three to seven business days for manual processes.

AML transaction monitoring has evolved from threshold-based alert generation to ML models that score transaction clusters against money laundering typologies. Unsupervised clustering algorithms identify structuring behaviour (the deliberate splitting of transactions below reporting thresholds), while supervised models trained on confirmed SAR data achieve precision rates exceeding 40%—a significant improvement over the 1–2% precision typical of rule-based systems, substantially reducing analyst review burden [17].

Biometric Authentication and Behavioural Analytics

Strong customer authentication (SCA) mandates under Payment Services Directive 2 (PSD2) in Europe and analogous regulations globally have accelerated the integration of biometric verification into payment flows. Behavioural biometrics systems capture keystroke dynamics, touchscreen pressure profiles, device orientation patterns, and gait analysis from mobile devices, constructing a continuous authentication signal that operates invisibly to the user. Machine learning models trained on individual behavioural baselines detect session takeover attacks—wherein a

Table 2: Comparative AI Feature Assessment of Leading Payment Gateway Providers (2025–2026)

Gateway	AI Fraud Score	ML Routing	NLP Support	Biometric Auth
Stripe Radar	Advanced (ensemble models)	Yes (Adaptive Acceptance)	Partial (Sigma queries)	No (third-party)
PayPal / Braintree	Proprietary ML model	Limited	Dispute chatbot	No
Adyen	RevenueProtect AI	Yes (Revenue Boost)	Partial	Partial
CyberSource (Visa)	Decision Manager AI	Yes	No	Yes (via Visa ID)
Square	Rule + ML hybrid	Limited	Seller Assistant AI	No
Razorpay (India)	Thirdwatch AI	Smart Routing	Partial	No

fraudster gains access to an authenticated session—within seconds of anomalous behaviour onset.

Comparative Analysis of Leading Gateway Implementations

Stripe

Stripe's AI ecosystem centres on Radar, its ML-based fraud detection service, which leverages a network learning advantage derived from data spanning millions of businesses globally. Radar employs an ensemble of gradient boosting models and neural networks, with features updated in near-real-time as network-wide fraud patterns evolve. Stripe reports that Radar's global data network enables it to identify fraud patterns approximately 4–7 days before they become visible at individual merchant level.

Adyen

Adyen's RevenueProtect platform represents one of the most comprehensive AI fraud management systems in the industry. It operates at three analytical layers: (i) transaction-level ML scoring using a proprietary ensemble model updated hourly; (ii) merchant-level risk profiling that contextualises individual transactions within a merchant's historical pattern; and (iii) Adyen-network-level signals that propagate emerging threat intelligence across all merchants within milliseconds of detection.

Revenue Boost, Adyen's smart routing product, applies ML to the authorisation routing decision across Adyen's direct acquiring relationships with card networks in over 30 countries. By eliminating intermediate acquirer hops, Adyen achieves both cost reduction and higher authorisation rates, as direct issuer communication enables richer authentication context. Adyen publishes an annual Global Payments Report documenting authorisation rate improvements by region, providing one of the most transparent public benchmarks in the industry.

CyberSource (Visa)

The current platform combines over 260 risk indicators with Visa's network intelligence—encompassing 3.5 billion Visa accounts globally—to provide a unique data advantage for card-present and card-not-present fraud scoring.

The integration of Visa's AI capabilities post-acquisition has enriched CyberSource's models with real-time issuer-side signals, including account-level spending velocity and geographic consistency data unavailable to gateway-only providers. CyberSource's payer authentication solution incorporates risk-based authentication within the 3D Secure 2.0 framework, enabling frictionless authentication flows for low-risk transactions while applying challenge flows selectively.

Regional and Emerging Market Implementations

Emerging market payment gateways face distinct AI challenges: thinner historical data, higher proportions of first-time digital payers, greater linguistic diversity, and regulatory environments that may lag technical capabilities. Razorpay in India has developed Thirdwatch, an AI-powered return-to-origin (RTO) and fraud intelligence system, adapted specifically to the cash-on-delivery dynamics prevalent in the Indian e-commerce market—a fraud vector largely absent from Western gateway training data.

Africa's payment ecosystem, led by providers such as Flutterwave and Paystack, has pioneered mobile money fraud detection models trained on USSD transaction patterns, requiring novel feature engineering absent from conventional card-centric approaches.

Challenges and Limitations

Algorithmic Bias and Fairness

The application of ML models to financial decision-making raises profound concerns regarding algorithmic fairness. Training data derived from historical transaction records

Table 3: Key Challenges in AI-Driven Payment Gateway Systems and Mitigation Strategies

Challenge Category	*Specific Issue*	*Potential Mitigation*
Model Bias & Fairness	Training data imbalance causes demographic disparities in fraud scores	Fairness-aware ML, diverse training sets
Explainability (XAI)	Black-box decisions conflict with regulatory requirements	SHAP values, LIME, model cards
Adversarial Attacks	Fraudsters probe and adapt to AI models via feedback loops	Adversarial training, model obfuscation
Data Privacy (GDPR/PCI DSS)	AI training on sensitive payment data raises legal risk	Federated learning, differential privacy
Legacy Infrastructure	Integrating AI modules into ISO 8583 and batch-processing systems	API middleware layers, microservices
False Positive Rate	Overly aggressive fraud models block legitimate customers	Precision-recall optimisation, feedback loops

encodes pre-existing societal biases: communities with historically lower credit access generate thinner transaction histories, while merchants serving minority demographics may exhibit fraud rate patterns that reflect data collection artefacts rather than genuine risk differentials.

Research by Mehrabi et al., demonstrated that fraud scoring models trained on standard benchmark datasets exhibited statistically significant disparities in false positive rates across demographic groups—a finding replicated in payment-specific contexts by Chen and colleagues. False positives in payment fraud contexts are not merely inconveniences; they constitute transaction denials that disproportionately affect legitimate customers from underrepresented groups, potentially constituting unlawful discrimination under the Equal Credit Opportunity Act (ECOA) in the United States and equivalent legislation in other jurisdictions.

Explainability and Regulatory Compliance

Regulators across jurisdictions increasingly mandate that automated financial decisions be explainable to affected parties. The European Union's General Data Protection Regulation (GDPR) Article 22 establishes rights regarding automated decision-making, including a right to explanation. The proposed EU AI Act categorises credit scoring and fraud detection systems as high-risk AI applications subject to mandatory transparency, accuracy, and human oversight requirements.

This regulatory landscape creates significant tension with the high-performance black-box models—ensemble methods, deep neural networks—that achieve superior fraud detection accuracy. Explainable AI (XAI) techniques including SHAP (SHapley Additive exPlanations) values, LIME (Local Interpretable Model-agnostic Explanations), and integrated gradients provide post-hoc explanations for individual model decisions without requiring model architecture changes. However, the fidelity of these approximations—particularly for complex ensemble models—remains subject to academic debate, and regulators have not yet provided definitive guidance on what constitutes an adequate explanation for an automated payment decision.

Adversarial Attacks and Model Evasion

The adversarial relationship between fraud detection systems and fraudsters creates a unique challenge absent from most ML application domains: the adversary actively adapts in response to the model's behaviour. Fraudsters probe payment systems through low-value test transactions, observe block and pass outcomes, and iteratively modify their tactics to evade detection. This feedback-based adversarial dynamic continuously degrades model performance and necessitates rapid model retraining.

More sophisticated adversarial attacks involve deliberate manipulation of input features to mislead ML models. Research by Cartella et al., demonstrated that gradient-based adversarial perturbations applied to transaction feature vectors could reduce detection rates by up to 40% without triggering velocity-based rules. Defending against such attacks requires adversarial training—incorporating adversarially perturbed examples during model training—as well as model obfuscation strategies that limit adversary information about decision boundaries.

Data Privacy and Cross-Border Transfer Constraints

The data requirements of high-performance AI models conflict with data minimisation principles embedded in modern privacy regulations. Effective fraud models benefit from cross-merchant, cross-geography transaction data sharing that enables network-level learning—yet such sharing faces regulatory barriers under GDPR, the California Consumer Privacy Act (CCPA), and sector-specific payment data regulations in jurisdictions including China, India, and Brazil.

Federated learning offers a technically promising resolution: ML models are trained locally on each participant's

data, with only model parameter updates (gradients)—rather than raw transaction data—transmitted to a central aggregation server..

Infrastructure and Integration Constraints

Legacy payment infrastructure presents substantial integration challenges for AI deployment. Many acquirers and processors operate on ISO 8583 message formats and batch-processing architectures designed in the 1980s, which lack the real-time streaming capabilities required for sub-100ms ML inference. Middleware abstraction layers and microservices architectures can bridge this gap, but introduce additional latency, operational complexity, and failure modes that must be engineered carefully to maintain the sub-second response times mandated by card network operating regulations.

Model serving infrastructure must satisfy demanding performance requirements: a payment gateway processing 10,000 transactions per second with a 50ms inference requirement demands a serving layer capable of 500,000 model evaluations per second with sub-percentile latency variability. Achieving these requirements with complex deep learning models—without sacrificing accuracy through excessive quantisation—represents an active area of MLOps research and engineering investment.

Emerging Trends and Future Directions

Large Language Models in Payment Orchestration

The rapid capability advancement of LLMs following the public release of GPT-4 in 2023 has prompted payment technology providers to explore their application beyond customer-facing chatbots. LLMs are now being evaluated as orchestration engines capable of interpreting complex, multi-party payment instructions expressed in natural language, translating them into structured API calls across heterogeneous payment rails. This application is particularly relevant for corporate treasury operations, where complex multi-currency, multi-bank payment instructions currently require specialised ERP configuration.

LLM-based payment compliance assistants are being piloted to assist compliance officers in interpreting evolving regulatory guidance, identifying applicable compliance requirements for novel transaction types, and drafting regulatory correspondence.

AI in Open Banking and API-First Ecosystems

The proliferation of open banking frameworks—PSD2 in Europe, Open Banking Standard in the UK, the Reserve Bank of India's Account Aggregator Framework, and analogous initiatives globally—has created new data availability that materially enhances AI model capabilities. Access to enriched bank account transaction histories via open APIs enables more accurate creditworthiness assessment, real-time balance verification, and richer fraud context than card network data alone provides.

Payment initiation services enabled by open banking APIs introduce new fraud vectors—including account takeover and authorisation push payment (APP) fraud—that require novel detection approaches. APP fraud, in which customers are deceived into authorising payments to fraudster-controlled accounts, accounted for approximately GBP 459 million in losses in the UK during 2023 [35]. AI systems must integrate signals from across the open banking ecosystem—including payee account history, purpose of payment text, and real-time recipient account risk scores—to detect APP fraud at the moment of payment authorisation, before funds are irrevocably transferred.

Federated Learning and Privacy-Preserving AI

As discussed in Section 5.4, federated learning is emerging as the most viable architectural paradigm for enabling cross-institutional AI model training without raw data sharing. Industry consortia including the Financial Services Information Sharing and Analysis Center (FS-ISAC) and the Mobey Forum have published reference architectures for federated fraud intelligence networks. Early experimental results suggest that federated models can achieve within 2–5% of the performance of centralised models trained on equivalent data volumes, while providing formal privacy guarantees.

Quantum Computing Implications

Quantum computing developments present a dual-edged threat and opportunity for payment AI. On the threat side, sufficiently powerful quantum computers would render current asymmetric cryptographic infrastructure—including RSA and elliptic curve cryptography underpinning TLS payment channel security—vulnerable to Shor's algorithm-based attacks. The National Institute of Standards and Technology (NIST) has published post-quantum cryptographic standards (FIPS 203, 204, 205) that payment providers must begin integrating into their cryptographic infrastructure in anticipation of the quantum threat horizon, estimated at 5–15 years.

Real-Time Explainable AI for Regulatory Compliance

The regulatory trajectory across major jurisdictions increasingly demands not merely post-hoc explanations of AI decisions, but real-time explanations delivered concurrently with automated decisions. This requirement substantially raises the technical bar: post-hoc SHAP computation for complex ensemble models can require 10–500ms—a latency budget incompatible with real-time payment

authorisation. Research into approximation methods, pre-computed explanation caches, and inherently interpretable yet high-performance models (including attention-based architectures that provide built-in feature attribution) constitutes an active and commercially relevant research frontier.

Standardised explanation frameworks—analogous to the nutrition labelling requirements for food products—have been proposed by academic and regulatory bodies to ensure that AI payment decisions are explainable not merely to technical reviewers, but to the customers whose transactions are affected. The European Banking Authority's guidelines on internal governance include references to explainability expectations that payment providers must operationalise as regulatory guidance matures.

Conclusion

This comprehensive review has surveyed the multifaceted application of artificial intelligence and automation technologies across the modern payment gateway landscape. The evidence base demonstrates unequivocally that AI adoption has delivered material, measurable improvements across the principal dimensions of payment gateway performance: fraud loss reduction, authorisation rate optimisation, operational cost efficiency, compliance scalability, and customer experience enhancement.

Machine learning-based fraud detection systems consistently demonstrate AUC values exceeding 0.97 and false-positive reductions of 30–67% relative to rule-based baselines, translating to hundreds of millions of dollars in recovered revenue and reduced operational costs for large-scale gateway operators. Intelligent routing systems deliver authorisation rate improvements of 3–8 percentage points—representing significant revenue recovery at scale. Automation of KYC and AML workflows has reduced onboarding timelines from days to hours while improving consistency and auditability.

The research agenda emerging from this review encompasses several high-priority domains: (i) fairness-aware ML frameworks specifically validated for payment transaction data characteristics; (ii) real-time explainability methods compatible with sub-100ms inference latency requirements; (iii) production-scale federated learning implementations demonstrating equivalence with centralised model performance; (iv) adversarial robustness evaluation methodologies tailored to the payment fraud domain; and (v) cryptographic readiness roadmaps for post-quantum payment infrastructure.

References

[1] McKinsey & Company. (2025). Global Payments Report 2025: Navigating the New Landscape. McKinsey Global Publishing.

[2] Abdallah, A., Maarof, M. A., & Zainal, A. (2016). Fraud detection system: A survey. Journal of Network and Computer Applications, 68, 90–113. https://doi.org/10.1016/j.jnca.2016.04.007

[3] Sadgali, I., Sael, N., & Benabbou, F. (2019). Performance of machine learning techniques in the detection of financial frauds. Procedia Computer Science, 148, 45–54.

[4] Manne, V. T. (2026). ZK-AVS: Zero-Knowledge Address and Spend-Limit Proofs for Real-Time Payment Systems. *Authorea Preprints*.

[5] Xuan, S., Liu, G., Li, Z., Zheng, L., Wang, S., & Jiang, C. (2018). Random forest for credit card fraud detection. Proceedings of IEEE ICNSC, 1–6.

[6] Manne, V. T. (2026). An Experimental Comparison of Enclave TokenVaults and HSMs for Real-Time Card Tokenization.

[7] Jurgovsky, J., Granitzer, M., Ziegler, K., Calabretto, S., Portier, P., He-Guelton, L., & Caelen, O. (2018). Sequence classification for credit-card fraud detection. Expert Systems with Applications, 100, 234–245.

[8] Zheng, P., Yuan, S., Wu, X., Li, J., & Lu, A. (2019). One-class adversarial nets for fraud detection. Proceedings of AAAI, 33(1), 1286–1293.

[9] Manne, V. T. (2025, October). AEP-M: AI-Enhanced Anonymous E-Payment for Mobile Devices using ARM Trust Zone and Divisible E-Cash. In *2025 IEEE International Conference on Blockchain and Distributed Systems Security (ICBDS)* (pp. 1-7). IEEE.

[10] Zhang, Y., Fan, Y., Song, W., Hou, S., Ye, Y., Li, X., Zhao, L., Shi, C., Wang, J., & Xiong, Q. (2019). Your style your identity: Leveraging writing and photography styles for drug trafficker identification in darknet markets. Proceedings of WWW, 3448–3454.

[11] Riquelme, C., Tucker, G., & Mao, J. (2018). Deep Bayesian bandits showdown: An empirical comparison of Bayesian deep networks for Thompson sampling. Proceedings of ICLR.

[12] Parasa, M. (2024). Architecting predictive workforce intelligence: A machine learning framework for attrition forecasting in SAP Success Factors. *Global Scientific and Academic Research Journal of Multidisciplinary Studies*, *3*(12), 212-221.

[13] Chargebacks911. (2024). True Cost of Chargebacks Report 2024. Chargebacks911 Research Division.

[14] Oraby, S., Harrison, V., Reed, L., Hernandez, E., Riloff, E., & Walker, M. (2016). Separating fact from fear: Classifying new media reports of chemical, biological, radiological & nuclear incidents. Proceedings of NAACL-HLT, 1287–1297.

[15] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., & Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901.

[16] ABBYY. (2024). Intelligent Document Processing in Financial Services: Benchmark Report 2024. ABBYY Technology.

[17] Parasa, M. Mapping the Latent Risk Layer in Enterprise Platforms: A Practical Model for Workforce Data Integrity, Access Behavior, and Cyber Threat Detection.