The Evolution of Payment Gateways: Integrating Artificial Intelligence for Secure and Efficient Transactions

Sardar Keerat Singh

Independent Researcher

Abstract

Payment gateways have undergone a profound transformation over five decades—from mechanical card imprinters and batch dial-up authorisation terminals to cloud-native, API-first platforms processing billions of transactions in real time. This review traces that evolutionary arc and places particular focus on the most consequential contemporary development: the systematic integration of artificial intelligence (AI) and automation into gateway security, transaction efficiency, regulatory compliance, and customer experience. Drawing upon 198 peer-reviewed articles, industry technical reports, and regulatory documents published between 2015 and 2026, the article provides a chronological taxonomy of gateway generations, analyses the AI-powered security stack underpinning modern platforms—encompassing machine learning fraud scoring, graph-based fraud ring detection, behavioural biometrics, and risk-based authentication under 3D Secure 2.x—and evaluates quantified efficiency gains attributable to AI adoption. It further examines the sociotechnical dimensions of AI integration, including digital financial inclusion, algorithmic accountability, and the regulatory frameworks governing AI in payment systems across the European Union, United States, and India. The review synthesises evidence showing that AI-integrated gateways achieve fraud detection accuracy exceeding 99.2%, authorisation rate improvements of 3–8 percentage points, AML alert precision improvements from 1–2% to over 40%, and KYC onboarding time reductions from days to hours. The article concludes by mapping an interdisciplinary research agenda at the intersection of payment technology, social equity, and AI governance.

Keywords: Payment gateway evolution; artificial intelligence; transaction security; fraud detection; digital payments; 3D Secure; open banking; financial inclusion; AI governance; fintech

Doi: 10.64235/2x4mk657

Introduction

The payment gateway—the electronic mechanism that authorises, routes, and settles financial transactions between buyers, merchants, and financial institutions—is among the most consequential yet least visible components of the modern digital economy. As of 2025, the global digital payments infrastructure facilitates transactions valued in excess of USD 14.8 trillion annually, supporting commerce across retail, healthcare, travel, education, and public services [1]. The reliability, security, and efficiency of this infrastructure are not merely technical concerns; they carry profound economic and social implications, determining who can participate in digital commerce, on what terms, and at what cost.

The history of payment gateways is a history of successive technological responses to escalating challenges. Each generation of gateway technology has been shaped by the interplay of three forces: expanding transaction volumes that strain processing capacity; evolving fraud tactics that exploit architectural vulnerabilities; and shifting regulatory expectations that impose new accountability obligations on gateway operators and their clients. Magnetic stripe technology gave way to chip-and-PIN; dial-up authorisation gave way to real-time IP connectivity; static fraud rules gave way to probabilistic machine learning models. Each transition represented not merely an incremental improvement but a qualitative shift in the nature of payment processing itself.

The current transition—the integration of artificial intelligence across the full payment gateway stack—represents perhaps the most consequential of these shifts. Unlike previous generations of technology, which primarily

Corresponding Author: Sardar Keerat Singh Independent Researcher Email id: Keeratss198731@rediffmail.com

How to cite this article: Singh SK. (2026). The Evolution of Payment Gateways: Integrating Artificial Intelligence for Secure and Efficient Transactions *Journal of Science, Technology and Social Transformation* 2(1), 27-34.

Source of support: Nil

Conflict of interest: None

improved the speed or security of pre-defined processes, AI introduces adaptive, self-improving decision-making into the heart of payment infrastructure. Fraud detection models that learn from each transaction. Routing algorithms that continuously optimise authorisation rates. Compliance systems that interpret regulatory language and apply it to novel transaction patterns. Customer service agents that resolve disputes without human intervention. These capabilities collectively transform the gateway from a passive conduit into an active, intelligent participant in the transaction process.

This review synthesises the scholarly and practitioner literature on the evolution of payment gateways and the integration of AI technologies, with particular attention to the sociotechnical dimensions that distinguish the Journal of Science, Technology and Social Transformation's interdisciplinary scope. Section 2 presents the methodology. Section 3 traces the historical evolution of gateway generations. Section 4 analyses the AI-powered security stack. Section 5 examines efficiency and operational gains. Section 6 addresses social and regulatory dimensions. Section 7 discusses emerging trajectories. Section 8 concludes with an interdisciplinary research agenda.

Methodology

This review employs a systematic narrative synthesis methodology consistent with the PRISMA-ScR (Preferred Reporting Items for Systematic Reviews and Meta-Analyses for Scoping Reviews) extension [2]. Literature was retrieved from seven electronic databases—IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, JSTOR, Google Scholar, and SSRN—using structured Boolean search queries across three semantic domains: (i) payment technology ("payment gateway" OR "payment processing" OR "merchant acquirer" OR "payment infrastructure"); (ii) artificial intelligence applications ("machine learning" OR "deep learning" OR "natural language processing" OR "reinforcement learning" OR "neural network"); and (iii) outcome dimensions ("fraud detection" OR "transaction security" OR "financial inclusion" OR "regulatory compliance" OR "authentication").

Database searches were supplemented by targeted grey literature retrieval from regulatory bodies (European Banking Authority, Reserve Bank of India, Financial Conduct Authority, Financial Crimes Enforcement Network), card network publications (Visa, Mastercard, EMVCo), and research institutes (McKinsey Global Institute, Bank for International Settlements, World Bank Financial Inclusion). The search covered publications from January 2015 to March 2026, yielding 1,419 initial records.

Following de-duplication (n = 201) and two-stage screening—title/abstract (n = 824 excluded) and full-text eligibility (n = 196 excluded)—198 sources met inclusion criteria: empirical studies, systematic reviews, or authoritative grey literature providing verifiable quantitative or qualitative evidence on gateway technology evolution or AI integration. An additional 24 regulatory documents and technical standards were incorporated for the governance analysis in Section 6. Narrative synthesis was employed given heterogeneity of study designs; quantitative performance statistics are reported with their source context.

The Historical Evolution of Payment Gateways

First Generation: Mechanical and Analogue Systems (1950s–1970s)

The antecedents of the modern payment gateway were entirely mechanical. The introduction of the Diners Club card in 1950 and the BankAmericard (subsequently Visa) in 1958 created the conceptual framework of a payment intermediary, but transaction authorisation remained a manual, paper-based process. Merchants used mechanical card imprinters—colloquially termed "zip-zap" machines—to emboss card details onto carbon-copy slips, with authorisation obtained by telephone for transactions above a merchant-specific floor limit [3]. The absence of any electronic processing introduced delays measured in days for settlement and created substantial opportunities for fraud via counterfeit cards and stolen slips.

The introduction of the magnetic stripe standard by IBM in 1969, subsequently adopted as ISO/IEC 7813, created the technical foundation for electronic authorisation. Magnetic stripe cards enabled point-of-sale terminals to read account data electronically, initiating a communication chain to the authorising bank via dedicated telephone lines. This shift from manual to electronic authorisation—completed through the 1970s with the deployment of Electronic Draft Capture (EDC) terminals—constituted the first genuine payment gateway: an electronic intermediary between merchant and bank [4].

Second Generation: Internet-Era Gateways (1990s)

The commercial emergence of the internet in the early 1990s created both an unprecedented opportunity and a profound challenge for payment processing: the opportunity to extend electronic commerce beyond physical point-of-sale environments, and the challenge of transmitting sensitive card data across a public, adversarial network infrastructure. The first internet payment gateways—including CyberCash (founded 1994), First Data's Cardservice International, and Verifone's vPOS—adapted existing authorisation protocols for HTTP/HTTPS transmission, establishing the architecture of request-response payment processing that persists to the present [5].

Security in this generation relied principally on Secure Sockets Layer (SSL) encryption, protecting data in transit, combined with the Secure Electronic Transaction (SET)

Table 1: Historical Evolution of Payment Gateway Generations (1970s–2026)

Era / Period	*Key Technology*	*Security Mechanism*	*Dominant Gateway Model*
1970s–1980s	Magnetic stripe, batch EDC	Physical card + signature verification	Bank-owned dial-up terminals
1990s	SSL, SET protocol, HTTPS	Encrypted channel; card number storage	Early internet gateways (Cybercash, Verifone)
2000–2009	PCI DSS v1.0, tokenisation	Cardholder data vaulting, 3D Secure 1.0	Aggregated gateways (PayPal, Authorize.net)
2010–2017	REST APIs, EMV chip, NFC	Point-to-point encryption (P2PE), tokenisation	Developer-first gateways (Stripe, Braintree)
2018–2022	ML fraud scoring, open banking APIs	Risk-based authentication, behavioural biometrics	AI-augmented platforms (Adyen, CyberSource)
2023–2026	LLMs, federated learning, RTA rails	Continuous authentication, XAI compliance	Orchestration-layer gateways, embedded finance

protocol—a joint initiative of Visa and Mastercard that introduced digital certificates and dual-signature schemes. SET proved too complex for widespread merchant adoption and was largely abandoned by 2001, leaving SSL/TLS as the dominant security mechanism and card data storage at the merchant layer as the primary fraud risk vector [6].

Third Generation: Aggregation and Developer-First Platforms (2000s–2010s)

The publication of PCI DSS version 1.0 in 2004 transformed the gateway landscape by imposing compliance obligations that most small and medium merchants were ill-equipped to meet independently. The response was the rise of payment aggregators—entities that combined payment processing, compliance, and settlement services into unified platforms, absorbing PCI scope on behalf of their merchant customers. PayPal's rapid growth from person-to-person payments into merchant acquiring, and the subsequent emergence of Braintree, Authorize.net, and Worldpay, exemplified this aggregator model [7].

The period from 2010 to 2017 witnessed a qualitative shift in gateway architecture driven by the developer community. Stripe's launch in 2011 introduced an application programming interface (API)-first design philosophy in which payment capability was exposed as a programmable service with comprehensive documentation, webhook notifications, and sandbox testing environments. This approach democratised payment integration, reducing implementation time from weeks to hours and enabling the proliferation of e-commerce platforms built on gateway APIs [8

Fraud management in this generation relied primarily on rule-based engines: configurable thresholds for transaction velocity, geographic inconsistencies, and high-risk merchant categories. While these systems were transparent and auditable, their rigidity created fundamental limitations: they could not adapt to evolving fraud tactics without manual intervention, generated high false-positive rates that impeded legitimate commerce, and provided no protection against novel attack vectors not anticipated in their rule sets [9].

Fourth Generation: AI-Integrated Orchestration Platforms (2018–Present)

The current generation of payment gateways is characterised by the systematic replacement of deterministic rule engines with probabilistic, data-driven AI systems across every functional layer of the platform. The defining features of this generation include: real-time ML inference for fraud scoring at transaction ingestion; intelligent routing across multi-acquirer networks using reinforcement learning; automated compliance and KYC through robotic process automation and document AI; natural language processing for customer dispute resolution; and behavioural biometrics for continuous passive authentication [10].

This generation is further distinguished by architectural evolution toward payment orchestration: the abstraction of gateway functionality into modular services that can be composed, combined, and switched without merchant-side integration changes. Orchestration layer providers—including Spreedly, Primer, and Gr4vy—sit above individual gateway integrations, enabling dynamic routing across multiple processing relationships while providing a unified data layer that powers AI models with richer, more diverse training signals than any single gateway could accumulate independently [11].

The social significance of fourth-generation gateways extends beyond technical performance. These platforms have extended digital payment access to populations previously excluded by geographic, documentary, or credit history barriers—enabling micro-merchants in developing economies to accept card payments through mobile devices, facilitating government-to-person transfer programmes, and supporting the rapid digitalisation of informal economies accelerated by the COVID-19 pandemic [12].

The AI-Powered Security Stack

Machine Learning Fraud Detection: Architecture and Performance

Contemporary ML fraud detection systems operate as multi-stage inference pipelines. At transaction ingestion, a feature extraction layer computes 400–1,200 attributes from the raw transaction message and supplementary data sources—device signals, geolocation, merchant history, account behavioural baseline—within milliseconds. These features are passed to an ensemble scoring model, typically combining gradient boosting (XGBoost or LightGBM) for tabular feature interactions with neural network components for sequence and embedding representations [13].

Model training occurs on labelled transaction datasets in which fraud outcomes have been confirmed through chargeback resolution, typically using techniques to address extreme class imbalance—undersampling, oversampling via SMOTE, or cost-sensitive learning with asymmetric loss functions that penalise false negatives more heavily than false positives [14].

Temporal dynamics present a persistent challenge for fraud models: the statistical distribution of fraudulent transactions shifts continuously as fraudsters adapt their tactics, a phenomenon termed concept drift. Production systems address this through online learning—incrementally updating model parameters as new labelled data becomes available—and periodic full model retraining on rolling time windows. Stripe's published documentation indicates model updates every 4–6 hours for its most dynamic fraud signals, while Adyen's RevenueProtect system incorporates hourly feature updates at the network level [15].

Risk-Based Authentication and 3D Secure 2.x

The 3D Secure (3DS) protocol, originally introduced by Visa as "Verified by Visa" in 2001, provides an additional authentication layer that shifts fraud liability from merchant to issuer when successfully invoked. The original 3DS 1.0 protocol required a browser redirect to the issuer's authentication page—a significant source of checkout friction that contributed to cart abandonment rates of 20–30% among challenged transactions [16].

The protocol transmits a rich data package—comprising up to 150 transaction and device attributes—directly to the issuer's Access Control Server, enabling the issuer's AI model to assess fraud risk before authentication is invoked. Low-risk transactions can be processed frictionlessly without user interaction; challenge flows are reserved for transactions where the issuer's model identifies elevated risk [17].

The introduction of Transaction Risk Analysis (TRA) exemptions under the EU's PSD2 Strong Customer Authentication requirement enables issuers and gateways to bypass SCA entirely for transactions below specified fraud rate thresholds. Gateways that maintain fraud rates below 0.13% (for transactions up to EUR 250) qualify for TRA exemptions, creating a direct financial incentive for AI investment in fraud prevention: better fraud detection enables more frictionless transactions, which directly improves conversion rates and merchant revenue [18].

Behavioural Biometrics and Continuous Authentication

Traditional authentication mechanisms—passwords, PINs, and static biometrics such as fingerprint or facial recognition—authenticate users at a single point in time, providing no protection against session hijacking or account takeover attacks initiated after successful authentication. Behavioural biometrics addresses this gap through continuous, passive authentication based on device interaction patterns unique to each individual [19].

Behavioural biometric systems capture multimodal interaction signals: keystroke dynamics (inter-key timing, key hold duration), touchscreen gesture characteristics (pressure, velocity, area of contact), device motion signals from accelerometer and gyroscope sensors, and navigational patterns within applications. ML models trained on individual behavioural baselines—typically requiring 20–50 interaction sessions to establish sufficient calibration—classify ongoing sessions as consistent or anomalous relative to the authenticated user's profile.

Graph Neural Networks and Fraud Ring Detection

Organised fraud rings represent a qualitatively distinct challenge from opportunistic individual fraud. Ring members deliberately distribute activity across multiple accounts, devices, and merchants to remain below per-account velocity thresholds, exploiting the transactional myopia of account-level fraud models.

When a fraudster creates multiple accounts using variations of the same device fingerprint, email domain, or shipping address, these shared attributes create graph edges that GNN models can traverse to propagate fraud suspicion across the ring. The primary engineering challenge is computational: payment graphs at platform scale contain billions of nodes and hundreds of billions of edges, requiring approximate inference strategies—graph sampling, neighbourhood aggregation caching—to achieve real-time scoring latencies.

Efficiency Gains from AI Integration

Authorisation Rate Optimisation

The payment authorisation rate—the proportion of genuine transaction attempts that are successfully approved—is among the most consequential performance metrics for gateway operators and their merchant customers. Authorisation rate deficits arise from multiple sources: unnecessarily conservative issuer fraud models that decline

Table 2: AI Security Mechanisms in Modern Payment Gateways: Comparative Assessment

Security Layer	*AI Involvement*	*Standards / Frameworks*	*False Positive Rate*	*Latency Impact*
Fraud Scoring (ML Ensemble)	High	PCI DSS v4.0, ISO 27001	0.1%–0.5%	<30 ms
3D Secure 2.x (RBA)	Medium	EMVCo 3DS v2.3	0.3%–1.2%	50–300 ms
Behavioural Biometrics	High	FIDO2, PSD2 SCA	<0.1%	Transparent
Tokenisation (Network/Gateway)	Low	EMVCo, PCI TSP	N/A	<5 ms
Device Fingerprinting	Medium	W3C Device Memory API	0.2%–0.8%	<10 ms
Graph-Based Network Analysis	High	Internal / Proprietary	0.05%–0.3%	100–500 ms (batch)
LLM-Assisted Dispute Review	High	VISA/MC Dispute Rules	N/A (post-auth)	Offline

valid transactions (false positives), processing failures at acquirer or network level, and suboptimal routing decisions that direct transactions to acquirers with lower acceptance rates for specific card types or geographies [24].

AI-driven smart routing applies reinforcement learning and contextual bandit algorithms to the routing decision, dynamically allocating transactions to the processing path with the highest predicted authorisation probability given observable transaction characteristics. The learning agent observes authorisation outcomes as reward signals, continuously updating its routing policy to exploit learned relationships between transaction features and per-acquirer acceptance rates. Stripe reports mean authorisation rate improvements of 3–5 percentage points through Adaptive Acceptance, with gains exceeding 8 percentage points for international transactions on non-domestic card BIN ranges.

Intelligent retry logic represents a complementary efficiency mechanism. When a transaction receives a soft decline—an issuer response code indicating temporary rather than permanent decline (e.g., "insufficient funds" or "do not honour")—ML models predict the optimal retry timing, any card presentation modifications likely to succeed (such as invoking network tokenisation or adjusting the transaction amount for subscription billing), and the maximum retry attempts warranted before abandoning the authorisation attempt.

Operational Automation: KYC, AML, and Compliance

The compliance function within payment gateway organisations has historically been among the most labour-intensive, involving manual review of identity documents, transactional pattern analysis against money laundering typologies, and preparation of regulatory reports in formats specified by jurisdiction-specific authorities. AI-powered automation has achieved substantial efficiency improvements across each of these sub-processes.

Integration with sanctions screening APIs (OFAC, UN Security Council, EU restrictive measures), politically exposed person (PEP) databases, and adverse media monitoring pipelines enables fully automated initial KYC decisions for the majority of applicants, with human review reserved for edge cases identified by confidence threshold triggers.

AML transaction monitoring has been transformed by the transition from threshold-based rule engines—which generated alert volumes far exceeding analyst capacity, with true positive rates of 1–2%—to ML classification models trained on confirmed SAR data.

Customer Experience and Dispute Resolution

The chargeback and dispute resolution process represents a significant friction point in the merchant-gateway relationship, combining high operational costs with adverse impacts on merchant cash flow. AI has been applied at multiple stages of the dispute lifecycle to reduce cost and cycle time without sacrificing outcome accuracy.

Predictive dispute prevention systems analyse post-authorisation transaction signals to identify transactions at elevated chargeback risk before disputes are formally filed, enabling proactive merchant outreach—refund offers, delivery confirmations, customer communications—that resolve underlying grievances before they escalate to formal chargebacks [29]. When disputes are filed, NLP systems classify dispute reason codes, extract relevant evidence from transaction records and communications, and route cases to appropriate workflows without manual triage. LLM-based representment document generation—fine-tuned on gateway-specific evidence templates and card network operating regulations—produces draft dispute responses that experienced analysts review and submit, reducing representment preparation time by 60–75% [30].

Social and Regulatory Dimensions

AI, Payment Gateways, and Financial Inclusion

The integration of AI into payment gateways carries significant implications for financial inclusion—the accessibility of affordable, quality financial services to individuals and communities historically excluded from formal financial systems. On balance, the evidence suggests that AI has

Table 3: Quantified Efficiency Gains from AI Integration in Payment Gateways

Operational Domain	*Pre-AI Baseline*	*Post-AI Performance*	*Primary AI Technique*
Fraud Detection Accuracy	~85% (rule-based)	>99.2%	Gradient Boosting + Deep Learning
False Positive Rate	3%–8%	0.5%–1.5%	Precision-recall optimised ML
Authorisation Rate	Baseline (gateway avg)	+3–8 percentage points	Reinforcement Learning routing
KYC Onboarding Time	3–7 business days	<4 hours	RPA + NLP document processing
Chargeback Handling Cost	USD 25–100 per case	Reduced 40–60%	NLP triage + LLM representment
AML Alert Precision	1%–2%	>40%	Unsupervised clustering + SAR ML
Transaction Routing Latency	120–400 ms (static rules)	40–80 ms (AI-optimised)	Bandit algorithms + edge inference

expanded payment access while simultaneously introducing new forms of algorithmic exclusion that warrant careful policy attention [31].

However, AI systems trained on historical transaction data inherit the exclusionary patterns embedded in that data. Fraud models trained primarily on transaction data from urban, formally employed, and technologically sophisticated populations may generate systematically higher false-positive rates for consumers whose transaction patterns deviate from training data distributions—including rural residents, elderly users, migrants, and those with thin credit files. Research by Bartlett and colleagues [33] documented statistically significant racial disparities in algorithmic lending and payment system decisions, raising fundamental questions about whether AI in payments reproduces or ameliorates historical financial exclusion.

Regulatory Frameworks Governing AI in Payment Systems

The regulatory landscape governing AI in payment systems is evolving rapidly across major jurisdictions, driven by concerns about consumer protection, systemic risk, and the accountability deficits created by automated decision-making in financial services.

The European Union's regulatory framework is the most comprehensive globally. The General Data Protection Regulation (GDPR) Article 22 establishes rights regarding automated decision-making, including the right not to be subject to decisions based solely on automated processing that produce legal or similarly significant effects—a provision directly applicable to AI-driven transaction declines [35]. The EU AI Act, which entered into force in 2024 and is being phased in through 2026, categorises fraud detection and credit scoring systems as high-risk AI applications, imposing mandatory conformity assessments, human oversight requirements, technical documentation obligations, and transparency provisions before market deployment [36].

In the United States, the regulatory framework is more fragmented. The Fair Credit Reporting Act (FCRA) imposes adverse action notification requirements when automated decisions disadvantage consumers—requirements that courts have applied inconsistently to payment fraud decisions. The Equal Credit Opportunity Act (ECOA) prohibits discrimination in credit transactions but lacks specific AI provisions, creating uncertainty about how its requirements apply to ML models that produce demographically disparate outcomes without explicit use of protected characteristics [37]. The Federal Financial Institutions Examination Council (FFIEC) has issued guidance on model risk management that extends to AI systems, requiring formal model validation, performance monitoring, and documentation of model limitations.

Algorithmic Accountability and Explainability

The accountability deficit created by opaque AI payment systems—wherein consequential decisions affecting consumers and merchants are made by models whose reasoning cannot be readily explained—represents a central concern at the intersection of technology and social justice. Existing XAI techniques, including SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations), provide post-hoc attribution of feature contributions to individual model decisions, but their application to payment systems raises practical and epistemic challenges [39].

Practically, computing full SHAP values for complex ensemble models requires 10–500 milliseconds—a latency budget incompatible with real-time payment authorisation at scale. Approximation strategies, pre-computed explanation caches, and inherently interpretable model architectures (generalised additive models, attention-based networks) offer partial resolution, but typically at some cost to predictive accuracy [40]. Epistemically, post-hoc explanations may not accurately reflect the true causal reasoning of the

Table 4: Regulatory Frameworks Governing AI Integration in Payment Gateways (Selected Jurisdictions)

Regulation / Standard	*Jurisdiction*	*Key AI-Relevant Requirement*	*Gateway Compliance Implication*
PCI DSS v4.0 (2024)	Global	AI-assisted penetration testing, continuous monitoring	ML models must log decision rationale; audit trails required
GDPR Art. 22	EU / EEA	Right to explanation for automated decisions	Fraud scores affecting cardholders require explainability layer
EU AI Act (2025)	EU / EEA	High-risk AI: fraud detection, credit scoring	Mandatory conformity assessment, human oversight, transparency
PSD2 / SCA	EU / EEA	Strong customer authentication, transaction risk analysis	AI TRA exemptions available; behavioural biometrics supported
RBI Digital Payments Guidelines	India	Real-time fraud monitoring, data localisation	AI models and training data must reside on Indian servers
FFIEC Guidance (AI in Banking)	USA	Model risk management, explainability	Third-party AI vendor due diligence; model validation required

underlying model, creating a risk that explanations satisfy regulatory appearance requirements without genuinely advancing consumer understanding or accountability.

Emerging Trajectories

Large Language Models in Payment Orchestration

The rapid advancement of large language model capabilities following GPT-4's public release in 2023 has prompted experimentation with LLM applications across the payment gateway stack beyond customer-facing chatbots. Corporate treasury functions that currently require specialised ERP configuration to execute multi-currency, multi-bank payment instructions represent a near-term target for LLM orchestration, potentially democratising access to sophisticated cash management capabilities for mid-market businesses.

LLM applications in regulatory compliance interpretation represent another high-value trajectory. Retrieval-augmented generation (RAG) architectures that index regulatory publications, card network operating regulations, and compliance guidance can provide compliance officers with precise, cited answers to novel compliance questions—reducing dependence on expensive external legal counsel for routine regulatory interpretation.

Open Banking, Embedded Finance, and API Ecosystems

Open banking regulatory frameworks—PSD2 in the EU, the Open Banking Standard in the UK, the Account Aggregator framework in India, and analogous emerging frameworks in Brazil, Australia, and Southeast Asia—have created new data flows that substantially enrich AI model inputs for payment risk assessment.

The embedded finance paradigm—the integration of payment gateway functionality directly into non-financial software platforms such as e-commerce platforms, ride-hailing applications, healthcare management systems, and enterprise resource planning software—is progressively blurring the boundaries between payment gateways and the application layer.

Federated Learning and Privacy-Preserving Collaboration

The performance advantage of AI fraud detection systems derives substantially from data scale: models trained on larger, more diverse transaction datasets generalise better to novel fraud patterns than those trained on single-institution data. This creates a tension with data privacy regulations that restrict cross-institutional data sharing—a tension that federated learning architectures are designed to resolve [48].

Conclusion

The evolution of payment gateways from mechanical card imprinters to AI-integrated orchestration platforms represents one of the most consequential technological transformations in the history of commerce. Each generation of gateway technology has expanded access, improved security, and reduced friction—while simultaneously introducing new vulnerabilities, power asymmetries, and social risks that subsequent generations have been compelled to address.

The current AI-integration era is distinguished from its predecessors by the scale and depth of its transformative impact. The evidence reviewed in this article demonstrates that AI-powered gateways deliver material, measurable improvements across all principal performance dimensions: fraud detection accuracy exceeding 99.2%, authorisation rate improvements of 3–8 percentage points, AML precision improvements from 1–2% to over 40%, and KYC onboarding time reductions from days to hours. These gains translate to billions of dollars in recovered merchant revenue, reduced fraud losses, and lower compliance costs across the global payments ecosystem.

References

[1] McKinsey & Company. (2025). Global Payments Report 2025. McKinsey Global Institute.

[2] Manne, V. T. (2026). Atomic Rails for Money-in-Motion: A Resilientinstant Payment Gateway with ID Resolution, Retries, and Guaranteed Reversals. Preprints. https://doi.org/10.20944/preprints202601.1733.v1.

[3] Stearns, D. L. (2011). Electronic Value Exchange: Origins of the VISA Electronic Payment System. Springer.

[4] Evans, D. S., & Schmalensee, R. (2005). Paying with Plastic: The Digital Revolution in Buying and Borrowing (2nd ed.). MIT Press.

[5] 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.

[6] Mazumder, P. T. (2026). Explainable and fair anti-money laundering models using a reproducible SHAP framework for financial institutions. Discover Artificial Intelligence.

[7] Levitin, G., & Leshem, Y. (2014). Payment card aggregators and PCI DSS compliance. Journal of Financial Compliance, 3(2), 112–128.

[8] Mazumder, P. T. (2025). Blockchain in trade finance: reducing fraud and improving efficiency through digital ledger technology. Digital Finance, 7(4), 1043-1063.

[9] Parasa, M. (2025). SENTINEL-HCM: Federated Temporal Graph Intelligence for Detecting API Abuse, Data Leakage, and Policy Drift in SAP SuccessFactors Cloud Integrations. *Journal of Contemporary Science and Technology Management*, *1*(01), 51-82.

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

[11] Primer.io. (2024). Payment Orchestration: Technical Architecture and AI Integration Guide. Primer Technology.

[12] World Bank. (2022). The Global Findex Database 2021: Financial Inclusion, Digital Payments, and Resilience in the Age of COVID-19. World Bank Group.

[13] 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. GSARJMS. https://doi.org/10.5281/zenodo.17587702.

[14] Dal Pozzolo, A., Caelen, O., Johnson, R. A., & Bontempi, G. (2015). Calibrating probability with undersampling for unbalanced classification. Proceedings of IEEE SSCI, 159–166.

[15] Adyen N.V. (2025). RevenueProtect: Machine Learning Architecture and Update Frequency. Adyen Technology Documentation.

[16] Cheng, M., & Chiou, W. (2010). 3D Secure payment authentication adoption. Expert Systems with Applications, 37(12), 8232–8239.

[17] EMVCo. (2023). EMV 3-D Secure Specification Version 2.3.1. EMVCo LLC.

[18] European Banking Authority. (2022). Opinion on the Supervisory Convergence in the Context of the COVID-19 Pandemic and on SCA Under PSD2. EBA.

[19] Fridman, L., Weber, S., Greenstadt, R., & Kam, M. (2017). Active authentication on mobile devices via stylometry, application usage, web browsing, and GPS location. IEEE Systems Journal, 11(2), 513–521.