Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

AI-Driven KYC and AML Solutions: Improving Efficiency and Risk Management in FinTech

Conference Research Paper

Nitin M
BMS College of Commerce and Management, Bengaluru
Email: nitinmahesh0@gmail.com

Conference Theme

Artificial Intelligence and Emerging Technologies in Financial Services — Innovation, Compliance and Risk Management

Keywords: Artificial Intelligence, KYC, AML, FinTech, Risk Management
October 2026

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

Abstract

The rapid growth of financial technology has transformed digital banking and financial services, while simultaneously creating new challenges related to identity verification, fraud, money laundering, and regulatory compliance. Know Your Customer (KYC) and Anti-Money Laundering (AML) processes are traditionally dependent on extensive documentation, manual verification, and rule-based monitoring, which can result in delays, operational costs, and difficulties in identifying sophisticated financial risks. This paper examines the potential of Artificial Intelligence (AI) to improve KYC and AML processes within the FinTech ecosystem. AI-based technologies, including machine learning, natural language processing, automated document verification, anomaly detection, and predictive analytics, can support faster customer onboarding, continuous transaction monitoring, risk classification, and identification of suspicious activities.

The study explores how AI-driven solutions can contribute to operational efficiency while strengthening financial risk management and regulatory compliance. It also considers challenges associated with data privacy, algorithmic bias, explainability, cybersecurity, and the need for appropriate human oversight. The paper proposes an integrated perspective in which AI complements rather than completely replaces human decision-making in financial compliance. By examining the opportunities and challenges associated with AI-enabled KYC and AML, the study highlights how technological innovation can support safer, more efficient, and scalable digital financial services.

Keywords: Artificial Intelligence, Know Your Customer (KYC), Anti-Money Laundering (AML), FinTech, Risk Management, Regulatory Compliance.

1. Introduction

Financial technology (FinTech) has fundamentally reshaped how individuals and businesses access banking, payments, lending, and investment services. Digital-first banks, payment platforms, and lending applications now onboard customers entirely online, often within minutes of a user downloading an app. This convenience, however, has widened the attack surface for identity fraud and has made financial systems more attractive conduits for money laundering, terrorist financing, and other illicit financial flows. Regulators worldwide have responded by tightening Know Your Customer (KYC) and Anti-Money Laundering (AML) obligations, requiring institutions to verify customer identity, assess risk, and continuously monitor transactions for suspicious activity.

Historically, KYC and AML compliance has been a labour-intensive function built around paper documentation, manual data entry, and static, rule-based transaction-monitoring engines. While such systems satisfy a baseline regulatory requirement, they scale poorly: as transaction volumes grow, rule-based engines tend to generate overwhelming numbers of alerts, the large majority of which turn out to be false positives, consuming analyst time that could otherwise be directed toward genuine threats (Youverify, 2026; Delos Intelligence, 2026). At the same time, customers

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

increasingly expect onboarding to be instantaneous, creating tension between the thoroughness regulators demand and the speed customers expect.

Artificial Intelligence offers a route out of this tension. Machine learning models can score customer and transaction risk dynamically rather than against static thresholds; natural language processing can parse unstructured data such as adverse-media reports and sanctions-list entries; computer vision and optical character recognition can verify identity documents automatically; and anomaly-detection algorithms can flag behavioural deviations that rule-based systems would miss entirely. Industry surveys suggest that this shift is already well underway: more than 60 per cent of financial institutions now use AI in some part of their AML or KYC function (StealthAgents Research, 2026, citing McKinsey, 2025).

This paper examines how AI-driven technologies are reshaping KYC and AML processes within the FinTech ecosystem, with a specific focus on operational efficiency and risk management. It reviews the relevant literature and industry evidence, outlines the specific AI techniques in use, quantifies the efficiency gains reported across multiple industry studies, and critically examines the risks — data privacy, algorithmic bias, explainability, and cybersecurity — that accompany AI adoption in a regulated environment. The paper concludes by proposing an integrated framework in which AI augments, rather than replaces, human compliance judgment.

2. Literature Review

A growing body of industry and academic research documents the efficiency gains associated with AI-driven KYC and AML systems, while also flagging the governance challenges that accompany them.

1. On transaction monitoring, research compiled by Deloitte and ACAMS indicates that AI-based monitoring can reduce AML false-positive alert volumes by 50 to 80 per cent relative to legacy rule-based engines, directly freeing analyst capacity for genuine investigations (StealthAgents Research, 2026).
2. Industry case evidence from Delos Intelligence (2026) reports that AI-powered Know-Your-Business (KYB) verification can cut onboarding cycle time from roughly eighteen business days to about two-and-a-half days, while reducing false-positive AML alerts by close to 78 per cent through contextual, graph-based analytics.
3. A separate Delos Intelligence (2026) study of AI-powered AML and KYC screening found that complex KYC cases that previously took eight to twelve business days, and forty-five to ninety minutes of analyst time per file, could be processed in under two hours using AI-assisted screening — with suspicious-activity-report drafting time falling from several hours to roughly fifteen to twenty minutes.
4. Wecan Group (2026) similarly documents standard KYC onboarding falling from fifteen to twenty-one days to two to three hours, and sanctions-screening false positives falling from the industry-typical 95–99 per cent range to 20–25 per cent, after automation.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

5. At an industry-wide level, McKinsey's 2025 research estimates that AI adoption across AML and KYC functions can lower total financial-crime compliance costs by 20 to 35 per cent, with document-verification automation and AI-based risk scoring already in use at 52 per cent and 44 per cent of surveyed institutions respectively.
6. Youverify (2026), drawing on Fenergo's 2025 Financial Crime Industry Trends report, notes that seventy per cent of financial institutions reported losing clients in the preceding year because of slow onboarding — the highest rate on record — underscoring that onboarding speed is now a competitive, and not merely operational, concern.
7. Research on false-positive drivers (Youverify, 2026) attributes high legacy alert volumes to four compounding factors: overly broad detection rules, weak name-screening logic, static risk models that fail to adapt to customer behaviour, and incomplete data collected at onboarding — each of which AI-based, adaptive risk-scoring approaches are designed to address.
8. Napier AI's published research, cited in sector analysis, reports that machine-learning-augmented monitoring delivers up to 40 per cent reductions in false-positive rates compared with rules-only approaches, with comparable gains observed across different regulatory jurisdictions, including South African banking markets (Youverify, 2026).
9. Taken together, this literature converges on a consistent picture: AI materially improves the speed and precision of KYC/AML operations, but almost every source also stresses that AI functions as a triage and prioritisation layer feeding into, rather than replacing, human compliance review — a theme this paper develops further in Sections 6 and 7.

A clear gap in the existing commentary is that most of the cited evidence comes from vendor and industry-research sources rather than independent academic study, and reported efficiency figures vary considerably by institution size, geography, and baseline process maturity. This paper treats the figures as indicative of direction and order of magnitude rather than as precise, universally applicable benchmarks, and it draws on evidence from several independent sources to triangulate a reasonable estimate of the efficiency gains achievable (see Section 5).

3. Traditional KYC/AML Processes and Their Limitations

Traditional KYC workflows begin with the customer submitting identity and address proof — typically scanned or photographed documents — which compliance staff then re-key into the institution's systems, cross-check manually against the submitted documents, and screen against sanctions and politically-exposed-person watchlists using largely exact-match logic. AML transaction monitoring, in parallel, relies on static rule sets (for example, flagging any transaction above a fixed value or any rapid sequence of transfers) that do not adapt to an individual customer's normal behaviour.

This architecture produces two structural problems. First, because rules are deliberately set broad to avoid missing genuine risk, they generate very large numbers of alerts, the overwhelming

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

majority of which are false positives — published estimates place legacy false-positive rates at 90 to 99 per cent of all flagged alerts (Delos Intelligence, 2026; Youverify, 2026). Second, document checks, data entry, and alert review are manual and sequential, so onboarding and investigation timelines stretch into days or weeks even for low-risk, routine cases. Figure 1 below contrasts the two workflows step by step.

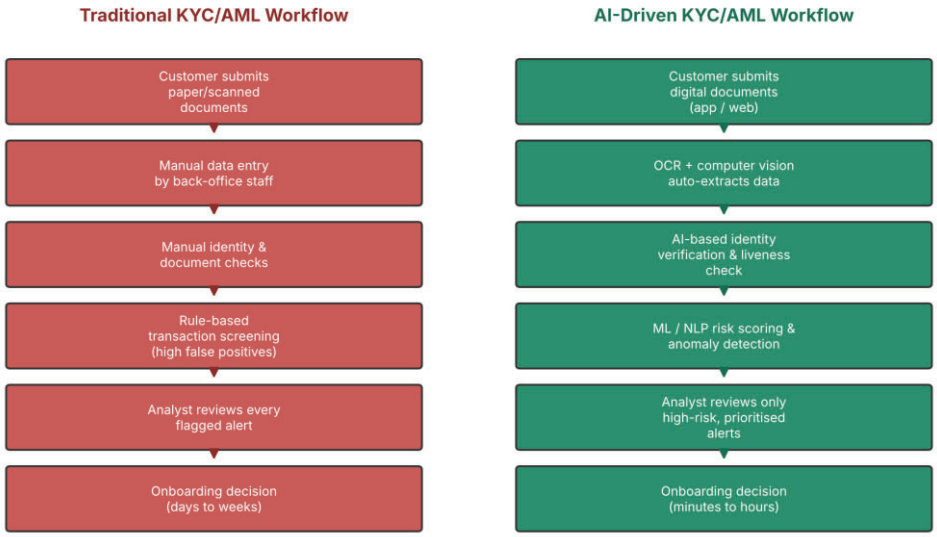


Figure 1. Traditional KYC/AML workflow (manual, sequential) compared with an AI-driven KYC/AML workflow (automated, risk-prioritised).

As the comparison illustrates, the AI-driven path does not eliminate human judgment — it relocates it to the point where it adds the most value, namely the review of genuinely high-risk, prioritised cases, rather than the exhaustive manual processing of every case regardless of risk.

4. AI-Driven Technologies Applied to KYC and AML

A number of distinct but complementary AI techniques are now embedded across the KYC/AML lifecycle. This section outlines the five technologies identified as most material to FinTech compliance operations.

4.1 Machine Learning for Risk Scoring

Supervised and unsupervised machine-learning models replace static, rule-based thresholds with dynamic risk scores learned from historical patterns of legitimate and suspicious behaviour. Because these models can weigh dozens of variables simultaneously — transaction velocity, counterparty network structure, geographic risk, and behavioural history — they are able to flag subtler patterns of potential money laundering that a fixed rule would miss, while also suppressing alerts on transactions that merely resemble a rule's trigger condition without being genuinely suspicious.

4.2 Natural Language Processing (NLP)

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

NLP techniques allow institutions to process unstructured text at scale — adverse media articles, regulatory enforcement notices, sanctions-list entries, and free-text transaction narratives — to identify relevant risk signals about a customer or counterparty. This is particularly valuable for ongoing monitoring, where new adverse information about an existing customer may emerge well after onboarding was first completed.

4.3 Automated Document Verification

Optical character recognition (OCR) combined with computer vision enables automatic extraction and validation of data from identity documents, alongside liveness detection to confirm that a submitted photograph corresponds to a genuine, present individual rather than a static image or a deepfake. Reported accuracy for automated data extraction in commercial deployments exceeds 98 per cent (Delos Intelligence, 2026), substantially reducing manual re-keying and the transcription errors that accompany it.

4.4 Anomaly Detection

Unsupervised anomaly-detection models establish a behavioural baseline for each customer and flag statistically significant deviations from that baseline, rather than comparing every transaction against a single fixed threshold applied uniformly across all customers. This behavioural, rather than rule-based, approach is central to the false-positive reductions documented in Section 2, since a transaction that is unusual for one customer may be entirely normal for another.

4.5 Predictive Analytics

Predictive models extend beyond flagging past or present anomalies to estimating the likelihood of future risk events — for instance, the probability that a given customer relationship will require escalation within a defined period. This allows compliance teams to allocate review resources proactively, concentrating senior analyst attention on relationships most likely to require it rather than distributing review effort uniformly across the customer base.

5. Impact on Operational Efficiency and Risk Management

The combined effect of the technologies described in Section 4 is visible in the efficiency metrics reported across multiple independent industry studies. Three effects stand out: a sharp reduction in false-positive alert volumes, a dramatic reduction in onboarding time, and widening adoption of AI across the core compliance functions.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

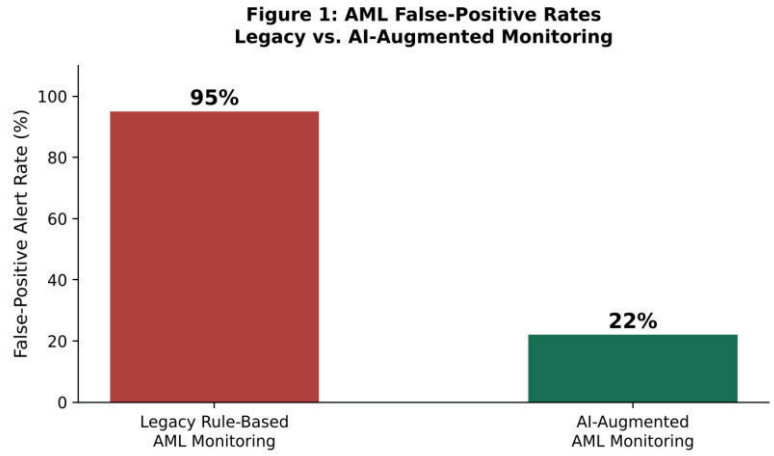


Figure 2. Representative AML false-positive alert rates before and after AI-augmented monitoring (Delos Intelligence, 2026; Youverify, 2026).

As Figure 2 shows, legacy rule-based AML monitoring commonly produces false-positive rates in the 90 to 99 per cent range, meaning that only a small fraction of flagged alerts represent genuine risk. AI-augmented monitoring, by contrast, has been reported to bring false-positive rates down to roughly 20 to 25 per cent in several independent deployments — a reduction that directly translates into recovered analyst capacity for substantive investigation work.

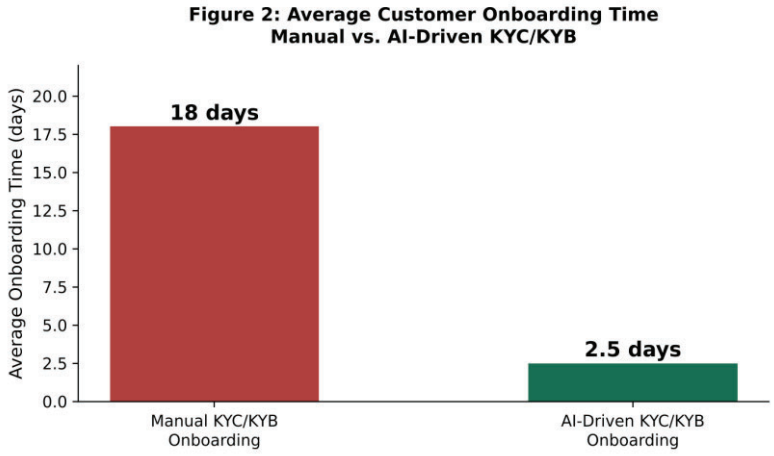


Figure 3. Average customer onboarding time, manual versus AI-driven KYC/KYB processes (Delos Intelligence, 2026; Wecan Group, 2026).

Onboarding speed shows a comparable pattern. Manual Know-Your-Business verification has been reported to take approximately eighteen business days on average, with more than half of business applicants abandoning the process before completion; AI-powered verification compresses this to roughly two-and-a-half days (Delos Intelligence, 2026). Separately, Deloitte's 2025 Financial Crime Technology Survey found that AI-based KYC reduced average time-to-

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

onboard by 60 to 70 per cent across both retail and small-business segments, with customer drop-off during onboarding falling by 22 to 31 per cent.

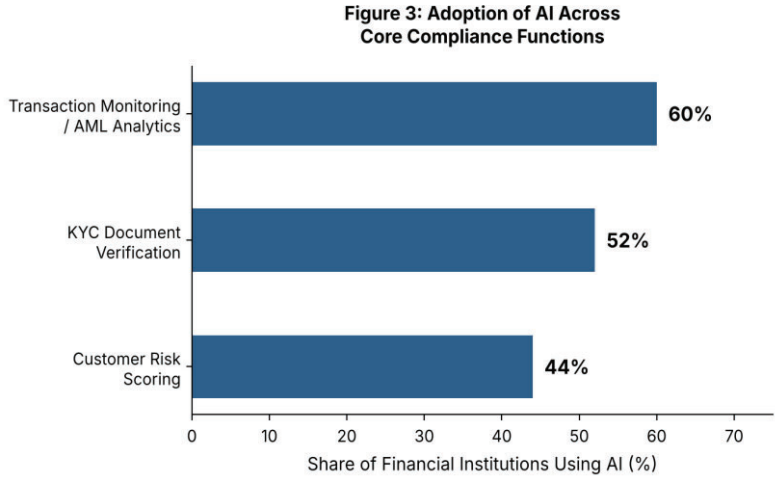


Figure 4. Share of financial institutions reporting active use of AI across core compliance functions (McKinsey, 2025; Deloitte, 2025).

Adoption data suggests this is no longer an early-stage experiment: AI-based transaction-monitoring analytics is already in use at over 60 per cent of surveyed financial institutions, document-verification automation at 52 per cent, and AI-based customer risk scoring at 44 per cent (StealthAgents Research, 2026, citing McKinsey and Deloitte, 2025).

Table 1 summarises the principal efficiency metrics drawn from the sources reviewed in this paper.

Metric	Before AI (Legacy / Manual)	After AI Adoption	Source
AML alert false-positive rate	90% – 99%	**20% – 25%**	Delos Intelligence, 2026; Youverify, 2026
Standard KYC onboarding time	15 – 21 days	**2 – 3 hours**	Wecan Group, 2026
KYB (business) onboarding time	~18 business days	**~2.5 days**	Delos Intelligence, 2026
Complex KYC case completion	8 – 12 business days	**Under 2 hours**	Delos Intelligence, 2026
Time-to-onboard (retail / SME)	Baseline	**60% – 70% faster**	Deloitte, 2025
Total financial-crime compliance cost	Baseline	**20% – 35% lower**	McKinsey, 2025
Onboarding customer drop-off rate	Baseline	**22% – 31% lower**	Celent, 2025

Beyond speed, AI-driven risk management also improves the quality of risk coverage. Because behavioural models continuously adapt to new data, they are better positioned than static rule sets to detect novel laundering typologies as they emerge, rather than only the patterns anticipated at the time rules were written. This adaptability is particularly valuable in FinTech settings, where product innovation (for example, embedded finance, buy-now-pay-later, and crypto-linked payment rails) regularly creates new channels that legacy rule sets were never designed to cover.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

6. Challenges and Risks in AI-Driven Compliance

The efficiency gains documented in Section 5 do not come without cost. Deploying AI within a regulated compliance function introduces its own category of operational and governance risk, five of which merit particular attention.

- Data privacy: KYC/AML systems process highly sensitive personal and financial data, and training or operating AI models on this data expands the volume of personal information processed, stored, and potentially exposed, raising obligations under data-protection regimes such as the GDPR and India's Digital Personal Data Protection Act.
- Algorithmic bias: Risk-scoring models trained on historical data can inherit and amplify biases present in that data, potentially resulting in disproportionate flagging or friction for customers from particular demographic, geographic, or socio-economic groups — a risk with direct implications for fair-lending and non-discrimination obligations.
- Explainability: Regulators generally require that a suspicious-activity determination, or a decision to decline or restrict a customer relationship, be justifiable in terms a human regulator can evaluate; complex machine-learning models, particularly deep-learning architectures, can be difficult to explain in these terms, creating tension between model sophistication and regulatory auditability.
- Cybersecurity: Centralising identity verification and risk data within AI systems creates a high-value target for attackers, and the same generative technologies that improve document verification (for example, computer vision) can also be used adversarially to produce convincing synthetic identities or deepfake documents designed to defeat automated checks.
- Over-reliance and reduced human oversight: An institution that treats AI output as determinative rather than advisory risks a loss of institutional compliance expertise over time, and risks systemic blind spots if a single model's weaknesses go undetected across a large share of the customer base.

None of these risks is a reason to avoid AI adoption outright; each is, instead, a design requirement. The following section proposes a framework intended to capture the efficiency benefits documented above while containing each of these five risk categories through deliberate human-oversight design.

7. Proposed Integrated AI-Human Oversight Framework

Building on the risk-tiering logic already used in several of the deployments reviewed in Section 2, this paper proposes a four-layer framework in which AI performs the high-volume analytical work of risk detection while human analysts retain authority over judgment-intensive and high-risk decisions.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

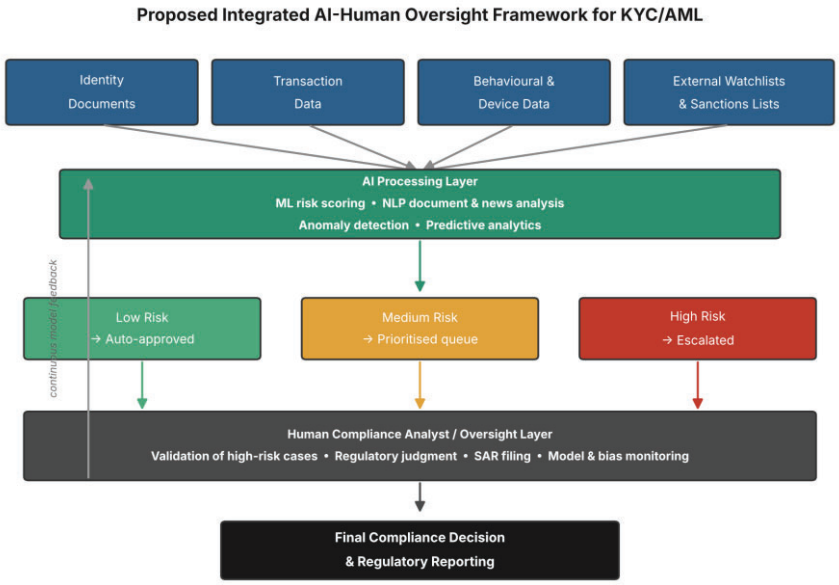


Figure 5. Proposed integrated AI–human oversight framework for KYC/AML compliance.

In the first layer, identity documents, transaction data, behavioural and device signals, and external watchlist data are ingested from across the customer relationship. In the second layer, the AI processing layer — combining the machine learning, NLP, anomaly-detection, and predictive-analytics techniques described in Section 4 — scores each case for risk. In the third layer, cases are triaged into low, medium, and high-risk categories rather than being treated uniformly: low-risk cases may be auto-approved, medium-risk cases are queued for prioritised review, and high-risk cases are escalated directly to senior analysts.

The fourth layer is the human compliance oversight function, which performs four roles that this paper argues should never be fully automated: validating high-risk and escalated cases before a final decision is recorded; exercising regulatory judgment in ambiguous cases that fall outside a model's training distribution; filing Suspicious Activity Reports, which carry legal accountability that should rest with a named individual; and monitoring the AI models themselves for drift, bias, and degraded performance over time. A continuous feedback loop from this human-oversight layer back into the AI processing layer — shown as the dashed path in Figure 5 — allows analyst decisions on escalated and disputed cases to retrain and recalibrate the underlying models, so that the system improves over time rather than remaining static.

This framework directly addresses the five risks identified in Section 6: data-minimisation and access controls can be enforced at the ingestion layer; bias is monitored explicitly as part of the human-oversight function rather than left implicit; explainability requirements are concentrated on the medium- and high-risk cases that actually reach a human reviewer, rather than demanded uniformly of every automated decision; cybersecurity controls can be layered around the centralised AI processing stage; and human judgment is preserved precisely at the points —

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

escalation, regulatory reporting, and model governance — where accountability law and prudent practice both require it.

8. Industry Adoption and Outlook

Adoption of AI-driven KYC and AML technology is already broad-based rather than confined to a handful of pioneering institutions. Survey evidence places AI use in at least one core AML or KYC function at more than 60 per cent of financial institutions (McKinsey, 2025, cited in StealthAgents Research, 2026), and vendor case studies report consistent double-digit-to-triple-digit percentage improvements in false-positive rates and onboarding speed across banks, FinTech lenders, and payment platforms operating in different regulatory jurisdictions, from the United Kingdom to South Africa (Youverify, 2026).

At the same time, the scale of the underlying problem remains large: the Nasdaq Global Financial Crime Report estimated that approximately USD 3 trillion in illicit funds moved through the global financial system in 2023, with fraud losses projected at roughly USD 485 billion, while the total global cost of financial-crime compliance has been estimated by LexisNexis Risk Solutions at more than USD 180 billion annually (cited in Youverify, 2026; Wecan Group, 2026). Compliance has also become a growing share of institutional overhead: compliance costs now average close to 19 per cent of annual revenue at some institutions, and senior bank executives report spending 42 per cent of their time on compliance matters, up from 24 per cent in 2016 (Youverify, 2026, citing the Bank Policy Institute).

Against this backdrop, the business case for AI-driven KYC/AML extends well beyond cost reduction. Fenergo's 2025 Financial Crime Industry Trends report found that seventy per cent of financial institutions lost clients in the preceding year due to slow onboarding — the highest such figure on record — indicating that onboarding speed, long treated as a secondary operational metric, has become directly material to customer retention and competitive positioning in the FinTech sector (Youverify, 2026). This suggests that AI adoption in compliance functions is likely to continue expanding, not purely as a cost-efficiency measure, but as a customer-experience and competitive necessity.

9. Conclusion and Future Scope

This paper has examined how AI-driven technologies — machine learning, natural language processing, automated document verification, anomaly detection, and predictive analytics — are reshaping KYC and AML compliance within the FinTech ecosystem. The evidence reviewed indicates substantial, consistently documented gains in operational efficiency: false-positive alert rates falling from the 90–99 per cent range common under legacy rule-based systems to roughly 20–25 per cent under AI-augmented monitoring, and onboarding timelines compressing from weeks to hours in several independently reported deployments.

These gains, however, are inseparable from a parallel set of governance obligations around data privacy, algorithmic bias, explainability, and cybersecurity. The framework proposed in Section 7

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 11, CONFLUENCE - 2026

argues that the appropriate response is not to choose between automation and human judgment, but to architect their relationship deliberately — using AI to triage volume and surface risk, while reserving final judgment on escalated, ambiguous, and legally consequential decisions for accountable human compliance professionals.

Future research in this area could usefully move beyond the vendor- and industry-reported efficiency metrics drawn upon in this paper toward independently audited, peer-reviewed evaluation of AI-driven KYC/AML systems — in particular, rigorous testing of these systems for demographic bias across the specific populations a given institution serves, and comparative study of explainability techniques that could reconcile the complexity of modern risk-scoring models with regulators' auditability requirements. As FinTech products continue to diversify into embedded finance, decentralised payment rails, and cross-border instant payments, the capacity of AI-driven compliance systems to adapt to genuinely novel risk typologies — rather than only to the typologies anticipated when a model was last trained — will likely determine how much of the promise documented in this paper is realised in practice.

References

Delos Intelligence. (2026). AI-powered KYB and customer onboarding: How banks and fintechs cut onboarding from weeks to hours while staying compliant. Delos Intelligence Blog.

Delos Intelligence. (2026). AI-powered AML and KYC screening: How financial institutions slash false positives by 75% and accelerate customer verification. Delos Intelligence Blog.

StealthAgents Research. (2026). AI KYC and AML automation statistics 2026, citing McKinsey & Company (2025) and Deloitte (2025). StealthAgents Research.

Wecan Group. (2026). KYC/AML compliance ROI: How AI reduces costs and processing time. Wecan Group Insights.

Youverify. (2026). Compliance automation: How banks are using AI to cut costs, reduce risk, and meet regulatory demands, citing Bank Policy Institute, Nasdaq Global Financial Crime Report (2024), and Fenergo Financial Crime Industry Trends Report (2025). Youverify Blogs.

Youverify. (2026). AML false positives: A practical guide for UK banks and fintechs under FCA guidance, citing Napier AI research. Youverify Blogs.

IndataLabs. (2026). Fraud detection and KYC automation: Where AI agents cut investigation time, citing Facctum (2026) and Sanction Scanner (2026).

AML Intelligence. (2026). Taktile launches 'Agentic AML Solution' platform, touts 75% reduction in false positives. AML Intelligence News.

PhotonPay. (2025). AI for KYC: Smart automation for global compliance. PhotonPay Industry Insights.

M, Nitin. (2026). AI-Driven KYC and AML Solutions: Improving Efficiency and Risk Management in FinTech [Conference abstract]. BMS College of Commerce and Management, Bengaluru.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)