pISSN:2349-5707 Volume:12, Issue 09, September-2025 eISSN:2349-5715

Artificial Intelligence and Cloud-Enabled Anti-Money Laundering: A Comprehensive Framework for Detection, Compliance, and Policy Optimization

Alexei M. Roy
Department of Financial Technologies, University of Lisbon

ABSTRACT:

Background: The accelerating sophistication of money laundering schemes, combined with the proliferation of fintech, digital remittances, and high-velocity transaction systems, has rendered legacy anti-money laundering (AML) controls insufficient. Modern AML demands an integrated approach that leverages advances in artificial intelligence (AI), cloud computing, machine learning (ML), and robust governance structures to detect, investigate, and prevent illicit financial flows in near real time (Agorbia-Atta & Atalor, 2024; Faccia et al., 2020).

Objective: This article develops a comprehensive, publication-quality framework for AI and cloud-enabled AML systems that harmonizes technical detection methods, explainability and legal constraints, operational processes, and policy optimization to strengthen financial institutions' compliance posture and investigative effectiveness.

Methods: Drawing strictly from the provided literature, the study synthesizes evidence from systematic reviews, empirical studies, technical reports, and legal analyses to construct a layered methodological approach. The framework integrates supervised and unsupervised ML techniques, deep learning anomaly detection, ensemble modeling, transaction monitoring architectures, cloud deployment strategies, automated robotic process automation (RPA) for compliance workflows, and policy optimization via reinforcement learning and rule-tuning methodologies (Alsuwailem & Saudagar, 2020; Dalal & Rele, 2018; Paula et al., 2016; Agorbia-Atta & Atalor, 2024; Singh, 2025).

Results: The paper presents (1) a taxonomy of laundering vectors in the digital era and their signal characteristics; (2) a modular system architecture combining data ingestion, feature engineering, model training, scoring, explainability modules, and case management; (3) algorithmic strategies to balance detection sensitivity and false positive rates; (4) cloud deployment and operationalization recommendations to support scalability and cross-border data handling; and (5) a policy optimization blueprint for aligning ML outputs with regulatory expectations. Each element is elaborated in detail, with practical recommendations for tuning, governance, auditability, and human–machine collaboration.

Conclusions: AI and cloud technologies offer transformative potential for AML, but realization depends on careful system design, legal and ethical safeguards, model transparency, continual policy feedback loops, and investment in investigatory capacity. The proposed framework reconciles technical capabilities with regulatory obligations and provides a roadmap for institutions seeking to modernize AML operations while minimizing harms from erroneous interventions.

Keywords: Anti-Money Laundering, Artificial Intelligence, Cloud Computing, Anomaly Detection, Compliance Optimization, Explainable AI, Transaction Monitoring

INTRODUCTION

Money laundering is a persistent and evolving threat to global financial integrity, public safety, and economic stability. The last decade has seen substantial shifts in laundering modalities: the migration of illicit flows into electronic payment rails, the rise of cross-border digital remittances, the criminal exploitation of decentralized fintech platforms, and the increasing use of layered transaction chains to obscure provenance (Faccia et al., 2020; Agorbia-Atta & Atalor, 2024). Traditional rule-based AML

systems—often reliant on threshold triggers, static rules, and manual reviews—are strained by volume, velocity, and complexity, generating an avalanche of false positives and straining investigation resources (Rohit & Patel, 2015; Alsuwailem & Saudagar, 2020). Concurrently, regulators and international bodies urge financial institutions to adopt risk-based approaches and technological innovation to detect and prevent illicit flows (FATF, 2014; Steiner & Pol, 2019).

The integration of AI, machine learning, and cloud computing into AML workflows presents both opportunity and challenge. Machine learning models, including supervised classifiers, unsupervised anomaly detectors, and deep neural networks, can identify complex behavioral patterns and adapt to evolving threats (Paula et al., 2016; Dalal & Rele, 2018). Cloud infrastructure enables scalable processing of high-velocity transaction streams, rapid model retraining, and cross-jurisdictional data sharing when permitted (Agorbia-Atta & Atalor, 2024; Josyula, 2024). However, AI adoption raises concerns regarding explainability, bias, legal accountability, data privacy, and the potential for automation to erode due process or generate discriminatory outcomes (Turksen, Benson & Adamyk, 2024; Oubari & Leontjeva, 2024). There is, therefore, a clear need for an integrative framework that articulates how technical components, operational processes, governance controls, and policy optimization mechanisms can be combined to realize AI's potential while respecting legal constraints.

This article responds to that need by synthesizing the extant literature supplied by the commissioning request to build a comprehensive, practical, and theoretically informed framework for AI and cloud-enabled AML systems. The central premise is that an effective system is not a single model or tool, but an orchestrated ecosystem of data, algorithms, human investigators, cloud platforms, compliance processes, and regulatory interfaces. The framework advances three core contributions: (1) a detailed taxonomy mapping laundering techniques to detectable transaction features; (2) an end-to-end technical architecture for real-time and batch AML monitoring using AI and cloud resources; and (3) a methodology for policy optimization that balances detection performance with operational and legal constraints. Each contribution is grounded in empirical and conceptual findings from the provided references and elaborated with practical guidance for deployment, governance, and continuous improvement.

METHODOLOGY

The methodological approach followed in this paper is a rigorous synthesis and analytic integration of the provided literature, structured to produce actionable system design prescriptions and normative guidance. The method comprises four interlocking strands: conceptual synthesis, technical mapping, architectural design, and policy optimization modeling. Each strand is described below.

Conceptual Synthesis. The study begins by extracting recurring themes, problem statements, and empirical findings from systematic reviews and domain studies. Significant contributions included the systematic review of AML systems (Alsuwailem & Saudagar, 2020), analyses of data mining frameworks for suspicious transaction detection (Rohit & Patel, 2015), and examinations of electronic money laundering in fintech contexts (Faccia et al., 2020). These sources provide foundational knowledge about the weaknesses of legacy AML, the forms of suspicious activity, and the institutional drivers for technological change (Alsuwailem & Saudagar, 2020; Faccia et al., 2020; Rohit & Patel, 2015). Additionally, recent works describing the strategic use of AI and cloud technologies informed high-level strategic choices (Agorbia-Atta & Atalor, 2024; Zhang & Chen, 2024).

Technical Mapping. Building on conceptual synthesis, the method maps common laundering behaviors to measurable transaction features, customer lifecycle events, and network relationships. This mapping derives from anomaly detection research and practical AML implementations that have demonstrated the utility of feature engineering and representation learning (Paula et al., 2016; Dalal & Rele, 2018). The mapping identifies candidate supervised labels (SARs, confirmed cases), unsupervised anomaly signals (novelty detection scores), and relational features (graph centrality, temporal motifs). It also codifies the typical data sources required for robust detection: transaction records, account metadata, customer due diligence (CDD), sanctions lists, device and channel telemetry, and external threat intelligence (Agorbia-Atta & Atalor, 2024; Josyula, 2024).

Architectural Design. The architecture strand synthesizes cloud design patterns, model lifecycle processes, and compliance workflows into a cohesive system architecture suited to modern AML use cases. This design emphasizes modular components—ingestion, normalization, feature store, model training and

deployment, scoring service, explainability and compliance layer, and case management—building on practical suggestions from RPA and AI deployment literature (Pingili, 2024; Balakrishnan, 2024). Special attention is given to data governance, secure cross-border processing, and audit logging to meet regulatory expectations (Turksen et al., 2024; Oubari & Leontjeva, 2024).

Policy Optimization Modeling. To reconcile model outputs with regulatory and operational realities, the framework describes methods for continuous policy optimization. This includes supervised threshold tuning, cost-sensitive learning to account for investigator effort and legal risk (Singh, 2025), reinforcement learning paradigms for adaptive rule tuning (Singh, 2025; Kaur, 2023), and human-in-the-loop mechanisms to correct model drift. Legal and ethical constraints—such as explainability mandates and data protection rules—are embedded into the optimization loop, following legal analyses of automated monitoring and AI governance (Turksen et al., 2024; Oubari & Leontjeva, 2024).

Validation Approach. Given the study's synthetic and prescriptive nature, validation is conceptual and comparative rather than empirical. The proposed framework's internal coherence is assessed by mapping its components onto successful elements reported in the literature—e.g., deep learning anomaly detection success in Brazilian export fraud investigations and supervised ML improvements in transaction monitoring (Paula et al., 2016; Zhang & Chen, 2024). Sensitivity analysis techniques for thresholding and cost modeling are described to assist institutions in locally validating model parameters using their proprietary data (Singh, 2025; Balakrishnan, 2024).

Ethical and Legal Integration. Across all methodological strands, legal and ethical considerations are treated as constraints shaping permissible algorithmic choices. This integration reflects normative guidance in the legal literature and policy analyses, emphasizing transparency, auditability, accountability, and proportionality (Turksen et al., 2024; Oubari & Leontjeva, 2024).

Throughout, every major claim and prescriptive recommendation is rooted in the cited literature to ensure fidelity to the provided evidence base and to maintain the strict constraint that this article be generated "based strictly on the References provided."

RESULTS

The results comprise the synthesized artifact—the comprehensive framework—and its component outputs: the laundering taxonomy, the modular system architecture, algorithmic recommendations, cloud deployment strategies, and the policy optimization blueprint. Each result is described in detailed prose, illustrating how it operationalizes findings from the literature.

Taxonomy of Laundering Vectors and Signal Characteristics. A foundational contribution is a taxonomy that organizes laundering techniques into detectable vectors and enumerates their typical signal characteristics. The taxonomy synthesizes observations from AML reviews, fintech laundering case studies, and anomaly detection research (Alsuwailem & Saudagar, 2020; Faccia et al., 2020; Paula et al., 2016).

1.Structuring and Smurfing. Characterized by numerous small transactions below reporting thresholds, often across many accounts or channels. Signals include repeated transactions at similar amounts, temporal clustering, and rapid fragmentation of funds. Detectable via aggregate features (transaction count, mean transaction value), temporal pattern analysis, and cohort comparisons (Rohit & Patel, 2015; Alsuwailem & Saudagar, 2020).

2.Layering through Cross-Border Transfers. Involves repeated movements across jurisdictions, sometimes exploiting remittance corridors or multiple foreign exchange conversions. Signals include rapid sequence of outbound and inbound transfers, inconsistency between declared purpose and flow structure, and use of high-risk correspondent banks. Detection requires enriched wire data, geographic tagging, and integration with remittance platforms (Faccia et al., 2020; Agorbia-Atta & Atalor, 2024).

3.Trade-Based Laundering. Uses misinvoicing, over/under-invoicing, or phantom shipments to move value. Signals are often outliers in unit pricing, mismatched shipping and billing entities, or unusually complex supply chains. Detection benefits from deep learning anomaly detection on structured trade data and text mining of invoices and shipping manifests (Paula et al., 2016; Faccia et al., 2020).

4.Virtual Asset and Fintech Exploitation. Involves cryptocurrencies, prepaid payment instruments, mobile

money, and fintech aggregation platforms. Signals include frequent on-chain-to-off-chain movements, transfers to mixing services, use of multiple wallets/accounts, and sudden spikes in conversion activity. Cloud-native monitoring and blockchain analytics tools are essential (Faccia et al., 2020; Agorbia-Atta & Atalor, 2024).

5.Account Takeover and Synthetic Identities. Fraudsters use stolen identities or fabricated profiles to open accounts and route illicit funds. Signals include inconsistencies in KYC data, device fingerprint anomalies, and rapid profile changes. Detection demands robust identity analytics, device telemetry, and cross-channel behavior modeling (Josyula, 2024; Balakrishnan, 2024).

6.Circular Transactions and Layered Networks. Criminals craft circular flows or complex multi-party networks to obfuscate the origin. Signals include repeating cyclical transaction patterns and high clustering coefficients in transactional graphs. Graph analytics and temporal motif detection are effective (Paula et al., 2016; Dalal & Rele, 2018).

For each vector, the taxonomy links recommended features and algorithmic approaches, enabling practitioners to choose appropriate detection techniques and prioritize data collection.
Modular System Architecture. The paper proposes a modular architecture aligned with cloud-native design principles and AI lifecycle management practices (Agorbia-Atta & Atalor, 2024; Pingili, 2024). The architecture's modules are:
1.Data Ingestion Layer. Streams and batches transaction records, KYC, sanctions lists, device telemetry, and external threat feeds into a secure raw data lake. Emphasizes schema versioning, encryption at rest, and PII minimization. This layer supports both high-velocity wire data and less frequent batch data like invoices (Agorbia-Atta & Atalor, 2024; Josyula, 2024).

2.Normalization and Enrichment. Applies cleaning, canonicalization, geo-resolution, AML-specific enrichment (PEP/sanctions matching), and risk scores. Natural language processing (NLP) is used to extract entities from free text fields such as remittance reasons and invoice descriptions (Paula et al., 2016; Zhang & Chen, 2024).

3.Feature Store. Stores reusable features with temporal validity to support consistent training and serving. Features include per-account aggregate metrics, velocity measures, graph features (degree, centrality), device fingerprints, and derived sentiment or text features (Dalal & Rele, 2018; Balakrishnan, 2024).

4.Model Training and Validation. Supports multiple learning paradigms: supervised classification (using historical SARs), unsupervised anomaly detection, semi-supervised approaches for scarce labeled data, and representation learning for feature extraction. Cross-validation, backtesting on historical windows, and drift detection are integrated (Paula et al., 2016; Dalal & Rele, 2018; Singh, 2025).

5.Model Registry and Deployment. Models are versioned, stored with metadata, and deployed as containerized scoring services. A/B testing and canary deployments allow safe rollouts. Model explainability snapshots are recorded with model versions to meet audit requirements (Turksen et al., 2024; Oubari & Leontjeva, 2024).

6.Scoring and Alerting Engine. Produces risk scores and triage levels; alerts feed into an automated case management queue. Scoring is accompanied by explainability artifacts—feature attributions, exemplar comparisons, and rule matches—to assist investigators (Paula et al., 2016; Turksen et al., 2024).

7.Explainability and Compliance Layer. Implements XAI techniques (feature importance, local explanations), compliance rules engine, and decision logs. Ensures that each alert includes an audit trail linking inputs, model versions, and investigator actions (Turksen et al., 2024; Oubari & Leontjeva, 2024).

8.Case Management and Investigation Workspace. Consolidates alerts, provides workflow automation (including RPA for routine document retrieval), and supports analyst feedback into the training pipeline.

Enables investigators to escalate to law enforcement or file SARs while capturing metadata for continuous learning (Pingili, 2024; Balakrishnan, 2024).

9.Governance and Monitoring. Monitors model performance metrics (precision, recall, false positive rate), system health, and compliance KPIs. Integrates privacy controls, role-based access, and legal hold mechanisms (Agorbia-Atta & Atalor, 2024; Turksen et al., 2024).

Algorithmic Recommendations and Trade-Offs. The literature indicates that no single algorithm dominates across AML contexts. Rather, ensemble approaches and hybrid pipelines outperform monolithic models (Paula et al., 2016; Dalal & Rele, 2018). Specific recommendations include:
1.Use supervised models (tree ensembles, gradient boosting) where quality labeled SAR data exist, but always complement them with unsupervised detectors to catch novel attack types (Paula et al., 2016; Dalal & Rele, 2018).

2.Apply deep representation learning (autoencoders, variational methods) to learn robust embeddings of transaction sequences and to detect subtle anomalies in high-dimensional data (Paula et al., 2016).

3.Employ graph neural networks (GNNs) or graph analytics to detect network-based laundering schemes and to surface suspicious clusters and intermediary entities (Dalal & Rele, 2018; Faccia et al., 2020).

4.Integrate cost-sensitive learning and calibration that explicitly incorporates the operational cost of false positives and the legal risk of missing true positives; this can be implemented via custom loss functions or post-hoc threshold optimization (Singh, 2025; Balakrishnan, 2024).

5.Maintain human-in-the-loop review and feedback loops to remediate bias, validate high-impact alerts, and provide labels for continual learning (Pingili, 2024; Kaur, 2023).

Cloud Deployment and Operationalization. Cloud platforms provide elasticity and managed services that are particularly valuable for AML workloads with bursty demand and heavy compute requirements (Agorbia-Atta & Atalor, 2024). Best practices from the literature include:
1. Hybrid Cloud for Data Residency. Use hybrid architectures to keep sensitive customer data within required jurisdictions while leveraging cloud compute and analytics where permissible (Agorbia-Atta & Atalor, 2024; Josyula, 2024).

2.Secure Key Management and Encryption. Employ robust encryption and key management practices, segregating duties and logging decryption events to satisfy auditors (Agorbia-Atta & Atalor, 2024).

3.Scalable Stream Processing. Use streaming platforms to support near real-time scoring for high-priority transaction types, and batch processing for periodic risk repricing (Zhang & Chen, 2024; Dalal & Rele, 2018).

4.Containerization and Orchestration. Containerize model scoring services and use orchestration to support automated scaling and rapid recovery (Pingili, 2024).

5.Data Lineage and Observability. Maintain full data lineage and observability for every model input and output to support explainability and regulatory inquiries (Turksen et al., 2024; Oubari & Leontjeva, 2024).

Policy Optimization Blueprint. The paper describes an iterative policy optimization loop that aligns model behavior with compliance objectives and legal constraints (Singh, 2025; Turksen et al., 2024). The loop cycles through:
1.Define Objectives and Constraints. Clarify operational KPIs (true positive rate, analyst workload), legal constraints (data protection, non-discrimination), and business outcomes (fraud losses, SAR filing timeliness) (Singh, 2025; Turksen et al., 2024).

2.Cost Modeling. Quantify costs for missed detection, false positives, investigation processing, and reputational risk. Use these costs in training objectives or for threshold calibration (Singh, 2025; Balakrishnan, 2024).

3.Simulation and Backtesting. Simulate model outputs and policy actions over historical windows to estimate impacts on KPIs and regulatory thresholds (Paula et al., 2016; Singh, 2025).

4.Human Feedback Integration. Integrate analyst labels and qualitative feedback to refine feature sets and model objectives; maintain a human-review buffer for high-risk decisions to satisfy legal checks (Pingili, 2024; Turksen et al., 2024).

5.Regulatory and Ethical Audit. Conduct periodic audits of model logic, data usage, explainability reports, and fairness assessments; adjust policies to reflect regulatory guidance (Oubari & Leontjeva, 2024).
Collectively, the architecture and policy loop produce a resilient AML ecosystem capable of adapting to new laundering tactics while providing auditability and legal defensibility.

DISCUSSION

The proposed framework synthesizes the literature into an actionable roadmap but introduces several tensions and practical considerations. This section interrogates those tensions, discusses limitations, and outlines future research and operational recommendations.
Balancing Detection Performance and False Positives. One of the most persistent operational challenges in AML is the trade-off between sensitivity (catching true illicit flows) and specificity (avoiding overwhelming false positives). Legacy rule-based systems were safe but inefficient, while naïve high-sensitivity ML models can saturate investigators with false alerts (Rohit & Patel, 2015; Alsuwailem & Saudagar, 2020). The literature recommends cost-sensitive optimization and hybrid models to manage this trade-off (Singh, 2025; Paula et al., 2016). A principled approach quantifies the marginal utility of additional true positives against the incremental investigation burden and reputational/legal risk of false positives. Operationally, institutions should implement graduated triage—high-confidence alerts trigger immediate action while lower-confidence scores pass through enrichment and secondary models before analyst review.
Human–Machine Collaboration and Investigator Capacity. AI tools do not eliminate the need for skilled investigators; rather, they change the nature of investigative work. The literature stresses automation for routine tasks (via RPA) to free investigators for complex analysis (Pingili, 2024). However, to be effective, AI systems must present clear, actionable explainability outputs and integrate seamlessly into investigators' workflows (Turksen et al., 2024). Training programs must evolve to build investigator proficiency in interpreting model outputs, understanding false positive patterns, and providing feedback for model refinement (Balakrishnan, 2024). Institutions should monitor analyst workload and model-to-investigator ratios as key governance metrics.
Explainability, Accountability, and Legal Constraints. The adoption of opaque models like deep neural networks raises legal and ethical issues. Regulators increasingly demand transparency regarding automated decisions that impact customers (Turksen et al., 2024; Oubari & Leontjeva, 2024). The proposed framework centers explainability modules and detailed decision logs to meet auditability and fairness assessments. Importantly, explainability does not mean simplification; instead, institutions must ensure that explanations are accurate, actionable, and legally sufficient—anchored in model versioning, input snapshots, and causal attributions where possible.
Data Quality, Label Scarcity, and Concept Drift. High-quality labeled data are often scarce in AML contexts because confirmed SARs are rare relative to transaction volumes and because labeling is costly (Paula et al., 2016). The literature suggests semi-supervised learning, transfer learning, and unsupervised anomaly detection as remedies (Paula et al., 2016; Dalal & Rele, 2018). Moreover, laundering methods evolve rapidly, producing concept drift; therefore, continuous monitoring, incremental retraining, and domain adaptation techniques are necessary. The framework advocates for automated drift detection, periodic retraining windows, and human feedback loops to update models effectively (Singh, 2025).
Cross-Border Data Sharing and Privacy. Effective detection of cross-border laundering often requires

sharing data across jurisdictions; yet data protection laws and sovereignty concerns limit such sharing (Agorbia-Atta & Atalor, 2024; Josyula, 2024). Hybrid cloud architectures and privacy-preserving techniques—such as secure multi-party computation, federated learning, and differential privacy—offer promising paths to collaborative detection while minimizing exposure of raw personal data. However, the literature shows these techniques are emergent and carry trade-offs in complexity and performance; implementation should be carefully scoped with legal input (Agorbia-Atta & Atalor, 2024; Oubari & Leontjeva, 2024).

Fintech, Virtual Assets, and Emerging Attack Surfaces. The rapid emergence of virtual asset service providers (VASPs), mobile money, and new payment primitives has created novel laundering vectors (Faccia et al., 2020; Agorbia-Atta & Atalor, 2024). Blockchain analytics and entity resolution across off-chain/on-chain datasets are critical for detecting virtual asset laundering. The literature highlights the need to integrate blockchain data enrichment and to update feature sets rapidly to reflect new platform behaviors (Faccia et al., 2020; Zhang & Chen, 2024).

Governance, Ethics, and Regulatory Alignment. Legal scholars highlight the importance of aligning automated AML with evolving AI governance frameworks and proposed regulatory instruments (Turksen et al., 2024; Oubari & Leontjeva, 2024). This requires organizations to embed legal constraints into model training, maintain transparent documentation, and engage proactively with regulators to define acceptable practices. The framework prescribes multidisciplinary governance councils composed of compliance, legal, data science, and operations to oversee model deployment and policy optimization.

Limitations of the Framework. While comprehensive, the framework has several limitations. First, its prescriptions are necessarily generic because institutions vary significantly in data availability, regulatory landscape, and risk appetite. Second, the framework relies on technical advances (e.g., robust federated learning) that, while promising, may be difficult to operationalize at scale today. Third, the validation is conceptual and reliant on reported case studies rather than live deployment results; institutions must therefore undertake localized backtesting and pilot studies before broad rollouts. Finally, there are unresolved normative challenges around algorithmic fairness in AML—for example, ensuring that models do not unduly target protected classes because of historical biases in SAR filing patterns (Turksen et al., 2024).

Future Research Directions. The literature points to several promising research trajectories. These include the development of standardized evaluation datasets for AML to enable reproducible benchmarking (Paula et al., 2016), the application of GNNs for large-scale network detection in near real time (Dalal & Rele, 2018), research on privacy-preserving cross-institutional learning (Agorbia-Atta & Atalor, 2024), and interdisciplinary work on aligning AI explainability with legal disclosure requirements (Turksen et al., 2024; Oubari & Leontjeva, 2024). Further empirical studies documenting the operational impact of AI AML deployments—reduction in laundering losses, investigator productivity gains, and the effect on false positive rates—would significantly strengthen the field's evidence base.

CONCLUSION

This article has presented a comprehensive framework for AI and cloud-enabled anti-money laundering systems, synthesizing the provided literature to produce a practical yet theoretically grounded roadmap. Key insights include the necessity of modular system design, the efficacy of hybrid algorithmic approaches combining supervised, unsupervised, and graph-based techniques, and the importance of policy optimization that internalizes legal, operational, and ethical constraints. Cloud platforms enable the scale and agility needed for modern AML tasks, but they must be deployed with careful attention to data residency, privacy, and auditability. Human investigators remain central to the loop; automation should amplify their capacity rather than replace their judgment.

Adoption of the framework requires iterative pilots, localized validation, and robust governance structures that bring together compliance, legal, data science, and operations. Institutions that implement the recommendations will be better positioned to detect sophisticated laundering techniques, streamline investigations, and meet evolving regulatory expectations. Nevertheless, the field faces open challenges—particularly in cross-border data sharing, fairness, and operationalizing privacy-preserving collaboration—that warrant continued research and dialogue among practitioners, regulators, and academics.

Ultimately, confronting money laundering in the digital era is not solely a technological problem; it is a

socio-technical challenge requiring coordination across institutions, regulators, technology vendors, and law enforcement. The framework advanced here aims to equip organizations with both the conceptual understanding and practical tools to modernize AML capabilities responsibly and effectively.

REFERENCES

1. A. A. S. Alsuwailem and A. K. J. Saudagar, "Anti-money laundering systems: A systematic literature review", Journal of Money Laundering Control, vol. 23, no. 4, pp. 833-848, May 2020.
2. Agorbia-Atta, C., & Atalor, I. (2024). Enhancing anti-money laundering capabilities: The Strategic Use of AI and Cloud Technologies in Financial Crime Prevention. World Journal of Advanced Research and Reviews, 23(2), 2035-2047.
3. Balakrishnan, A. (2024). Leveraging artificial intelligence for enhancing regulatory compliance in the financial sector. International Journal of Computer Trends and Technology.
4. Dalal, K. R., & Rele, M. (2018). Cyber Security: Threat Detection Model based on Machine learning Algorithm. 2018 3rd International Conference on Communication and Electronics Systems (ICCES), Coimbatore, India, pp. 239-243.
5. Faccia, A., Moçteanu, N. R., Pio, L., Cavaliere, L., & Mataruna-DosSantos, L. J. (2020). Electronic money laundering the dark side of fintech: an overview of the most recent cases. ICIME 2020: Proceedings of the 2020 12th International Conference on Information, pp. 29-34, 16 September 2020.
6. FATF (2014). Guidance: Customer due diligence approach; Financial Action Task Force.
7. Josyula, H. P. (2024). Enhancing Security and Compliance. In Redefining Cross-Border Financial Flows: Transforming Remittances with AI and Other Technologies (pp. 49-64). Berkeley, CA: Apress.
8. Kaur, G. (2023). Trust the Machine and Embrace Artificial Intelligence (AI) to Combat Money Laundering Activities. In Computational Intelligence for Modern Business Systems: Emerging Applications and Strategies (pp. 63-81). Singapore: Springer Nature Singapore.
9. Oubari, Z., & Leontjeva, L. (2024). Maximizing Anti Money Laundering Compliance through AI: Assessing the Obligations and Responsibilities of Financial Institutions under the Proposed EU AI Act.
10. Paula, E. L., Ladeira, M., Carvalho, R. N., & Marzagão, T. (2016). Deep learning anomaly detection as support fraud investigation in Brazilian exports and anti-money laundering. Proceedings of the 15th IEEE International Conference on Machine Learning and Applications (ICMLA), Anaheim, CA, USA, 18–20 Dec. 2016, pp. 954-960.
11. Pingili, R. (2024). Transforming Anti-Money Laundering Compliance in Banking with AI-Driven Robotic Process Automation.
12. Rohit, K. D., & Patel, D. B. (2015). Review on detection of suspicious transaction in anti-money laundering using data mining framework. Journal for Innovative Research in Science and Technology, 1, 129–133.
13. Sadiya, H., & Shah, H. Predictive Analytics and AI Integration: Revolutionizing AML and Fraud Detection in Financial Services.
14. Singh, V. (2025). Policy Optimization for Anti-Money Laundering (AML) Compliance using AI Techniques: A Machine Learning Approach to Enhance Banking Regulatory Compliance. International Journal of Engineering Research & Technology (IJERT), 14.
15. Steiner, M., & Pol, A. (2019). Can artificial intelligence defeat financial intelligence? Public and Accounting Journal.
16. Turksen, U., Benson, V., & Adamyk, B. (2024). Legal implications of automated suspicious transaction monitoring: enhancing integrity of AI. Journal of Banking Regulation, pp. 1-19.
17. W. Zhu, L. Xu, W. Fan (2019). Anti-money laundering policy and public accounting. Journal of Policy.
18. Y. Liu, Y. Mao, M. Chen (2018). A survey: Big data for mobile networks and systems information. Systems Information and Knowledge.
19. Zhang, W., & Chen, L. (2024). Real-Time Transaction Monitoring Using AI: Detecting Suspicious Activities and Money Laundering in Banking. Asian American Research Letters Journal, 1(3).