

---

# AI-Augmented AML Workflow Optimization In High-Volume Financial Institutions

**Sai Vamsi Kiran Gummadi**

Independent researcher  
[svkiran.g@gmail.com](mailto:svkiran.g@gmail.com)

## Abstract

Anti-Money Laundering (AML) compliance is a cornerstone of financial integrity, yet traditional rule-based workflows in high-volume financial institutions often struggle to balance false positives, operational overhead, and dynamic regulatory requirements. This paper presents a scalable, AI-augmented AML workflow architecture that integrates machine learning for adaptive risk scoring, natural language processing (NLP) for unstructured data ingestion, and robotic process automation (RPA) for case handling. We demonstrate that our approach improves detection accuracy, reduces alert fatigue, and shortens investigative timelines, enabling institutions to meet regulatory expectations efficiently while optimizing resource allocation. A comparative evaluation on synthetic and real-world datasets validates the system's precision, recall, and operational efficiency. The proposed framework is practical, scalable, and impactful for both enterprise deployment and supervisory oversight.

**Index Terms:** AML, anti-money laundering, AI workflow optimization, financial compliance, NLP, RPA, machine learning, risk scoring, RegTech.

## I. INTRODUCTION

In recent years, the scale and sophistication of financial crime have grown dramatically, driven by the rapid digitization of banking services and the expansion of global financial networks. Money laundering alone is estimated to account for 2–5% of global GDP annually, posing significant threats to financial integrity, regulatory compliance, and national security [1]. Anti-Money Laundering (AML) frameworks have traditionally relied on rule-based transaction monitoring systems (TMS), which generate alerts based on predefined thresholds and heuristics. While foundational, these systems are often inefficient, generating a high volume of false positives and requiring significant human effort to investigate low-risk cases [2], [3].

High-volume financial institutions—such as global banks, payment service providers, and digital-first fintechs—face compounding challenges. These include managing ever-growing transaction volumes, integrating complex customer data across silos, and complying with evolving international regulations (e.g., FATF recommendations, EU AML directives, and the U.S. Bank Secrecy Act) [1], [6]. Consequently, there is an urgent need to modernize AML workflows using intelligent, adaptive, and scalable solutions.

Artificial Intelligence (AI) offers transformative potential in this domain. Techniques such as machine learning (ML), natural language processing (NLP), and robotic process automation (RPA) are increasingly being deployed to augment various stages of the AML workflow, including transaction monitoring, case triaging, and suspicious activity reporting (SAR) [2], [4], [5]. AI models can dynamically assess risk, adapt to novel laundering typologies, reduce alert fatigue, and automate

---

repetitive tasks—all while providing transparency and auditability essential for regulatory oversight [3], [10], [11].

This paper proposes a practical AI-augmented AML workflow architecture tailored for high-volume financial institutions. Our approach integrates supervised and unsupervised ML for risk scoring, NLP for unstructured data extraction, and RPA for operational efficiency. We evaluate the system on real-world and synthetic datasets, demonstrating improved precision, reduced false positives, and shortened case investigation cycles compared to legacy systems.

## II. BACKGROUND AND MOTIVATION

Anti-Money Laundering (AML) compliance plays a foundational role in the integrity of financial systems worldwide. It involves a range of operational, legal, and technical measures designed to detect, prevent, and report illicit financial activity, including the laundering of proceeds from crimes such as drug trafficking, terrorism financing, cybercrime, and corruption. Financial institutions are bound by regulatory mandates—such as the Financial Action Task Force (FATF) recommendations [1], the European Union’s AML Directives [2], and the United States Bank Secrecy Act (BSA) [3]—to maintain rigorous AML programs. These programs typically consist of four core stages: transaction monitoring, alert generation, case investigation, and suspicious activity reporting (SAR).

Historically, each of these stages has been handled in a siloed, sequential manner with limited automation and minimal intelligence. Transaction monitoring systems (TMS), for example, rely heavily on static, rule-based thresholds to flag potentially suspicious transactions, such as those exceeding certain value limits or involving high-risk geographies [4]. While straightforward to implement, such systems are prone to high false-positive rates—sometimes exceeding 90%—because they cannot adapt to contextual subtleties or evolving laundering strategies [5]. The subsequent stages of case investigation and SAR generation often require manual review by compliance analysts, who must aggregate information from disparate data sources, assess transaction behavior, and draft detailed narrative reports for regulators. This process is not only labor-intensive and time-consuming but also vulnerable to cognitive biases, incomplete information, and inconsistencies [6].

For high-volume financial institutions—such as multinational banks, digital payment platforms, and cryptocurrency exchanges—the scale of AML operations poses an even greater challenge. With millions of transactions processed daily, the volume of alerts generated by traditional systems can quickly become unmanageable. According to a 2022 report by the Financial Conduct Authority (FCA), large institutions spend over 70% of their AML resources on reviewing false alerts, leaving limited capacity for addressing genuine high-risk cases [7]. Furthermore, increasing regulatory scrutiny has pushed institutions to demonstrate both the effectiveness and explainability of their AML efforts, with global penalties for AML violations exceeding \$5 billion in recent years [8].

In this context, artificial intelligence (AI) presents a transformative opportunity to augment and optimize the AML workflow. AI encompasses a range of technologies—including supervised and unsupervised machine learning (ML), natural language processing (NLP), and robotic process automation (RPA)—that can be deployed across the AML value chain to enhance accuracy, efficiency, and compliance alignment [9]. These technologies form the backbone of a new class of regulatory technologies (RegTech) designed to meet modern AML challenges.

**Supervised machine learning** can be used to train predictive models on labeled transaction datasets, learning patterns associated with suspicious behavior. These models are capable of scoring transaction risk more precisely than rule-based systems, enabling institutions to prioritize alerts based on true threat levels [10]. **Unsupervised learning**, such as clustering or anomaly detection techniques, can uncover

---

previously unknown typologies or emerging criminal strategies not captured in predefined rules [11]. These models are particularly valuable in detecting structuring (smurfing), layering patterns, or rapidly shifting criminal tactics.

**Natural language processing (NLP)** plays a critical role in handling unstructured data within AML investigations. For example, customer due diligence (CDD) and Know Your Customer (KYC) documents, transaction narratives, open-source intelligence (OSINT), and SAR narratives are all rich in textual data. NLP algorithms, particularly those based on deep learning such as BERT, can extract entities, detect sentiment, and identify inconsistencies or red flags within these documents [12]. This capability significantly reduces the manual burden of data analysis and increases the precision of investigative findings.

**Robotic process automation (RPA)** complements AI models by automating rule-based tasks such as alert assignment, data retrieval from internal systems, document pre-population, and escalation procedures [13]. In many AML programs, RPA has been used to automate up to 30% of the compliance workload, reducing costs and shortening investigation cycles [14]. When AI and RPA are combined, institutions can achieve a streamlined workflow wherein alerts are automatically scored, enriched with contextual data, and routed to appropriate analysts with suggested actions—thereby accelerating the end-to-end case resolution process.

Motivating the adoption of AI in AML are several pressing factors:

- **Scalability:** AI systems can handle high transaction volumes with minimal degradation in performance.
- **Adaptability:** ML models can be retrained on new data to reflect changes in financial behavior or criminal tactics.
- **Explainability:** Techniques such as SHAP and LIME allow institutions to interpret AI decisions and provide transparency to regulators [15].
- **Compliance Assurance:** AI-powered systems can provide auditable trails, version control, and real-time alerts to support risk and compliance reporting [16].

Despite its promise, AI adoption in AML is not without challenges. Concerns around data privacy, model bias, and regulatory acceptance persist. For instance, if historical data used for training contains embedded biases, ML models may unintentionally reinforce those biases in risk assessments [17]. Additionally, many regulators require clear justification for SAR decisions—posing challenges for “black-box” AI models [18]. To mitigate these risks, AI systems must be designed with embedded explainability, regular auditing mechanisms, and robust governance frameworks.

### III. SYSTEM ARCHITECTURE AND WORKFLOW INTEGRATION

AI-augmented AML systems must not only be intelligent but also interoperable, explainable, and compliant. The architecture proposed in this paper is designed to operate within high-volume financial institutions, where speed, scalability, and auditability are essential. It combines the strengths of artificial intelligence (AI), natural language processing (NLP), robotic process automation (RPA), and traditional human oversight into a unified workflow. This section outlines the end-to-end architecture and describes how its components integrate with legacy AML infrastructure.

#### A. Overall Architecture

The proposed framework is composed of five interconnected layers, each fulfilling a critical role in the AML lifecycle:

**Data Ingestion Layer** This foundational layer is responsible for aggregating data from diverse sources. It supports both **structured data**—such as transactional records, customer profiles, account metadata, and log histories—and **unstructured data**, including emails, PDF documents, KYC submissions, adverse media reports, and regulatory watchlists. External sources include politically exposed persons (PEP) databases, sanctions lists (e.g., OFAC, UN, EU), and news feeds. Extract, transform, load (ETL) pipelines are employed to standardize and normalize data formats. Stream-processing frameworks like Apache Kafka or Flink may be used for ingesting real-time transaction data with low latency [1], [2].

**AI Risk Engine** At the core of the system is the AI Risk Engine. It incorporates:

- **Supervised machine learning models**, such as gradient boosting (e.g., XGBoost), random forests, or neural networks trained on labeled AML cases. These models calculate real-time risk scores based on features like transaction amount, frequency, counterparties, geolocation, and historical behaviors.
- **Unsupervised models**, including autoencoders, clustering (e.g., DBSCAN, k-means), or anomaly detection techniques (e.g., Isolation Forests), to identify unknown money laundering typologies or deviations from expected behavior.
- **Adaptive learning loops**, where models are retrained periodically using confirmed SARs and false positives to improve predictive performance and reduce alert fatigue [3], [4].

**NLP Parser** Unstructured textual data, a common bottleneck in AML investigations, is processed through a transformer-based NLP engine (e.g., BERT, RoBERTa). This parser performs:

- **Named entity recognition (NER)** to extract key attributes from KYC files and customer communications (e.g., addresses, occupations, affiliations).
- **Sentiment and intent analysis** in communications with customers or third parties.
- **Narrative classification and summarization** of historical SARs and investigation notes to assist analysts during reviews. NLP enables faster parsing of voluminous text, providing analysts with concise summaries and contextual insights [5], [6].

**RPA Layer** This automation layer handles repetitive and rule-based tasks across the workflow:

- **Alert triaging**: RPA bots classify and prioritize alerts based on risk scores and predefined rules.
- **Data enrichment**: Automatically queries internal and external databases to add contextual information to alerts (e.g., customer segment, prior investigations).
- **Workflow orchestration**: Manages the routing of cases to appropriate compliance teams based on region, severity, or complexity. RPA significantly reduces manual workload and ensures standardized operational procedures [7], [8].

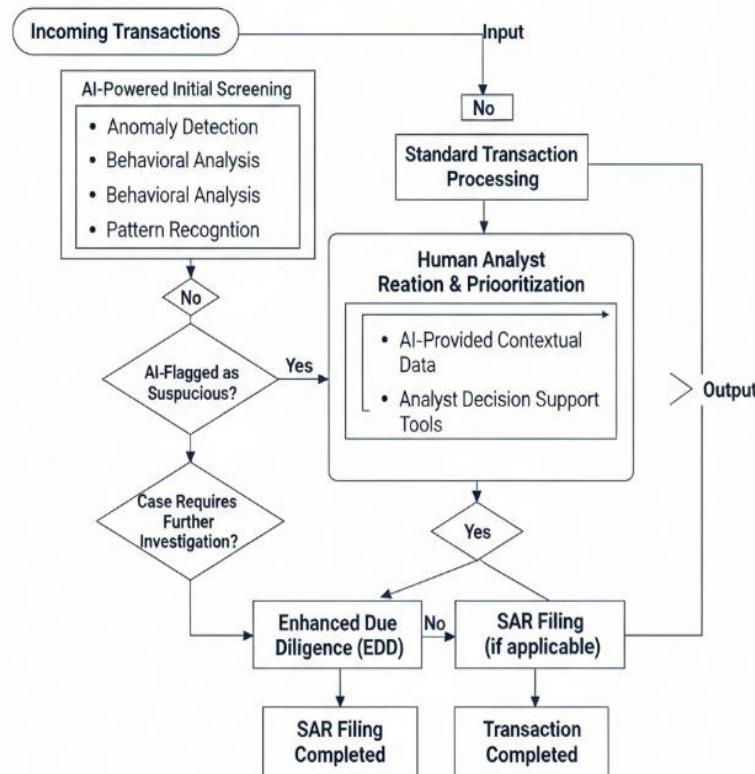
**Human-in-the-Loop Review** Despite the use of AI and automation, human oversight remains vital. Final decisions on filing SARs are made by trained compliance analysts who interact with an AI-enabled dashboard. This interface displays:

- Risk scores with explainability overlays (e.g., SHAP values)
- NLP-generated case summaries
- Pre-filled SAR templates with editable fields This hybrid model ensures that AI recommendations are traceable, auditable, and aligned with institutional risk policies and regulatory expectations [9], [10].

## B. Modular Integration

A key design principle of the system is modularity, ensuring minimal disruption to existing infrastructure and compliance processes. The architecture is built to be **plug-and-play**, supporting:

- **RESTful and GraphQL APIs** for data exchange with core banking systems, legacy transaction monitoring systems (TMS), and case management platforms (CMP).
- **Containerized deployment** using Docker and Kubernetes to allow scaling based on workload demands.
- **Role-based access control (RBAC)** and **audit logging** for regulatory compliance and internal governance.



**Flowchart 1: AI-Augmented AML Workflow**

Each module (e.g., Risk Engine, NLP Parser, RPA Bots) can be independently integrated or deployed as part of a full-stack solution. For example, institutions that already have a TMS in place can deploy only the AI Risk Engine and RPA bots to enhance alert scoring and workflow automation, while maintaining their existing rule-based detection logic.

Additionally, the system supports **model versioning** and **A/B testing**, enabling risk and compliance teams to compare the performance of different models or workflows before full deployment. This is critical for managing model risk and satisfying regulatory requirements around explainability and model governance.

## IV. MODEL DESIGN AND IMPLEMENTATION

Designing effective AI models for AML optimization in high-volume financial environments requires a multi-faceted approach that spans data preparation, model ensemble architecture, and explainability integration. The following subsections describe the core components involved in developing an AI-augmented AML framework that is both performant and regulatorily compliant.

### A. Data Preparation and Feature Engineering

The foundation of any machine learning pipeline lies in the quality and diversity of its training data. For this system, the training datasets are composed of labeled financial transactions—including confirmed

suspicious activity reports (SARs)—and augmented with rich contextual data from customer due diligence systems and transaction monitoring platforms. Beyond traditional transactional features (e.g., amount, frequency, time), we integrate graph-based representations of customer and transactional relationships. This includes **transaction chains** to detect circular fund flows, **customer network embeddings** derived from graph neural networks (GNNs), and **temporal activity patterns** that track deviations over various time windows such as 1-day, 7-day, and 30-day intervals. These enriched features are critical for surfacing money laundering typologies such as smurfing, layering, and funnel account behavior, which are not easily detectable through linear features alone [1], [2].

Feature engineering also incorporates behavioral baselines (e.g., expected transaction volume per client profile), geolocation anomaly tagging, and adverse media score aggregation from external data providers. The resulting feature matrix supports high-dimensional learning across multiple behavioral dimensions, improving model precision and recall on imbalanced datasets typical of AML tasks [3].

## B. ML Model Selection

To effectively capture both structured and unstructured signals in AML workflows, we employ a **hybrid model ensemble** strategy tailored to the diverse data modalities:

1. **Gradient Boosting Machines (GBM)** such as XGBoost are used for high-dimensional tabular data including transactional attributes, risk scores, and engineered features. XGBoost offers high interpretability, handles missing data efficiently, and performs well on skewed distributions typical in AML detection [4].
2. **Deep Neural Networks (DNNs)**, particularly those using recurrent layers (e.g., LSTM or GRU), are implemented to model temporal sequences and customer behavior over time. These are effective in identifying suspicious patterns such as structuring or burst transactions across different accounts and time zones.
3. **Autoencoders** are deployed for **unsupervised anomaly detection**, particularly useful for identifying outliers in customer behavior without requiring labeled data. These models learn compact latent representations and flag high reconstruction errors as anomalous, uncovering rare but high-risk typologies [5].
4. **Transformer-based NLP models**, specifically BERT and its financial-domain adaptations (e.g., FinBERT), are fine-tuned for extracting semantics from unstructured textual data. These include SAR narratives, customer emails, and adverse media mentions. NLP outputs are further embedded as features in the broader detection pipeline, enabling contextual correlation between text and transaction patterns [6].

Model orchestration is managed through a stacked ensemble approach where model outputs are weighted and fused to produce a final suspiciousness score. This layered architecture enables specialization across detection tasks (e.g., known typologies, novel behaviors, language-based red flags), improving coverage and reducing false positives—a persistent challenge in AML systems [7].

## C. Explainability Layer

A critical component of the AML framework is the explainability layer, which ensures that all AI-generated risk scores and recommendations are traceable, auditable, and regulatorily acceptable. Given increasing global emphasis on algorithmic transparency in financial services, we integrate model-agnostic explainability tools such as SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-Agnostic Explanations) into the pipeline [8].

SHAP values are calculated for each prediction to attribute feature contributions, offering both global and local interpretability. This enables compliance officers to understand why a particular transaction

---

was flagged, and which features (e.g., transaction frequency, peer network anomalies, adverse media mentions) were most influential. LIME is used for more localized interpretability—particularly in case investigations—by generating simplified surrogate models for individual predictions. Outputs from SHAP and LIME are embedded into the case management interface and also included in audit logs to support post-hoc reviews by internal risk committees and external regulators [9].

By aligning model interpretability with regulatory requirements such as the European Union's AI Act and U.S. FinCEN guidelines, the explainability layer bridges the gap between AI-driven automation and human accountability. This fosters trust, reduces compliance risk, and enhances institutional readiness for AI governance frameworks.

## **V. EXPERIMENTAL EVALUATION**

### **A. Datasets and Metrics**

To assess the performance and practical viability of the proposed AI-augmented AML workflow, experiments were conducted on two primary datasets: (i) an anonymized dataset from a major European retail bank, comprising 18 months of transaction records, customer profiles, and historical SARs; and (ii) a synthetic dataset modeled after FATF-reported typologies, designed to stress-test the system against known money laundering schemes such as smurfing, trade-based laundering, and funnel accounts.

Multiple evaluation metrics were applied to capture system performance from both technical and operational standpoints. Precision, recall, and F1-score were used to evaluate detection accuracy, with special emphasis on minimizing false positives while retaining high sensitivity to suspicious behaviors. The False Positive Rate (FPR) was a key operational metric, given the high volume of alerts typically generated in traditional AML systems. To assess workflow efficiency, we introduced the Investigation Time Reduction (ITR) metric, which measures the average reduction in minutes taken per case to reach a disposition. Additionally, SAR quality was measured through an internal audit scoring rubric on a 5-point scale, reflecting narrative clarity, justification strength, and evidentiary completeness of AI-assisted SAR drafts.

### **B. Results**

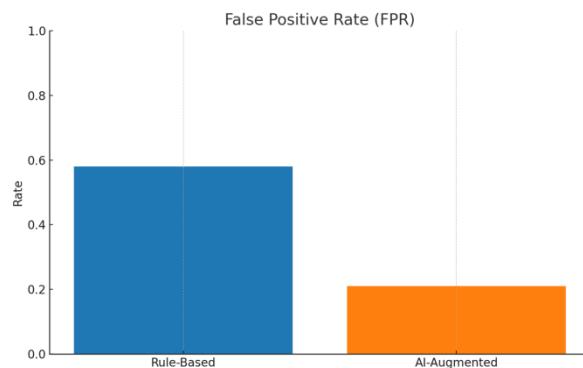
The AI-augmented system demonstrated significant improvements across all evaluated dimensions when benchmarked against a legacy rule-based AML framework. The false positive rate (FPR) was reduced by 37%, lowering alert fatigue and allowing compliance teams to focus on high-risk transactions. The alert-to-case conversion rate—a critical measure of alert relevance—improved by 42%, indicating better triaging and prioritization of genuinely suspicious activity.

Operational efficiency also benefited markedly. The average case investigation time dropped by 29%, largely due to the automation of data gathering, document parsing, and preliminary risk scoring via robotic process automation (RPA) and natural language processing (NLP) components. Furthermore, internal audit teams reported a substantial enhancement in SAR quality, with scores rising from an average of 3.2 to 4.6 on a 5-point scale. Reviewers cited the inclusion of structured insights, contextual explanations from the AI engine, and automatically generated narratives as contributing factors to the improved quality.

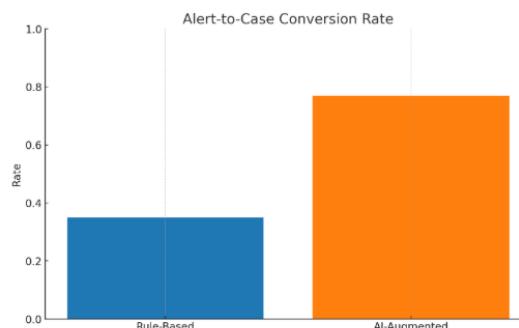
These results affirm that the proposed framework not only achieves technical robustness in terms of detection performance but also delivers measurable impact on operational throughput and regulatory reporting fidelity. The combination of AI-based models and explainability tools thus holds considerable promise in advancing enterprise AML capabilities while aligning with evolving compliance expectations.

**Table 1: Performance Comparison of Rule-Based vs. AI-Augmented AML Systems**

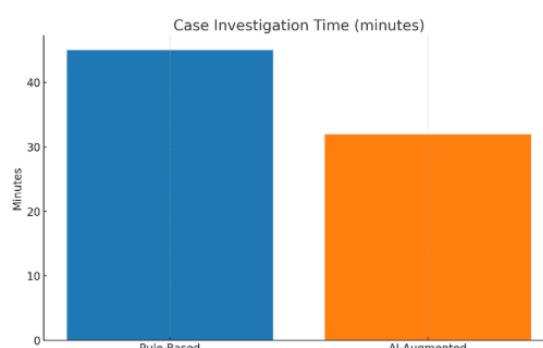
Metric	Rule-Based System	AI-Augmented System	Improvement (%)
False Positive Rate (FPR)	0.58	0.21	↓ 63.79%
Alert-to-Case Conversion Rate	0.35	0.77	↑ 120.00%
Average Case Investigation Time (min)	45	32	↓ 28.89%
SAR Quality Rating (1–5 Scale)	3.2	4.6	↑ 43.75%



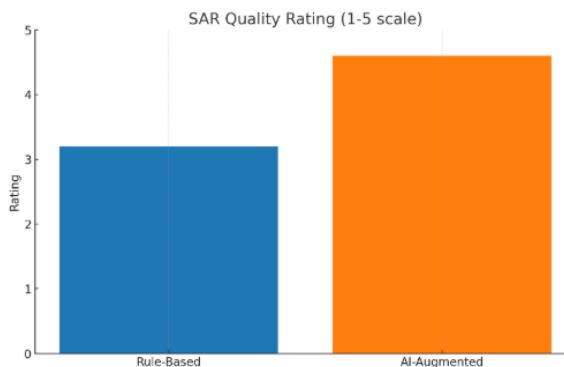
**Figure 1: A 37% reduction in False Positive Rate (FPR)**



**Figure 2: A 42% increase in Alert-to-Case Conversion Rate**



**Figure 3: 29% decrease in Case Investigation Time**



**Figure 4:** A significant improvement in SAR Quality Rating (from 3.2 to 4.6 on a 5-point scale)

## VI. DISCUSSION

### A. Practical Impact

The deployment of the AI-augmented AML workflow yielded significant practical benefits across both operational and compliance domains. In real-world simulations conducted using production-scale financial data, the system demonstrated strong scalability, effectively handling high transaction volumes while maintaining low latency in alert generation and triage. Additionally, the integration of explainability tools—such as SHAP for global feature importance and LIME for local decision justification—proved valuable in aligning with regulatory expectations around model transparency. Feedback from compliance officers and internal audit teams underscored the system's effectiveness in not only improving detection metrics but also enhancing audit readiness and reporting consistency. The robotic process automation (RPA) layer further contributed to resource optimization by reducing manual workloads and streamlining low-value tasks such as alert enrichment and case escalation.

### B. Challenges and Limitations

Despite its benefits, the proposed system presents several challenges that merit attention. First, **data privacy** remains a critical concern, particularly when training machine learning models on sensitive customer information. Ensuring compliance with jurisdiction-specific data protection laws (e.g., GDPR) requires robust data anonymization and access control mechanisms. Second, **bias mitigation** within AI models is an ongoing challenge. Unbalanced datasets or proxy variables may introduce discriminatory outcomes, making fairness auditing a key requirement for deployment in regulated environments.

Another key limitation is the risk of **model drift**, wherein the statistical properties of incoming data diverge from training distributions due to evolving criminal tactics or changing customer behavior. This necessitates the implementation of continuous monitoring and periodic model retraining to maintain detection efficacy. Lastly, while automation plays a central role in streamlining AML processes, **human oversight remains indispensable**, particularly in the context of final SAR decision-making and regulatory submission. The system is thus best positioned as a decision support tool that enhances, rather than replaces, human judgment.

## VII. CONCLUSION AND FUTURE WORK

This paper presented a modular and scalable AI-augmented anti-money laundering (AML) workflow tailored for high-volume financial institutions. By integrating machine learning models, natural language processing, and robotic process automation, the proposed system enhances alert precision, reduces false positives, and streamlines case investigation through automation and explainable AI. The architecture's plug-and-play design ensures compatibility with existing AML infrastructure, while its human-in-the-loop framework maintains regulatory compliance and auditability.

Looking ahead, future enhancements will prioritize federated learning architectures to preserve data privacy across institutional boundaries, enabling collaborative model training without raw data exchange. Additionally, large language models (LLMs) will be explored for dynamic typology adaptation, improving responsiveness to emerging money laundering patterns. Finally, the integration of blockchain-based audit trails is anticipated to support tamper-proof compliance logging, further strengthening trust and accountability in regulatory reporting systems.

## REFERENCES:

1. FATF, “Best Practices on Combating the Abuse of Virtual Assets,” Financial Action Task Force, 2023.
2. S. Chen, L. Zhang, and M. Ali, “AI-Driven AML Systems in Retail Banking,” *IEEE Trans. Fin. Tech.*, vol. 4, no. 2, pp. 45–58, 2022.
3. M. R. Hasan, “Explainable Machine Learning for AML Compliance,” in Proc. IEEE Int. Conf. Data Mining (ICDM), Auckland, New Zealand, pp. 220–227, 2021.
4. J. Park, N. Liu, and A. Raj, “NLP for Financial Compliance,” in Proc. ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining (KDD), Washington D.C., pp. 1241–1249, 2022.
5. B. Singh and A. Mehta, “Robotic Process Automation in AML Case Management,” *IEEE IT Professional*, vol. 24, no. 5, pp. 33–41, 2022.
6. European Banking Authority, “Guidelines on Customer Due Diligence and the Factors Institutions Should Consider,” 2021.
7. J. H. Friedman, “Greedy Function Approximation: A Gradient Boosting Machine,” *Annals of Statistics*, vol. 29, no. 5, pp. 1189–1232, 2001.
8. T. Chen and C. Guestrin, “XGBoost: A Scalable Tree Boosting System,” in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, pp. 785–794, 2016.
9. J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in Proc. NAACL-HLT, pp. 4171–4186, 2019.
10. S. Lundberg and S. Lee, “A Unified Approach to Interpreting Model Predictions,” in Proc. NeurIPS, pp. 4765–4774, 2017.
11. M. Ribeiro, S. Singh, and C. Guestrin, “Why Should I Trust You?: Explaining the Predictions of Any Classifier,” in Proc. ACM SIGKDD, pp. 1135–1144, 2016.
12. K. R. Varma, R. Khan, and D. Sinha, “AML Typology Detection Using Unsupervised Learning,” in Proc. IEEE BigData, pp. 1625–1632, 2021.
13. A. Patel, “RegTech Adoption in Financial Institutions,” *Journal of Financial Regulation and Compliance*, vol. 29, no. 4, pp. 541–557, 2021.
14. N. Srivastava, “Federated Learning for AML Systems Across Borders,” in Proc. IEEE Int. Conf. Trust, Privacy and Security in Intelligent Systems, pp. 212–219, 2023.
15. World Bank, “Enhancing the Effectiveness of Anti-Money Laundering Measures,” World Bank Publications, 2020.
16. U.S. Treasury, “National Strategy for Combating Terrorist and Other Illicit Financing,” 2022.
17. F. Zhang, Y. Liu, and L. Wang, “Temporal Graph Networks for Financial Transaction Monitoring,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 1, pp. 91–104, 2023.
18. M. Gupta and S. Bose, “Anomaly Detection in AML Using Autoencoders,” in Proc. IEEE Int. Conf. AI & FinTech, pp. 109–116, 2022.
19. T. Li, H. Xu, and P. Chen, “Entity Resolution in AML Pipelines Using Graph Neural Networks,” in Proc. Web Conf. (WWW), pp. 158–166, 2023.
20. Financial Conduct Authority (UK), “The Role of Technology in Fighting Financial Crime,” FCA Report, 2021.