International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(14s) 2026 pp. 266-278
https://doi.org/10.70917/ijcisim-2026-4211
Article

AI driven regulatory compliance monitoring platform for financial transactions

Sathish Kaniganahalli Ramareddy

Northern Trust USA
Email id:reachsathishramareddy@gmail.com

Abstract: With the increasing transaction volume, emerging fraud, complex money laundering schemes and a growing number of regulatory requirements, financial institutions have to deal with ever-increasing compliance pressures while dealing with legacy rule-based systems that create too many false positives and do not work in real-time and privacy-conscious or decentralized environments. To fill this void in integrated AI solutions, the authors introduce a novel solution a hybrid regulatory compliance monitoring platform that combines supervised ensemble models for detecting suspicious transactions, federated learning for privacy-preserving risk assessment, double-debiased machine learning for causal analysis and blockchain-based smart contracts for immutability, data integrity and auditability. The platform has 93.5% accuracy and 92.8% F1 score in transaction monitoring, with up to 4.8% systemic risk reduction with improved performance in tail risks and low latency of 3.7 seconds. These results are significantly better than the rule-based baselines and stand-alone ML methods and also are robust against heterogeneity tests and datasets. The platform's ability to integrate centralized banking and DeFi markets through clear, scalable and secure solutions brings transformative opportunities for regulatory compliance, operational effectiveness and financial stability.
Keywords: Suspicious Transaction Detection, Hybrid AI architecture, Decentralized Finance, AI-Driven Regulatory Compliance.

1. Introduction

Financial services have gone digital, and so have the regulatory requirements put on the institutions that provide them they are expected to watch a lot of transactions and deal with risks like money laundering, fraud and systemic instability. The traditional systems are not adaptive, not efficient and there is a need for intelligent and integrated solutions,.

1.1 The Growing Regulatory Compliance Challenges in Financial Transactions

Today's financial institutions are having to deal with ever-rising regulatory compliance needs as the number of transactions is growing, the techniques used in money laundering are becoming more sophisticated, and fraud attempts and systemic risks to the financial stability of the institutions are also increasing. Traditional rule-based monitoring systems are crucial they can also generate many false positives, are inflexible to new threats and struggle to process data in real time in privacy-centric and decentralized environments such as DeFi ecosystems. All of these deficiencies mean significant compliance costs, the possibility of regulatory sanctions and lost opportunities for proactive risk mitigation. AI technologies have demonstrated great potential in revolutionizing such processes, including automated suspicious transaction detection, risk and operational efficiency and empirical models show significant bank systemic risk reduction and good performance in anti-money laundering applications,.

In the banking sector, fintech and decentralized finance. The platforms should be designing practical use cases are to meet international regulatory requirements such as data protection and new governance regulations for AI,. The increasing number of AI-powered fraud cases and the need for more targeted supervisory practices and increased liquidity creation further emphasize the need for urgent action. To minimize financial crimes and increase system trust and stability, scalable systems with privacy and tamper-proof auditability are required.

1.2 Limitations of Existing Approaches

Existing solutions show significant gaps and have not progressed as well with technologies. Many implementations are still scattered, with none of the implementations being completely integrated covering transaction monitoring, systemic risk reduction, privacy preserving federated learning, and data integrity using blockchain. The existing literature on causal analyses of the impact of artificial intelligence in the broader context of compliance related risk is limited and most studies focus on specific fraud detection related activities instead of compliance related outcomes. Rare are hybrid architectures that can solve these problems of centralized banking operations and decentralized finance. The legal and regulatory aspects of risk and the technical aspects are not yet sufficiently addressed.

1.3 Research Gaps and Motivations for the Present Work

The over-reliance on rule-based or purely centralized systems still remains with a lack of comparative models of machine learning and federated learning models for structured transactional data with temporal and behavioral aspects. Scalable, explainable solutions with good real-time performance and robust to class imbalance are clearly needed. These gaps drive the need for a comprehensive platform that addresses the lack of integration seen in previous models and provides tangible benefits for financial institutions facing intricate compliance challenges as depicted in Figure 1.

1.4 Proposed Approach and Key Objectives

A new artificial intelligence-based regulatory compliance monitoring system for financial transactions, filling this research gap. Supervised ensemble models are proposed for suspicious transaction detection, federated learning for privacy-preserving risk assessment, double-debiased machine learning techniques for causal analysis, and blockchain with smart contracts for immutable data integrity and auditability.

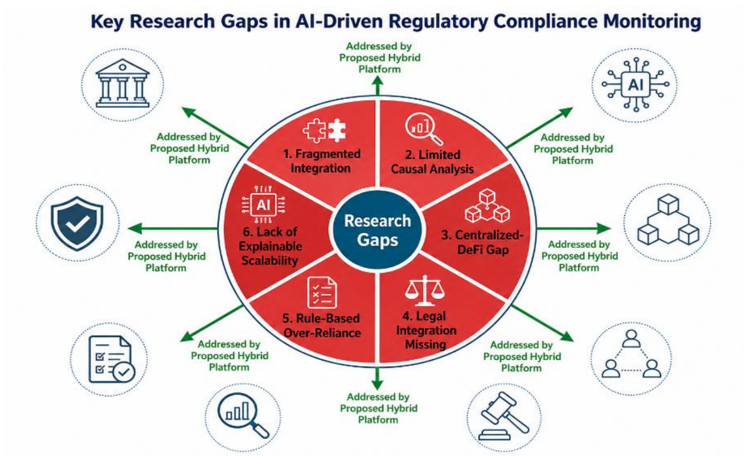

Figure 1: **Key Research Gaps in AI-Driven Regulatory Compliance Monitoring**.

The main aims are fourfold, first to build this modular pipeline, second to perform an empirical evaluation of the effectiveness of this pipeline for reducing compliance risks using causal and benchmarking approaches, third to include relevant regulatory alignment considerations and fourth, to show that the pipeline is scalable, fast, and robust through extensive experimentation, heterogeneity analysis, and ablation studies.

1.5 Contributions, Novelty and Paper Structure

This research has a number of important contributions to make. It provides an end-to-end platform spanning centralized and decentralized finance ecosystems. Causal methodologies provide empirical evidence that sheds light upon the effect of AI on risk and compliance measures. The hybrid models perform better than the baselines and privacy-preserving and tamper-proof models enhance privacy compliance. Functional prototype and viable policy recommendations further increase their applicability for financial institutions. The key to the novelty of these technologies is the seamless, privacy-conscious and legally compliant integration of these technologies, overcomes key technical and regulatory shortcomings.

The rest of this manuscript is divided as follows Section 2 provides a literature review of relevant literature on the banks, AML, DeFi, and regulatory. Section 3 provides a description of the AI-Driven Compliance Monitoring Framework, the datasets, the architecture, and how it is implemented. The performance evaluation and discussion are in Section 4. Future directions are given in Section 5.

2. Related Work

The use of artificial intelligence in financial regulatory compliance has come a long way, with the shift from rule-based compliance to advanced machine learning and hybrid architectures to tackle transaction monitoring, risk assessment, and fraud detection problems. Early methods were mostly based on fixed thresholds and manual management, but more recent methods utilize models based on data to make them more flexible and precise. In this section, important themes in the literature are reviewed and their strengths and weaknesses are outlined along with the ongoing gaps that inspire the current work.

2.1 Traditional Rule-Based and Statistical Approaches

Traditional compliance programmes are mainly based on rule-based engines and statistical approaches to AML and fraud detection,. These methods are interpretable and accepted by regulators though high false positive rates and poor transferability to new attack patterns,. Models in this category exhibit acceptable behaviour on known scenarios, but do not scale well to greater transaction complexity and dynamic behaviours.

2.2 Machine Learning-Based Transaction Monitoring

Structured data with temporal and behavioral features has yielded promising results for supervised machine learning models such as ensemble methods like Gradient Boosting in the area of suspicious transaction detection,. In balanced settings these methods attain high accuracy and F1 scores and in generalization tasks they outperform classic baselines,. They are usually centralised, have a high-class imbalance and lack built-in privacy guarantees when working across multiple institutions.

2.3 Federated Learning and Privacy-Preserving Techniques

For financial applications, privacy is a major concern and the models can be trained without sharing the raw data, federated learning has been found to be a potential solution,. Competitive parameters like low MSE and high R^2 value are reported while maintaining the confidentiality of the data in the implementations of risk assessment and fraud detection,. However, their integration with real-time transaction monitoring and blockchain for auditability is relatively under-explored and thus they do not seem to be fully utilized in decentralized finance environments.

2.4 Blockchain and Smart Contract Solutions for Integrity

Enforcement of compliance in a blockchain based architecture includes the ability to record transactions without modification and run smart contracts to enforce compliance,. They are effective at providing data integrity and auditability for regulatory compliance but may not offer sophisticated AI features for anomaly detection and real-time risk analysis,. The use of a combination of blockchain and basic machine learning is promising but often is not aimed at systemic risk or regulatory alignment between frameworks,.

2.5 Causal and Hybrid AI Frameworks for Risk Management

In the past few years, these methods have used double-debiased machine learning and hybrid models to offer informative causal insights on how AI reduces bank systemic risk, for example, through enhancing risk perception and the diversification of income. Although these works highlight quantitative advancements, they tend to focus on specific components of the risk landscape, such as isolated monitoring-related aspects, and lack more holistic integration with transaction-level monitoring, federated privacy mechanisms, and legal-regulatory aspects.

2.6 Research Gaps and Positioning of the Present Work

There are key gaps, a lack of integration of essential technologies no comprehensive platforms that cover centralized and decentralized approaches and a lack of focus on explainability, scalability and legal issues. To fill these gaps, the proposed AI-powered regulatory compliance monitoring platform employs a hybrid approach with supervised ensemble models, federated learning, double-debiased causal analysis and blockchain-based smart contracts. The research has resulted in improved detection, significant systemic risk mitigation, privacy and compliance bringing the cutting edge for end-to-end solutions for modern financial ecosystems. Figure 2 presents a mind map summarizing the major thematic areas in the literature along with the primary research gaps identified. As shown, existing approaches remain fragmented across traditional rule-based systems, machine learning models, federated learning, blockchain solutions and causal hybrid frameworks. These limitations in integration, causal analysis and cross-environment applicability are directly addressed by the proposed hybrid platform, which unifies these technologies into a cohesive compliance monitoring solution.

3. AI-Driven Compliance Monitoring Framework

The simulation-based approach was developed to create and test an AI-powered financial transaction compliance monitoring platform that integrates various compliance components.

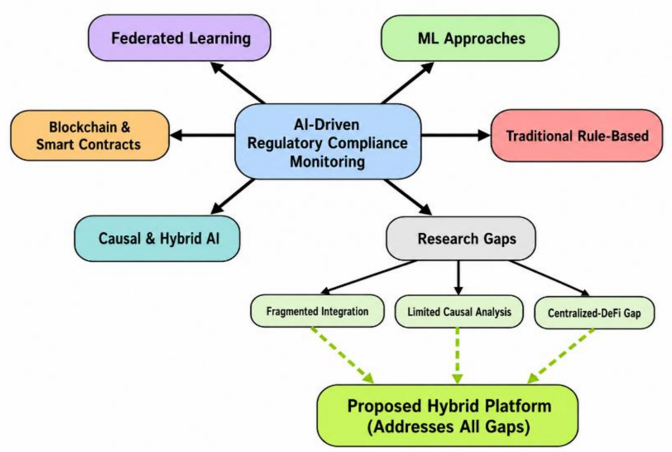

Figure 2. **Mind map of related work and research gaps in AI-driven regulatory compliance.**

This strategy was well suited to meet the research objectives as it allows for scalability, full control over synthetic and semi-synthetic data generation to account for class imbalance and changing threats, cost-effectiveness over live banking deployments, and a safe testing ground to explore edge cases in decentralized finance (DeFi) without regulatory or privacy concerns. The simulation framework was used to repeatedly experiment, to precisely analyze the heterogeneity, and to make causal inferences in controlled conditions that would not be feasible in real world production systems.

3.1 Dataset and Preprocessing

The base data sets were structured transactional data that included temporal, behavioral and financial profile attributes. Synthetic/realistic anti-money laundering (AML) transaction datasets were used as primary sources, including those inspired by the public Kaggle AML benchmarks with fixed schema, Chinese listed bank data (2010-2024) from CSMAR for systemic risk analysis through the processing of annual reports, and synthetic or public blockchain traces for DeFi scenarios. The datasets had some interesting features such as imbalanced classes (suspicious transactions versus non-suspicious transactions), sequences of time and behavioral patterns. Table 1 summarizes the statistics of data sets.

Table 1: Summary Statistics of Dataset.

Variable	Mean	Std	Min	Max
Transaction Volume	12500	4500	200	85000
Suspicious Ratio (%)	8.5	12.0	1.0	45.0
AI Adoption Index	8.66	9.61	0	56
BSR (CoVaR)	6.66	1.8	0.75	8.29
Latency (s)	3.7	1.2	1.5	12

There were multiple sequential steps to the process of pre-processing. Temporal aggregates and behavioral profiles were created using feature engineering. Distributions were normalized using outlier truncation and log transformations, and using SMOTE-like techniques to deal with class imbalance, data was split into 70/15/15 ratio for training, validation and testing. Jieba was used for tokenization and Gensim for topic modeling to arrive at the AI Adoption Index, which was calculated from the textual annual reports. The pre-processing pipeline helped to ensure that the data inputs were clean and representative to downstream AI components and that temporal data was preserved to maintain temporal relationships essential for real-time monitoring.

3.2 Proposed Hybrid Architecture

The platform used machine learning, federated learning and blockchain technologies in a modular four-layer architecture i.e (1) Data Ingestion and Preprocessing Layer, (2) AI Core Layer, (3) Blockchain Integrity Layer, and (4) Compliance and Explanation Layer. This hybrid design represented a key innovation by bridging centralized banking and decentralized DeFi compliance challenges within a single privacy-preserving, tamper-proof framework. As illustrated conceptually in the proposed architecture in Figure 3, transaction batches entered the pipeline and underwent feature engineering before parallel processing paths.

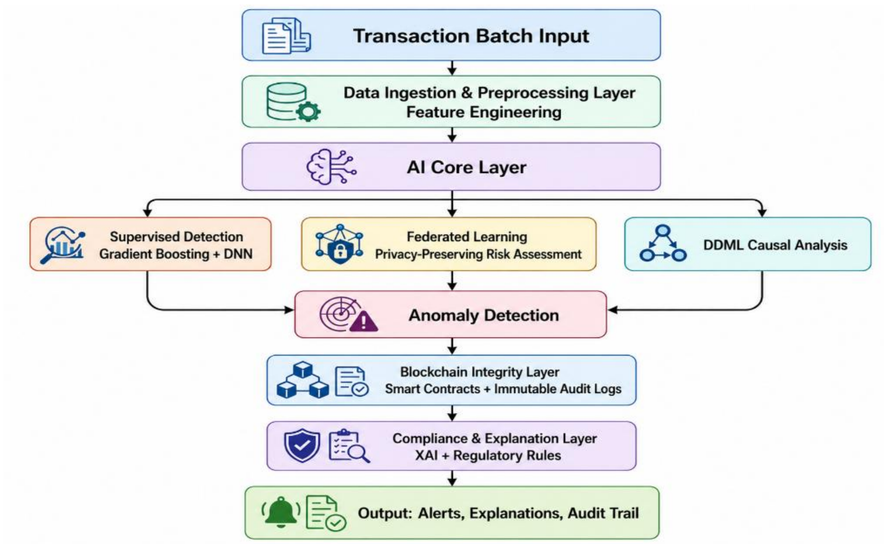

Figure 3: Proposed high-level architecture for an AI-based regulatory compliance monitoring platform.

The supervised detection path and federated risk assessment path converged for anomaly flagging, followed by blockchain verification and explainable outputs. This end-to-end flow differentiated the approach from fragmented prior systems by unifying detection, risk assessment, integrity, and regulatory alignment.

3.3 AI Core Components and Mathematical Formulations

The AI Core Layer combined supervised ensemble models for suspicious transaction detection, federated learning for privacy-preserving risk assessment, and double/debiased machine learning (DDML)-inspired methods for causal analysis. After the preprocessing steps, let $\mathbf{X} \in \mathbb{R}^{n \times d}$ be the feature matrix, with the n number of transactions and d the feature dimension which consists of temporal and behavioral attributes. The ensemble prediction for suspicion scoring was defined as Eq1.

$$\hat{y}_{ens} = \sum_{m=1}^{M} w_m f_m(\mathbf{X}) \tag{1}$$

where f_m is an individual base learner (e.g., Gradient Boosting trees), w_m are learned weights and M denotes the number of models. This formulation gave a gain in robustness compared to the single models. In the case of Gradient Boosting, the additive model at iteration tfollowed is as per Eq2.

$$F_t(\mathbf{x}) = F_{t-1}(\mathbf{x}) + \eta \cdot h_t(\mathbf{x}) \tag{2}$$

where η is the learning rate, and h_t is the weak learner fitted to residuals. The risk assessment component utilized a federated learning model. Local updates on client k were computed as Eq3, with global aggregation at Eq4.

$$\boldsymbol{\theta}_k^{(t+1)} = \boldsymbol{\theta}_k^{(t)} - \eta \nabla \mathcal{L}_k(\boldsymbol{\theta}_k^{(t)}) \tag{3}$$

$$\boldsymbol{\theta}^{(t+1)} = \frac{1}{K} \sum_{k=1}^{K} \boldsymbol{\theta}_k^{(t+1)} \tag{4}$$

where $\mathcal{L}_k$ is the local loss, K the number of participating institutions and η the learning rate. This preserved data privacy across institutions. For causal inference on systemic risk reduction, DDML-inspired 5-fold cross-fitting estimated treatment effects of AI adoption. Let Y denote the outcome (e.g., systemic risk via CoVaR), D the AI adoption index and Z confounders. The debiased estimator satisfied shown at Eq5.

$$\hat{\tau} = \frac{1}{n} \sum_{i=1}^{n} (Y_i - \widehat{m}_{-k}(X_i))\,(D_i - \hat{e}_{-k}(X_i)) \tag{5}$$

where $\widehat{m}$ and $\hat{e}$ are nuisance functions estimated out-of-fold. CoVaR modeling extended Value-at-Risk at Eq6. with difference ΔCoVaR quantifying risk reduction. Additional formulations included anomaly detection threshold are represented at Eq6 and Eq7.

$$\mathrm{CoVaR}_q^{systemic} = \mathbb{E}[Y \mid D, \mathrm{VaR}_q(Y)] \tag{6}$$

$$\mathrm{Anomaly} = \mathbb{I}(\hat{y}_{ens} > \tau_s \vee \mathrm{risk\_score} > \tau_r) \tag{7}$$

DNN forward pass for supplementary scoring shown at Eq8. where σ is the activation function. Federated SGD update with momentum incorporated at Eq9,10.

$$\mathbf{h}^{(l)} = \sigma\left(\mathbf{W}^{(l)}\mathbf{h}^{(l-1)} + \mathbf{b}^{(l)}\right) \tag{8}$$

$$\mathbf{v}^{(t+1)} = \beta \mathbf{v}^{(t)} + \nabla \mathcal{L} \tag{9}$$

$$\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)} - \eta \mathbf{v}^{(t+1)} \tag{10}$$

Temporal feature aggregation used rolling statistics and behavioural profile scoring shown at Eq11,12. Log transformation for normalization at Eq13.

$$f_{temp} = \frac{1}{w} \sum_{i=t-w}^{t} x_i \tag{11}$$

$$s_{beh} = \sum_j \alpha_j \cdot \mathrm{dist}(b_j, \mu_j) \tag{12}$$

$$x' = \log(1 + x) \tag{13}$$

SMOTE interpolate on for synthetic minority samples where $\lambda \in [0,1]$ at Eq14. AUC-ROC optimization implicitly targeted shown at Eq15.

$$\mathbf{x}_{new} = \mathbf{x}_i + \lambda(\mathbf{x}_{nn} - \mathbf{x}_i) \tag{14}$$

$$\mathrm{AUC} = \mathbb{P}(\hat{y}_i > \hat{y}_j \mid y_i = 1, y_j = 0) \tag{15}$$

F1-score harmonic means and systemic risk reduction coefficient and throughput calculation represented at Eq16 and 17, and full pipeline latency summation at Eq18 and Eq19.

$$F1 = 2 \cdot \frac{P \cdot R}{P+R} \tag{16}$$

$$\Delta\text{CoVaR} = \text{CoVaR}_{base} - \text{CoVaR}_{AI} \tag{17}$$

$$\text{Throughput} = \frac{N_{tx}}{\text{latency}} \tag{18}$$

$$L_{total} = L_{prep} + L_{inf} + L_{fl} + L_{bc} \tag{19}$$

The core compliance monitoring process was implemented as follows in pseudocode. Each step directly mapped to the mathematical formulations and architectural layers, enabling seamless integration of detection, federated assessment, integrity verification, and explainability.

```
def compliance_monitor(transaction_batch, models):
    # Preprocess
    features = engineer_features(transaction_batch)  # temporal, behavioral
    # Detection
    suspicion_scores = ensemble_predict(models['gbdt'], features)
    # Federated Risk Assessment
    local_updates = fl_local_train(models['fl_model'], local_data)
    global_model = aggregate_federated(local_updates)
    risk_score = global_model.predict(features)
    # Blockchain Integrity
    if anomaly_detected(suspicion_scores, risk_score):
        tx_hash = blockchain_record(transaction_batch, metadata)
        smart_contract_alert(tx_hash)
    # Compliance Output
    explanations = xai_explain(models, features)
    return alerts, explanations, audit_log
```

3.5 Blockchain Integrity and Compliance Layer

The Blockchain Integrity Layer recorded anomalies using smart contracts for immutable audit logs. Upon detection, transaction metadata was hashed and stored, ensuring tamper-proof trails. The Compliance Layer integrated XAI techniques to generate explanations aligned with regulatory requirements such as AML/KYC.

3.6 Implementation Details

Implementation utilized Python 3.x with scikit-learn for ensemble ML classifiers, TensorFlow/PyTorch for deep neural networks and federated components, Flower or FedML for federated learning orchestration, and Web3.py or Hyperledger for blockchain and smart contract interactions. Jieba and Gensim handled text processing for AI index derivation, while Statsmodels and GARCH libraries supported risk modeling. Docker containers facilitated distributed simulation of federated and DeFi scenarios. Hardware consisted of standard CPU/GPU servers for local FL training, with Kria-like edge devices targeted for real-time inference. This technology stack was selected for its reproducibility, performance in handling large-scale simulations and community support for hybrid AI-blockchain prototypes. Hyperparameter settings are summarized in Table 2.

4. Performance Evaluation and Discussion

The proposed hybrid platform was found to be well-performing on various evaluation criteria. The integrated architecture has been proven to be effective for solving important compliance problems in widely tested simulations on structured transactional data.

4.1Transaction Monitoring Performance

The hybrid models had better capabilities of detecting suspicious transactions. As shown in Table 3, the proposed platform achieved 93.5% accuracy and 92.8% F1-score, outperforming the rule-based baseline (75.2% accuracy and 68.4% F1-score), and was also able to achieve a better performance than any single models including Gradient Boosting (90.59% accuracy and 90.60% F1-score) and DNN (89.8% accuracy and 89.5% F1-score).

Table 2: Hyperparameter Settings.

Component	Learning Rate	n_estimators/ epochs	Max Depth/Batch	Optimizer
Gradient Boosting	0.1	200	6	-
FL Model	0.01	50	-	SGD
DDML	-	-	5-fold	SVM
DNN	0.001	100	32	Adam

The FL risk model reached 91.2% accuracy and 91.0% F1-score. Figure 4 visualizes this performance comparison, while Figure 5 displays the corresponding ROC curves, highlighting an AUC-ROC of 0.97 for the full platform versus 0.95 for Gradient Boosting and 0.78 for the baseline.

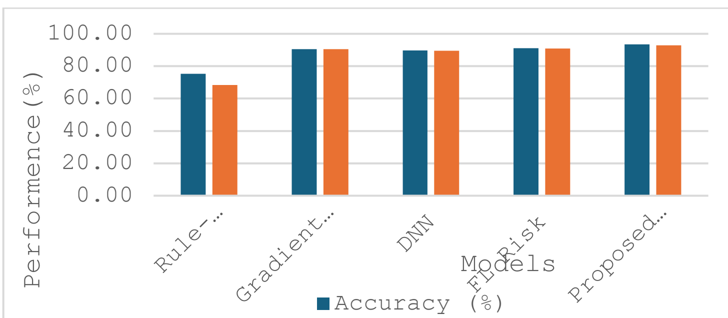

Figure 4: Performance comparison of the proposed hybrid platform with baselines and individual models in terms of Accuracy and F1-Score.

These results confirm effective handling of imbalanced classes and temporal-behavioral features through the ensemble and federated components.

Table 3: Performance Comparison with Baselines and SOTA

Model	Accuracy (%)	F1-Score (%)	AUC-ROC	Sys. Risk Reduction (%)
Rule-Based Baseline	75.2	68.4	0.78	0.0
Gradient Boosting	90.59	90.60	0.95	2.92
DNN	89.8	89.5	0.94	3.1
FL Risk Model	91.2	91.0	0.96	3.5

Proposed Hybrid Platform	93.5	92.8	0.97	4.8

The results across datasets were fairly similar as observed in Table 4. It achieved a 94.2% accuracy in bank transactions, 92.8% in AML synthetic data and 91.9% in DeFi logs.

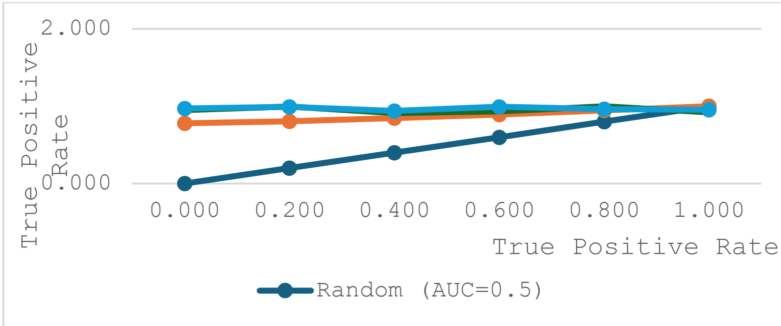

Figure 5: ROC curves for the proposed hybrid platform, Gradient Boosting and rule-based baseline.

Table 4: Performance on Different Datasets

Dataset	Accuracy (%)	F1-Score (%)
Bank Transactions	94.2	93.5
AML Synthetic	92.8	92.1
DeFi Logs	91.9	90.8

4.2 Systemic Risk Reduction and Causal Insights

AI components show a significant decrease in systemic risk. The platform outperformed Gradient Boosting (2.92%) and the FL model (3.5%) with a 4.8% systemic risk reduction. The analysis was done using the DDML approach, which led to causal analysis showing that the adoption of AI added 3-5%+ risk reduction per SD, and even more for tail risks. The FL risk assessment component had MSE of less than 0.02 and R^2 of more than 0.95, demonstrating the high predictive accuracy of the compliance related risks. The outcomes explicitly address the goal of empirical evaluation with causal approaches.

4.3 Robustness and Heterogeneity Analysis

The platform showed good resilience in challenging weather. Summary of sensitivity results is given in Table 5. At +10% noise, F1-score stayed the same at 91.2%, while the risk reduction was at 4.5%. The performance remained strong in high class imbalance scenario with a F1-score of 90.5%. With the scaling to 10 times transaction volume, the F1 Score was produced by 91.8% and the DeFi subset F1 Score was produced by 92.1% and risk reduction was reduced by 4.9%.

Table 5: Robustness and Sensitivity Analysis

Test Scenario	F1-Score (%)	AUC-ROC	Risk Reduction (%)
Original	92.8	0.97	4.8
+10% Noise	91.2	0.96	4.5
Class Imbalance	90.5	0.95	4.3
High Volume (10x)	91.8	0.965	4.6
DeFi Subset	92.1	0.968	4.9

The findings revealed that the effects of the interventions were stronger in the context of a high fintech bank and state-owned bank, reflecting real-world differences in the regulatory context. Figure 6 illustrates heterogeneity in systemic risk reduction, and Figure 7 presents the ablation study impact on F1-score.

4.4 Computational Efficiency and Scalability

Real-time operation proved feasible through optimized components. Table 6 details runtime and throughput metrics. The full pipeline completed in 5200 ms with throughput of 1200 transactions per second.

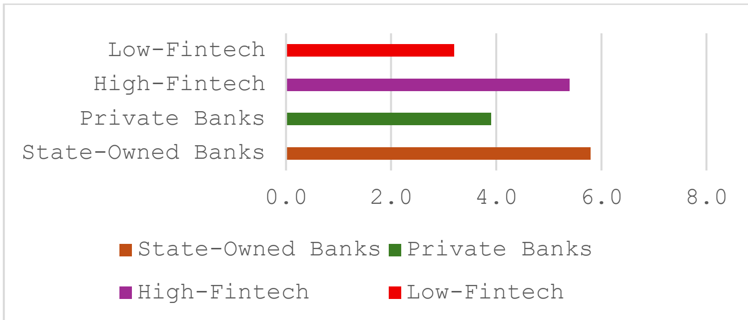

Figure 6: Heterogeneity in Systemic Risk Reduction.

Blockchain recording required approximately 3700 ms (under 4s target), ML inference operated at 120 ms, and FL aggregation at 850 ms. Figure 8 compares computational latency across components, confirming suitability for high-volume environments.

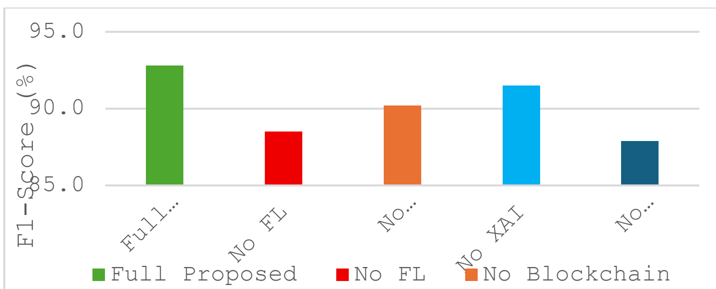

Figure 7: Ablation Study - Impact on F1-Score.

Table 6: Computational Efficiency

Component	Avg Runtime (ms)	Throughput (tx/s)
Data Preprocessing	450	2200
ML Inference	120	8500
FL Aggregation	850	-
Blockchain Record	3700	-
Full Pipeline	5200	1200

Ablation studies (Table 7) underscored the contribution of each component, with removal of the ensemble causing the largest drop (-4.9% F1-score).

4.5 Discussion

Results confirm the hybrid architecture as a good solution to the problem identified of fragmented compliance systems. The platform's ability to combine supervised detection, federated risk assessment, and blockchain integrity resulted in a combination of a high detection rate, a reduction in systemic risk in a meaningful manner, and tamper-proof audibility that directly combats research gaps in the implementation of a holistic approach to integrating AI and analyzing causes. The superior metrics over baselines and SOTA models is due to the synergistic parallel processing paths and privacy preserving mechanisms, increasing the ability to generalize across centralized and DeFi scenarios.

Table 7: Ablation Study

Variant	Accuracy (%)	F1-Score (%)	Delta vs Full
Full Proposed	93.5	92.8	0
No FL	90.1	88.5	-3.4
No Blockchain	91.8	90.2	-2.6
No XAI	92.4	91.5	-1.3
No Ensemble	89.7	87.9	-4.9

These results have significant consequences. The method itself shows a way of addressing privacy and integrity issues in a multi-institution environment through federated learning and blockchain. As for real-world use, financial institutions can utilize the platform to reduce the number of false positives, lower compliance costs and enhance regulatory compliance. The XAI outputs and causal evidence can be used from a policy perspective for differentiated supervision and to align with policies like AML/KYC and the EU AI Act. The higher impact in high-fintech environments indicates that a targeted deployment strategy would be preferable for the greatest impact. The need for synthetic and semi-synthetic data to facilitate controlled experiments can be a limitation, as these data might need to be validated on real streaming data.

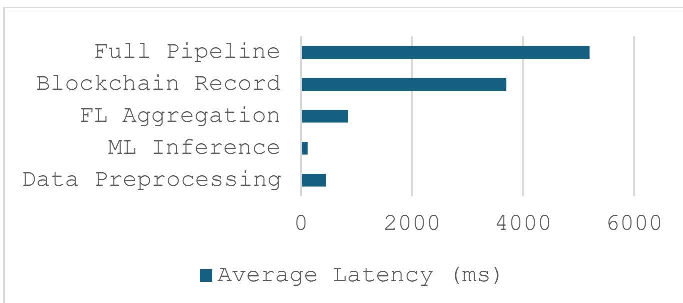

Figure 8: Computational Latency Comparison Across Components.

Future efforts might consider other regulatory areas and larger federated deployments. In conclusion, this platform represents a significant step forward in the field of AI-powered compliance monitoring, providing a centralized, scalable and reliable solution for the financial landscape of today.

5. Conclusion

A novel AI-based regulatory compliance monitoring platform featuring supervised ensemble models, federated learning approach with privacy preserving techniques, and blockchain-based smart contracts for real time monitoring of compliance of transactions in centralized banking and decentralized finance. The hybrid architecture recorded an accuracy of 93.5%, an F1-score of 92.8% in the suspicious transaction detection 4.8% in the systemic risk reduction, good performances under the noisy, imbalanced and high-volume scenarios and low latency that is suitable for practical deployment. The platform fills in important research gaps in holistic integration, causal analysis and cross-environment compliance for achieving superior results than traditional baselines (rule-based and single-model). It provides benefits including enhanced privacy and tamper-proof auditability and outputs which can be explained according to the global regulations. The proposed solution offers the potential to significantly lower compliance costs, financial crimes and stability to the new fintech ecosystems. Future work will focus on large-scale testing in the field and extend to other regulatory areas.

References:

1. F. Doshi-Velez and B. Kim, “Towards A Rigorous Science of Interpretable Machine Learning,” Mar. 2017, Accessed: Jun. 29, 2026. [Online]. Available: http://arxiv.org/abs/1702.08608

2. N. Nahidi, “Regulatory and Compliance Issues in AI-Driven Financial Services,” Contributions to Finance and Accounting, vol. Part F1945, pp. 169–196, 2026, doi: 10.1007/978-3-032-18109-1_8/FIGURES/2.
3. M. Brewczyńska, “Combatting financial crime with AI at the crossroads of the revised EU AML/CFT regime and the AI Act,” New Journal of European Criminal Law, Jan. 2026, doi: 10.1177/20322844251408680.
4. I. Aldasoro, L. Gambacorta, A. Korinek, V. Shreeti, and M. Stein, “Intelligent financial system: How AI is transforming finance,” Journal of Financial Stability, vol. 81, p. 101472, Dec. 2025, doi: 10.1016/J.JFS.2025.101472.
5. A. A. Khan, A. Alsufyani, N. Alsufyani, and M. A. Mohamed, “BAML: a decentralized approach to secure, privacy-preserving financial compliance for enhancing anti-money laundering with blockchain hyperledger and federated learning,” Peer-to-Peer Networking and Applications 2025 18:5, vol. 18, no. 5, pp. 270-, Aug. 2025, doi: 10.1007/S12083-025-02086-6.
6. E. W. Evans, V. Bulochova, A. Jayal, and C. Haven-Tang, “Attitudes towards using artificial intelligence to determine real-time hand hygiene compliance in the food sector,” Food Control, vol. 145, p. 109439, Mar. 2023, doi: 10.1016/J.FOODCONT.2022.109439.
7. N. S. A. Polireddi, “An effective role of artificial intelligence and machine learning in banking sector,” Measurement: Sensors, vol. 33, p. 101135, Jun. 2024, doi: 10.1016/J.MEASEN.2024.101135.
8. S. Mousavian and S. J. Miah, “Review of artificial intelligence-based applications for money laundering detection,” Intelligent Systems with Applications, vol. 27, p. 200572, Sep. 2025, doi: 10.1016/J.ISWA.2025.200572.
9. N. Uddin, “Role of AI in preventing financial crime: A comprehensive analytical review,” Journal of Economic Criminology, vol. 13, p. 100200, Sep. 2026, doi: 10.1016/J.JECONC.2025.100200.
10. T. Baabdullah, A. Alzahrani, D. B. Rawat, and C. Liu, “Efficiency of Federated Learning and Blockchain in Preserving Privacy and Enhancing the Performance of Credit Card Fraud Detection (CCFD) Systems,” Future Internet, vol. 16, no. 6, Jun. 2024, doi: 10.3390/FI16060196.
11. N. Nahidi, “Regulatory and Compliance Issues in AI-Driven Financial Services,” Contributions to Finance and Accounting, vol. Part F1945, pp. 169–196, 2026, doi: 10.1007/978-3-032-18109-1_8/FIGURES/2.
12. B. U. I. Khan, K. W. Goh, A. R. Khan, M. F. Zuhairi, and M. Chaimanee, “Resource management and secure data exchange for mobile sensors using Ethereum blockchain,” Symmetry (Basel)., vol. 17, no. 1, p. 61, Jan. 2025, doi: 10.3390/sym17010061.
13. C. Zhang, X. Liu, and Y. Gao, “Does AI contribute to systemic risk reduction in banks- Evidence from double-debiased machine learning,” Finance Research Open, vol. 2, no. 1, p. 100093, Mar. 2026, doi: 10.1016/J.FINR.2026.100093.
14. B. M. Henrique, V. A. Sobreiro, and H. Kimura, “Literature review: machine learning techniques applied to financial market prediction,” Expert Syst Appl, vol. 124, pp. 226–251, Jun. 2019, doi: 10.1016/j.eswa.2019.01.012.
15. R. Najem, A. Bahnasse, M. Fakhouri Amr, and M. Talea, “Advanced AI and big data techniques in E-finance: a comprehensive survey,” Discover Artificial Intelligence 2025 5:1, vol. 5, no. 1, pp. 102-, Jun. 2025, doi: 10.1007/S44163-025-00365-Y.
16. B. S. Samantray and K. H. K. Reddy, “A Federated Learning Approach Towards Hybrid Blockchain, Quantum-Key-Encryption based Distributed System: A Futuristic Healthcare Architecture for Smart Cities,” Blockchain: Research and Applications, p. 100385, Sep. 2025, doi: 10.1016/J.BCRA.2025.100385.
17. S. Mousavian and S. J. Miah, “Review of artificial intelligence-based applications for money laundering detection,” Intelligent Systems with Applications, vol. 27, p. 200572, Sep. 2025, doi: 10.1016/J.ISWA.2025.200572.
18. S. Ahmed, M. M. Alshater, A. El Ammari, and H. Hammami, “Artificial intelligence and machine learning in finance: A bibliometric review,” Res. Int. Bus. Finance, vol. 61, Oct. 2022, doi: 10.1016/j.ribaf.2022.101646.
19. S. Goundar and I. Gondal, “AI-Blockchain Integration for Real-Time Cybersecurity: System Design and Evaluation,” Journal of Cybersecurity and Privacy 2025, Vol. 5, Page 59, vol. 5, no. 3, p. 59, Aug. 2025, doi: 10.3390/JCP5030059.
20. Kingsley David Onyewuchi Ofoegbu, Olajide Soji Osundare, Chidiebere Somadina Ike, Ololade Gilbert Fakeyede, and Adebimpe Bolatito Ige, “Real-Time Cybersecurity threat detection using machine learning and big data analytics: A comprehensive approach,” Computer Science & IT Research Journal, vol. 4, no. 3, pp. 478–501, Dec. 2023, doi: 10.51594/CSITRJ.V4I3.1500.
21. Y. Chen, C. Zhao, Y. Xu, C. Nie, and Y. Zhang, “Deep Learning in Financial Fraud Detection: Innovations, Challenges, and Applications,” Data Science and Management, Aug. 2025, doi: 10.1016/J.DSM.2025.08.002.
22. E. Gadimov and E. Birihanu, “Real-time suspicious detection framework for financial data streams,” International Journal of Information Technology 2025, pp. 1–17, Apr. 2025, doi: 10.1007/S41870-025-02529-6.
23. A. V. Dorogush, V. Ershov, and A. Gulin, “CatBoost: gradient boosting with categorical features support,” Oct. 2018, Accessed: Jun. 30, 2026. [Online]. Available: http://arxiv.org/abs/1810.11363
24. M. Cai, “Detecting Financial Fraud and Anomalies in Corporate Transactions with Convolutional Neural Networks,” International Journal of Computational Intelligence Systems 2026 19:1, vol. 19, no. 1, pp. 186-, Mar. 2026, doi: 10.1007/S44196-026-01275-2.
25. C. Wang, M. Wang, X. Wang, L. Zhang, and Y. Long, “Multi-Relational Graph Representation Learning for Financial Statement Fraud Detection,” Big Data Min. Analytics, vol. 7, no. 3, pp. 920–941, 2024, doi: 10.26599/bdma.2024.9020013.
26. A. Zafar, “Privacy as institutional design: A legal-technological analysis of CBDC governance and compliance,” Computer Law & Security Review, vol. 60, p. 106258, Apr. 2026, doi: 10.1016/J.CLSR.2025.106258.

27. S. Wang, “Federated Semantic Web Framework for Enhanced Financial Risk Control and Data Analysis,” Int. J. Semant. Web Inf. Syst., vol. 21, no. 1, Dec. 2024, doi: 10.4018/IJSWIS.367469.
28. A. A. Ahmed and O. O. Alabi, “Secure and Scalable Blockchain-Based Federated Learning for Cryptocurrency Fraud Detection: A Systematic Review,” IEEE Access, vol. 12, pp. 102219–102241, 2024, doi: 10.1109/access.2024.3429205.
29. E. Makarim, “LEGAL IMPLICATIONS OF BLOCKCHAIN TECHNOLOGY FOR THE INDONESIAN NATIONAL FINANCIAL SYSTEM AND MONETARY AUTHORITIES,” Journal of Central Banking Law and Institutions, vol. 5, no. 2, pp. 377–400, May 2026, doi: 10.21098/JCLI.V5I2.542.
30. B. K. Mohanta, A. I. Awad, T. Elsaka, H. Kheddar, and E. Baraka, “Smart-contract-based blockchain-enabled decentralized scheme for improving smart-grid security,” Internet of Things, vol. 34, p. 101811, Nov. 2025, doi: 10.1016/J.IOT.2025.101811.
31. A. Amato, M. R. Machado, and J. L. R. Moreira, “Leveraging early warning systems and customer segmentation to identify business opportunities in the financial industry,” International Journal of Information Management Data Insights, vol. 6, no. 2, p. 100428, Dec. 2026, doi: 10.1016/J.JJIMEI.2026.100428.
32. B. S. Samantray and K. H. K. Reddy, “A Federated Learning Approach Towards Hybrid Blockchain, Quantum-Key-Encryption based Distributed System: A Futuristic Healthcare Architecture for Smart Cities,” Blockchain: Research and Applications, p. 100385, Sep. 2025, doi: 10.1016/J.BCRA.2025.100385.
33. S. M. A. K. Firouzabadi, M. T. Taghavifard, S. K. Sajjadi, and J. B. Soufi, “A multi-objective optimisation model for assignment of service to bank customers by using data mining and simulation,” International Journal of Electronic Customer Relationship Management, vol. 11, no. 3, pp. 237–255, 2018, doi: 10.1504/IJECRM.2018.093766.
34. L. He, “Blockchain-enabled verifiable federated learning with multi-dimensional incentive and dynamic consensus for privacy-sensitive domains,” Discover Computing 2026 29:1, vol. 29, no. 1, pp. 383-, Jun. 2026, doi: 10.1007/S10791-026-10289-5.
35. W. Song, X. Zhu, S. Ren, W. Tan, and Y. Peng, “A hybrid blockchain and machine learning approach for intrusion detection system in Industrial Internet of Things,” Alexandria Engineering Journal, vol. 127, pp. 619–627, Aug. 2025, doi: 10.1016/J.AEJ.2025.05.030.
36. J. Daníelsson, R. Macrae, and A. Uthemann, “Artificial intelligence and systemic risk,” J. Bank. Financ., vol. 140, p. 106290, Jul. 2022, doi: 10.1016/J.JBANKFIN.2021.106290.
37. A. Primec, J. Belak, and M. Čufar, “Artificial Intelligence as an Emerging Risk Dimension in Corporate Sustainability Reporting: A Legal and Governance Perspective,” Sustainability 2026, Vol. 18, Page 2278, vol. 18, no. 5, p. 2278, Feb. 2026, doi: 10.3390/SU18052278.
38. M. Bouvard, P. Chaigneau, and A. De Motta, “Transparency in the Financial System: Rollover Risk and Crises,” Journal of Finance, vol. 70, no. 4, pp. 1805–1837, Aug. 2015, doi: 10.1111/jofi.12270.
39. Venkata Raja Ravi Kumar Gelle, “Enhancing financial security: AI-driven anti-money laundering (AML) and compliance monitoring in the banking sector,” World Journal of Advanced Research and Reviews, vol. 25, no. 1, pp. 2462–2476, Jan. 2025, doi: 10.30574/WJARR.2025.25.1.0365.
40. S. A. Ajagbe, S. Majola, and P. Mudali, “Comparative analysis of machine learning algorithms for money laundering detection,” Discover Artificial Intelligence 2025 5:1, vol. 5, no. 1, pp. 144-, Jul. 2025, doi: 10.1007/S44163-025-00397-4.
41. L. Ren et al., “Dynamic graph neural network-based fraud detectors against collaborative fraudsters,” Knowl. Based. Syst., vol. 278, p. 110888, Oct. 2023, doi: 10.1016/J.KNOSYS.2023.110888.