Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

Autonomous Federated Compliance Intelligence for Global Anti-Financial Crime Networks

Mallikarjun Reddy Gouni

University of Illinois Springfield, USA

ARTICLE INFO

Received: 01 Nov 2025

Revised: 16 Dec 2025

Accepted: 27 Dec 2025

ABSTRACT

Financial crime detection faces unparalleled challenges as criminal networks exploit digital payment channels, cryptocurrency platforms, and cross-border transaction systems outside traditional monitoring frameworks. In this respect, AFCI introduces a novel framework for federated machine learning, regulatory reasoning engines, and real-time risk propagation analytics to build unified global privacy-preserving anti-crime intelligence ecosystems. The framework lets organizations train collaborative models with decentralized institutions, safely aggregating information from multiple parties without sharing sensitive transaction data by means of secure aggregation protocols and differential privacy mechanisms. Large language models coupled with knowledge graphs automate the processes of regulatory interpretation and rule generation, and graph neural networks enable the detection of coordinated criminal activities on a large scale in transaction networks through temporal message passing mechanisms. Reinforcement learning agents continuously optimize detection policies to balance the identification of genuine threats against the goal of minimizing false alarms. The framework bridged critical gaps in cross-border compliance coordination and empowered institutions to develop shared detection capabilities in support of data localization requirements and an array of diverse regulatory frameworks. Long-term security of privacy-preserving federated computation would be guaranteed with post-quantum cryptography. This convergence of advanced technologies allows next-generation financial crime prevention systems to remain effective against evolving criminal methodologies while preserving fundamental privacy rights.

Keywords: Federated Learning, Financial Crime Detection, Graph Neural Networks, Regulatory Intelligence Engines, Privacy-Preserving Machine Learning

1. Introduction to Challenges of Financial Crime Detection in Modern Banking Ecosystems

Detection of financial crimes within contemporary banking environments is faced with complex challenges due to the emerging sophistication of criminal methodologies and the fragmented nature of compliance systems across institutions. Modern money laundering operations have advanced activities that have blinded sight on digital payment channels, cryptocurrency platforms, and cross-border transactions operating outside the reach of conventional monitoring frameworks. The integration stage of money laundering, where illicit funds are reintroduced into the legitimate economy through complex layering techniques, presents a particular challenge in terms of detection, since criminals use multiple intermediary accounts and jurisdictions in order to obscure the origin of funds [1]. Detection systems need to handle vast transaction volumes and identify subtle patterns that distinguish legitimate business activities from structured attempts at evading regulatory scrutiny.

Where digital transformation has elevated the art of detection by allowing the analysis of rich metadata and behavioral patterns derived from transactions, it has also created new attack vectors that criminals can exploit through synthetic identity fraud, account takeover schemes, and automated

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

bot networks operating thousands of fraudulent transactions before any detection system could trigger a response. The proliferation of mobile banking applications and instant payment systems has also shrunk timelines for transactions, demanding detection capabilities in near real time, capable of assessing risk in milliseconds without compromising accuracy standards that minimize friction for customers due to false-positive alerts [2]. Financial institutions face the twin imperative of keeping customers safe from fraud while delivering seamless digital experiences that meet consumer expectations created by technology leaders in other industries.

Traditional compliance architectures are focused on isolated institutional boundaries that preclude the collaborative intelligence sharing necessary for the detection of sophisticated criminal networks operating across multiple financial institutions. Regulatory frameworks often prohibit the direct sharing of customer transaction data between competing institutions, creating information asymmetries that criminals actively exploit by distributing their activities across multiple banks to remain below individual institutional detection thresholds. The lack of standardized data schema and similarity in the manner in which risk is classified by institutions further complicates the ability to come up with a coordinated response to new typologies of financial crime. These organizational constraints create the need to employ new methods that allow collaborative detection without violating privacy restrictions, competitive sensitivities, and regulatory demands that inform information sharing in the financial services industry.

2. Privacy-Preserving Compliance Intelligence - Federated Learning Architecture

Federated learning is a paradigm of distributed machine learning, whereby several organizations can jointly train models without necessarily having to centralize sensitive data. This represents a solution to the fundamental challenges of privacy and data governance in financial services. Iterative rounds are the basis of the architectural framework within which the participating institutions train local models on proprietary transaction data, share only model parameters or gradient updates with a central coordination server that aggregates these contributions into a global model. In contrast to centralized machine learning, which requires that all training data be collected into a single repository, the approach keeps data resident across institutional boundaries, although the models can leverage patterns observed across the full network [3]. The federated approach marries up naturally with the regulatory and competitive constraints inhibiting the pooling of customer data by financial institutions, while they collectively enable the improvement of detection capabilities.

The aim is to apply advanced cryptographic protocols that prohibit any participating institution from inducing information about the data of other participants from shared model updates. The purpose of differential privacy mechanisms is to add calibrated statistical noise to model updates before aggregation, providing mathematical guarantees that individual transactions or customer patterns cannot be reconstructed from the global model parameters. The privacy budget, quantified through epsilon and delta parameters in differential privacy theory, is thus a key determinant of the trade-off between model utility and privacy protection, where small values of epsilon yield stronger privacy guarantees at higher costs in terms of model accuracy. Secure aggregation protocols typically use cryptographic techniques like secret sharing and homomorphic encryption to let a coordination server compute aggregate model updates without observing individual institutional contributions, thereby ensuring no party, including the coordination infrastructure, can access raw parameter updates from any single participant.

Deploying a realistic federated learning architecture against financial crime challenges such as statistical heterogeneity of participating institutions, communication efficiency in a geographically distributed network, and resilience to adversarial participants who may try poisoning the global model, all stand before financial institutions serving diverse customer populations with different

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

transaction patterns, maintaining different core banking systems with different data representations, and heterogeneous fraud rates generating non-identically distributed data across the federation. Adaptive optimization algorithms and personalized model layers enable the latter to maintain institution-specific model customizations for their operating environment, while contributing to shared pattern recognition capabilities valuable to all participants for the detection of novel methodologies employed by criminals anywhere in the network.

Component	Function	Privacy Protection Method	Key Benefit
Local Model Training	Train on institutional data	Data remains at the source	Preserves data sovereignty
Parameter Aggregation	Combine model updates	Differential privacy noise	Prevents pattern reconstruction
Secure Coordination	Aggregate without exposure	Homomorphic encryption	Zero raw data access
Adaptive Optimization	Handle data heterogeneity	Personalized model layers	Institution-specific customization

Table 1: Federated Learning Architecture Components and Privacy Mechanisms [3, 4]

3. Large Language Model-based and Knowledge Graph-based Regulatory Intelligence Engines

The complexity of international financial regulations, comprising several thousand jurisdictional requirements that are continuously updated using domain-specific terminology, presents enormous challenges to operating compliantly across international banking networks. Large language models have impressive capabilities in processing regulatory text, extracting compliance obligations, and mapping these abstract requirements to operational controls, thanks to training on extensive corpora of legal and regulatory documents. These models can analyze newly published regulatory guidance for the extraction of relevant provisions that apply to specific transaction types or customer segments and generate interpretations to support automated rule generation for transaction monitoring systems [5]. Large language models are able to process regulatory text in multiple languages and reconcile the varying terminologies used across different jurisdictions when describing similar compliance concepts.

Knowledge graphs provide structured representations that encode the relationships between regulatory concepts, risk factors, transaction types, and compliance obligations in formats supportive of sophisticated reasoning operations and consistency validation. Large language model-based regulatory parsing integrated with knowledge graph structures yields a hybrid system where the language model extracts information from unstructured regulatory text and the knowledge graph maintains logical relationships that enable automated inference of compliance implications [6]. If new guidance documents are published by regulatory authorities, the language model processes the text to identify new compliance requirements, risk indicators, or prohibited activities, then updates the knowledge graph with nodes representing these concepts and edges encoding their relationships to existing regulatory structures and operational controls.

Regulatory intelligence needs mechanisms for translating these abstract compliance requirements into executable detection rules and monitoring parameters within the institutional transaction surveillance systems. Language models generate candidate rule specifications by analyzing regulatory descriptions of suspicious activity patterns and then map these to available transaction attributes in

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

institutional databases, such as transaction amounts, relationships between counterparties, geographic locations, and temporal patterns. Knowledge graph reasoning validates proposed rules against existing controls for logical consistency, determines any potential conflicts or coverage gaps, and recommends optimal parameter configurations based on historical effectiveness data and false positive rates. This automated interpretation capability reduces implementation timelines from months to weeks while improving consistency in how similar requirements are operationalized across different jurisdictions within a multi-national institution's compliance program.

Processing Stage	Technology Component	Input Type	Output Type	Implementation Impact
Regulatory Parsing	Large Language Models	Unstructured regulatory text	Extracted compliance obligations	Multi-language processing
Knowledge Encoding	Knowledge Graphs	Regulatory concepts	Structured relationships	Automated inference
Rule Generation	LLM Analysis	Suspicious activity descriptions	Candidate detection rules	Transaction attribute mapping
Validation	Graph Reasoning	Proposed rules	Validated parameters	Consistency checking

Table 2: Regulatory Intelligence Engine Processing Workflow [5, 6]

4. Real-time Risk Propagation Analysis using Graph Neural Networks

Financial crime often materializes in the form of coordinated activities across networks of entities that look legitimate as individual units but together reveal patterns indicative of money laundering, fraud rings, or sanctions evasion schemes. Graph neural networks directly operate on the graph-structured data representation of transaction networks and learn node embeddings that capture the characteristics of the individual entity, as well as its structural relationships to the wider network through message-passing mechanisms aggregating information from neighboring nodes. The detection of suspicious subgraphs representing criminal networks that would remain invisible to traditional machine learning approaches analyzing transactions in isolation is thus enabled by such architectures. The core novelty of graph neural networks lies in their ability to propagate information along multiple hops in transaction graphs, showing how entities separated by several intermediary nodes show correlated suspicious behaviors indicative of coordinated criminal activity.

Temporal graph neural networks extend static graph architectures with time-aware message-passing mechanisms to model how transaction networks evolve, and suspicious patterns propagate through these dynamic structures. These temporal models leverage recurrent neural network components or temporal attention mechanisms to maintain representations of network state history, which allows them to foresee emerging concentrations of risk, predict the likelihood of entities becoming implicated in suspicious activity based on their network position and temporal behavior, and estimate propagation velocities indicating the speed at which illicit funds may traverse financial systems. It is in this regard that the temporal dimension becomes very important for the detection of techniques such as trade-based money laundering, in which value transfer occurs through extended time periods via multiple sequences of transactions among networks of complicit entities at mispriced values, and

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

the identification of normally dormant accounts suddenly becoming active as part of money mule networks.

Real-time deployment of graph neural network-based risk propagation systems is extremely challenging due to the fact that financial transaction networks contain millions of entities that are interrelated through billions of edges in continuously changing payment relationships as new transactions come in. Economical architectures make use of neighborhood sampling strategies, which concentrate message passing operations on the most relevant subgraphs surrounding entities of interest; incremental update mechanisms that recompute embeddings only for affected network regions when new transactions arrive; and distributed computing frameworks that parallelize graph operations across multiple processing nodes. Attention mechanisms within graph neural network architectures automatically learn to weight network relationships by their relevance for risk assessment; hence, models can scale to massive networks by focusing computational resources on those connections that most strongly indicate coordinated suspicious activity while efficiently processing high-velocity transaction streams characteristic of modern payment systems.

Network Type	Detection Capability	Temporal Dimension	Computational Strategy	Application Scenario
Static GNN	Individual entity patterns	Single timepoint	Message passing	Basic fraud detection
Temporal GNN	Network evolution tracking	Time-series analysis	Recurrent components	Money laundering networks
Multi-hop GNN	Cross-entity correlation	Historical relationships	Neighborhood sampling	Sanctions evasion schemes
Attention-based GNN	Risk-weighted connections	Real-time streams	Incremental updates	Trade-based laundering

Table 3: Graph Neural Network Capabilities for Risk Detection [7, 8]

5. Adaptive Compliance Systems using Reinforcement Learning Mechanisms

Reinforcement learning provides frameworks for developing autonomous agents that learn optimal decision policies through interaction with their operating environments, receiving rewards or penalties based on action outcomes that guide iterative policy improvement. In compliance detection applications, reinforcement learning agents learn to balance competing objectives of maximizing detection of genuine suspicious activity while minimizing false positive alerts that consume investigative resources and potentially impact customer experience through unnecessary account restrictions or transaction delays [9].

The agent observes states representing transaction characteristics, entity profiles, historical behavior patterns, and current risk assessments, then selects actions related to alert generation, risk scoring, or investigative prioritization, receiving rewards that reflect detection accuracy, investigation efficiency, and regulatory compliance metrics. Formulation of reinforcement learning problems for compliance applications requires careful construction of state spaces, action spaces, and reward structures so that the learned policies are aligned with institutional risk tolerances and regulatory obligations. State representation needs to capture relevant features from the high-dimensional transaction data while keeping computation tractable; this is usually achieved by employing dimensionality reduction

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

techniques or learned representations through neural networks to compress raw transaction attributes into compact vectors of state. The action space includes, in general, discrete decisions over alert generation thresholds, continuous adjustments over parameters of risk scoring, and sequential decisions on investigation resource allocations across multiple pending cases [10].

The reward structure has to account not only for binary detection outcomes but also consider the costs of different types of errors, where missed suspicious activity may lead to regulatory penalties and reputational damage, and excessive false positives burden the investigative teams and potentially affect customer relationships. Safe exploration, sample efficiency, and integration with human expertise are some of the challenges that have to be addressed when reinforcement learning-based adaptive systems are deployed in production for compliance. Compliance applications cannot afford extended exploration phases that might miss critical suspicious activity while agents learn effective policies, so off-policy learning approaches are needed to train agents on historical data before deployment, along with safe exploration constraints that maintain minimum detection thresholds during online learning. Multi-armed bandit formulations provide frameworks for balancing exploration of alternative detection strategies against exploitation of known effective approaches and allow for continuous improvement while maintaining high levels of operational effectiveness. Humans-in-the-loop mechanisms allow compliance officers to provide feedback on agent decisions, correct misclassifications, and guide policy learning toward outcomes that respect institutional risk culture and regulatory expectations, all while reducing the manual burden of ongoing rule maintenance and threshold tuning.

System Element	Design Consideration	Implementation Approach	Optimization Goal	Human Integration
State Space	Transaction features	Dimensionality reduction	Computational tractability	Feature relevance
Action Space	Detection decisions	Threshold adjustments	Alert optimization	Decision validation
Reward Structure	Detection accuracy	Multi-objective balancing	Error cost minimization	Feedback incorporation
Exploration Strategy	Policy learning	Off-policy training	Safe detection thresholds	Guidance provision

Table 4: Reinforcement Learning Components in Adaptive Compliance Systems [9, 10]

6. Cross-Border Applications and Future Directions in Financial Crime Prevention

Cross-border financial crime prevention requires coordination mechanisms that enable international collaboration in respect of national sovereignty, data localization requirements, and varying regulatory frameworks across jurisdictions. Financial crime networks deliberately exploit regulatory differences and coordination challenges between national authorities by structuring activities across multiple jurisdictions to complicate investigation and prosecution efforts. International remittance networks pose particularly challenging environments for compliance monitoring, given that they serve as channels of legitimate fund transfers for millions of individuals supporting families across borders, while serving as a potential channel for illicit value movement that criminals exploit through structuring techniques and exploitation of informal value transfer systems [11].

Federated compliance intelligence frameworks allow institutions operating in multiple jurisdictions to develop shared detection capabilities without centralizing customer data in ways that would violate

Journal of Information Systems Engineering and Management
2025, 10(63s)
e-ISSN: 2468-4376
https://www.jisem-journal.com/

Research Article

data protection regulations or compromise competitive positions. Anti-money laundering measures have undergone a major transition to reflect the latest threats in the digital world in the form of digital assets, virtual currencies, and decentralized financial protocols that can be used to transfer value without necessarily going through a conventional banking system in a large financial hub.

Financial cryptocurrency exchanges, wallet providers, and some decentralized finance protocols have been added to the definition of the financial institutions to which the anti-money laundering requirements apply. These entities have been required to implement customer due diligence, transaction monitoring, and suspicious activity reporting programs comparable to those maintained by traditional banks [11]. The unique pseudonymous nature of blockchain transactions brings unique challenges for performing identity verification and beneficial ownership, while the global accessibility and rapid settlement speed of digital assets have empowered bad actors to move value across borders more quickly than traditional correspondent banking networks. Compliance systems must integrate blockchain analytics capabilities that track fund flows across multiple distributed ledger networks with traditional financial crime detection to maintain visibility of value movement across conventional and digital financial systems. Future steps in autonomous compliance intelligence will be toward post-quantum cryptographic techniques that retain security guarantees regarding privacy-preserving, federated computation in light of emerging quantum computing capabilities. Current cryptographic protocols forming the basis for secure aggregation and homomorphic encryption rely on computational hardness assumptions about factoring large integers or solving discrete logarithm problems that quantum computers could conceivably break using Shor's algorithm and other related quantum algorithms [12].

Alternative foundations are lattice-problem-based and code-based cryptography, as well as hash-based signatures, which have post-quantum cryptography security against quantum adversaries, thus making it possible to continue building privacy-conscious collaborative compliance systems as quantum computing technology advances. The combination of quantum-resistant cryptography and the development of explainable artificial intelligence, multi-dimensional data fusion, and adaptive regulatory technology will enable financial crime prevention systems of the next generation that will not only remain effective against the evolving criminal tactics but will also comply with the basic privacy rights and regulatory mandates.

Conclusion

Autonomous Federated Compliance Intelligence makes up a transformational paradigm for global financial crime prevention; through privacy-preserving collaborative intelligence, it solves fundamental limitations with traditional compliance architectures. Integration of federated learning lets institutions build shared detection capabilities without aggregating sensitive customer data in a central location, whereas differential privacy mechanisms give mathematical guarantees, protecting individual transaction information. Regulatory intelligence engines with large language models and knowledge graphs enable automated interpretation of complex, multi-jurisdictional compliance requirements. This reduces implementation timelines and increases consistency across the boundaries of the institution. Temporal-aware graph neural networks make it possible to detect more complex criminal networks spanning across institutions and jurisdictions and displaying coordinated suspicious activity otherwise unnoticed by isolated surveillance. The reinforcement learning processes provide ongoing adjustment to the changes in the criminal patterns. Optimal detection policies emerge by interaction with the operational environment, while maintaining human oversight. International remittance networks and digital asset ecosystems are emerging as some of the best fits for cross-border applications, where traditional compliance frameworks are finding it challenging due to jurisdictional complexity. The future developments had to be made in such a way that they ensure

that these systems, which are founded on post-quantum cryptography, explainable artificial intelligence, and multimodal data fusion, are in line with the threats that are emerging and that their services would continue to assure the evolution of privacy rights and regulatory demands within the global financial services.

References

[1] Mamunur R Raja et al., "Detecting and Preventing Money Laundering Using Deep Learning and Graph Analysis," International Journal of Advanced Computer Science and Applications, Vol. 16, No. 6, 2025. [Online]. Available: https://thesai.org/Downloads/Volume16No6/Paper_1-Detecting_and_Preventing_Money_Laundering.pdf

[2] Fraud.net, "Fraud detection in banking: Key challenges and solutions," 2024. [Online]. Available: https://www.fraud.net/resources/fraud-detection-in-banking-key-challenges-and-solutions#the-digital-era-has-become-a-curse-and-a-blessing-at-the-same-time-

[3] Qiang Yang et al., "Federated machine learning: Concept and applications," ACM Transactions on Intelligent Systems and Technology (TIST), Volume 10, Issue 2, 2019. [Online]. Available: https://dl.acm.org/doi/10.1145/3298981

[4] Kang Wei et al., "Federated Learning With Differential Privacy: Algorithms and Performance Analysis," IEEE Transactions on Information Forensics and Security, Volume 15, 2020. [Online]. Available: https://ieeexplore.ieee.org/document/9069945

[5] Ruben Kruiper et al., "A platform-based Natural Language processing-driven strategy for digitalising regulatory compliance processes for the built environment," Advanced Engineering Informatics, Volume 62, Part B, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S147403462400301X

[6] Zhuo Chen et al., "Knowledge Graphs Meet Multi-Modal Learning: A Comprehensive Survey," arXiv:2402.05391, 2024. [Online]. Available: https://arxiv.org/abs/2402.05391

[7] Jie Zhou et al., "Graph neural networks: A review of methods and applications," AI Open, Volume 1, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2666651021000012

[8] Emanuele Rossi et al., "Temporal Graph Networks for Deep Learning on Dynamic Graphs," arXiv:2006.10637, 2020. [Online]. Available: https://arxiv.org/abs/2006.10637

[9] Vincent Francois-Lavet et al., "An Introduction to Deep Reinforcement Learning," arXiv:1811.12560, 2018. [Online]. Available: https://arxiv.org/abs/1811.12560

[10] Olivier Sigaud and Olivier Buffet, "Markov Decision Processes in Artificial Intelligence," Wiley, 2013. [Online]. Available: https://onlinelibrary.wiley.com/doi/book/10.1002/9781118557426

[11] KYC360, "Anti-Money Laundering Regulations: A Comprehensive Guide," 2024. [Online]. Available: https://kyc360.com/knowledge-hub/resources/anti-money-laundering-regulations-a-comprehensive-guide

[12] European Union Agency for Cybersecurity (ENISA), "Post-quantum cryptography: Current state and quantum mitigation," 2021. [Online]. Available: https://www.enisa.europa.eu/sites/default/files/publications/ENISA%20Report%20-%20Post-Quantum%20Cryptography%20Current%20state%20and%20quantum%20mitigation-V2.pdf