

AI-Driven Financial Crime Analytics: Enhancing Compliance Through Predictive Modelling and Blockchain Forensics

Babul Bijan Mandal¹, Nikhil Teja Gurram², Dr. Aparna Pavani³, Appa Rao Nagubandi⁴, Dr. Roopajayasingh⁵

¹assistant Professor, Information Technology, Eshwantrao Chavan College Of Engineering, Nagpur, Maharashtra,

Email ID : mandalbabul2020@gmail.com

²Technical Manager, Software Engineer, GLA University, HCLTech,

Email ID : moneyprash@gmail.com

³Associate Professor, School of Business, Alliance University, Bangalore, Karnataka,

Email ID : aparnapavani@gmail.com

⁴Lead Software Engineer, Apparao.

Email ID : Nb@Gmail.Com, Orcid Id: 0009-0005-8424-7071

⁵asp/Ece, V.S.B.College Of Engineering Technical Campus, COIMBATORE

ABSTRACT

AI-driven financial crime analytics is redefining regulatory compliance by shifting detection from reactive rule-based monitoring to proactive probabilistic intelligence. Financial crime—money laundering, sanctions evasion, synthetic identity fraud, ransomware financing, and cross-chain asset obscuring has become faster, more automated, and technically sophisticated than current compliance infrastructure can handle. This study proposes a novel compliance-first analytics architecture that fuses predictive modeling, blockchain transaction forensics, graph intelligence, smart-contract tracing, risk-propagation modeling, and real-time anomaly profiling to strengthen compliance accuracy, auditability, and enforcement readiness. Using supervised learning, “graph neural networks (GNN)”, transformer-based behavioral profiling, and multi-chain forensic tagging, the framework detects illicit capital movement earlier, maps attribution failures across decentralized ledgers, and improves compliance decision quality while preserving explainability for regulators. The results indicate that AI embeddings improve fraud-pattern discovery by 220–300%, reduce false compliance flags by 45–55%, and enable 3–6-week earlier risk detection versus traditional compliance engines. The study contributes a scalable, regulator-friendly, automated compliance system capable of tracing illicit flows even when adversaries use mixers, privacy chains, or cross-chain bridges..

Keywords: Financial crime analytics; Predictive compliance; Blockchain forensics; GNN tracing; AML automation; Sanctions intelligence; Ransomware forensics; Crypto compliance..

1. INTRODUCTION:

Financial crime has evolved into a parallel financial system digitally native, decentralized, pseudonymous, automated, and engineered to bypass compliance. The assumption that financial crime leaves linear trails is dead. Modern crime capital moves through cross-chain bridges, smart-contract routers, mixers, privacy chains, layered wallets, and institutional liquidity pools at machine speed. Bitcoin, Ethereum, stablecoins, and exchange tokens are now default settlement layers for illicit finance, ransomware liquidity exit, sanctions evasion, dark-web commerce, and AI-assisted synthetic identity fraud.

Legacy compliance systems fail because they rely on static rules, delayed reporting, manual forensic tagging, and weak graph intelligence. They flag transactions, not behavioral entropy. They measure deviation, not criminal optimization strategies. Illicit finance today resembles engineered network attacks distributed, obfuscated, adaptive, self-healing, and feedback-driven. Computational intelligence is no longer optional; it is the only realistic compliance defense. This study argues a blunt position: compliance must adopt the same weapons

as criminals automation, prediction, probabilistic modeling, graph intelligence, identity embeddings, and real-time forensic state awareness. We introduce a compliance-optimized AI architecture that detects financial crime earlier, traces deeper, reduces false flags, and outputs regulator-ready evidence graphs. The objective is not investment modeling, sentiment analytics, or market forecasting. The objective is compliance dominance precision, attribution, auditability, and regulatory usability.

2. RELATED WORKS

Financial crime analytics research historically relied on rule-based compliance engines, threshold monitoring, and manually curated “suspicious-activity reports (SAR)”. These systems were effective when fraud followed linear banking rails, slower execution, and isolated jurisdictions, but collapse against modern digital crime automation [1], [2]. Early machine-learning applications introduced anomaly classifiers and clustering models to detect deviations in payment behavior, yielding improved detection but producing poor explainability and high false-positive rates making them unsuitable for regulatory evidence pipelines [3], [4].

The rise of cryptocurrencies forced compliance research into blockchain transaction tracing, wallet attribution heuristics, and graph-based link analysis. Studies demonstrated that Bitcoin and stablecoins frequently serve as liquidity exit layers for illicit capital, enabling ransomware cash-outs, sanctions evasion, darknet settlements, and cross-border laundering [5], [6]. Blockchain forensic research introduced address clustering, transaction-hop tracking, and entity tagging, confirming that crypto crime flows behave as engineered capital-routing attacks rather than statistical anomalies [7].

More advanced compliance research incorporated “Graph Neural Networks (GNNs)”, temporal propagation models, and deep-learning sequence encoders to track risk diffusion across wallets, bridges, and smart contracts. These models identified systemic crime hubs, evolving attack subgraphs, and obfuscation signatures such as mixers, cross-chain bridges, and contract layering [8], [9]. Transformer-based behavioral models further enhanced identity fraud detection and laundering fingerprinting, showing strong gains in pattern discovery but requiring hybrid evidence layers to maintain interpretability for regulators [10], [11].

Recent frameworks explored cross-chain forensic automation and probabilistic attribution modeling, reinforcing that illicit finance is adaptive, feedback-driven, and multi-layered across decentralized ledgers [12], [13]. Hybrid compliance-intelligence models attempted to balance predictive power and explainability, yet most systems remain post-execution detectors rather than pre-exit predictive compliance engines, leaving a major operational gap [14], [15]. This paper takes a firm stance: existing models detect crime late, trace chains manually, and fail to unify predictive compliance with automated blockchain forensic state awareness. The literature lacks a compliance-native system that predicts illicit capital movement before attribution collapse, regulatory breach, or liquidity exit. That gap is exactly what AI + GNN + blockchain forensics must fill next [16], [17]

3. METHODOLOGY

3.1 Research Design

This research adopts a **compliance-first computational architecture**, built on the premise that financial crime behaves as an adaptive, multi-layered, multi-chain network intrusion rather than a static anomaly. The study combines predictive illicit-flow intelligence, identity-level wallet embeddings, and blockchain forensic state tracking to generate outputs that can be audited and operationalized by compliance teams and regulators. Unlike legacy rule-centric fraud monitoring, the design prioritizes early attribution, structural risk discovery, evidentiary traceability, and temporal risk surface awareness. The computational pipeline and the evidence-output layer are jointly evaluated for performance improvements over traditional AML and transaction-flagging systems.

3.2 Illicit Capital Chain and Dataset Selection

Illicit capital movement was tracked across “*Bitcoin (BTC), Ethereum (ETH), Binance Smart Chain (BSC), Polygon, and stablecoin ledgers (USDT/USDC)*” to reflect real-world laundering preferences that favor liquidity depth, exchange off-ramps, and cross-jurisdiction bridges. The study’s crime-flow dataset includes regulator-flagged wallets, sanctions-linked address clusters, ransomware exit wallets, contract-level mixer fingerprints, multi-chain bridge transaction logs, and synthetic-identity address networks commonly associated with identity fraud and layered laundering. These assets were selected not for investment analysis, but for their forensic and compliance relevance liquidity centralization, regulatory tagging, and capital exit influence.

3.3 Wallet Behavior Transformation and Feature Engineering

Instead of price or financial returns, raw wallet activity was converted into **behavior-return embeddings**, capturing time-normalized patterns of transaction bursts, token swap sequences, smart-contract hops, gas-fee irregularity signatures, cross-chain routing frequency, and wallet age graphs that indicate identity stability or synthetic origins. Noise introduced by chain congestion was filtered using entropy-confidence normalization, ensuring that illicit flow patterns were separated from network delays or fee-spike distortions that typically inflate false compliance flags. Stationarity checks were applied to wallet-behavior sequences to ensure the model captures persistent crime-routing structure rather than transient system congestion.

3.4 Computational Forensic Modeling and Graph Construction

Directional crime influence among wallet clusters was modeled using **Vector Autoregression (VAR) on behavioral embeddings** to measure asymmetric capital shock transmission and illicit routing influence. To capture nonlinear diffusion and contract-level obscuring, **Graph Neural Networks (GNN)** were deployed to propagate compliance-risk signals across wallets, bridge contracts, and mixer subgraphs. Smart-contract fingerprinting was applied to identify mixer-level capital obscuring, while cross-chain bridge routing detection mapped jurisdiction-spanning layering trails. The final output layer constructs a **weighted illicit finance graph**, where nodes represent wallet clusters, exchanges, bridge contracts, and sanctioned identity groups, and edges represent directional influence intensity weighted by attribution confidence and spillover strength. Centrality scoring identifies compliance hubs, fragility nodes, and cascading risk potential when shocks originate from dominant liquidity layers.

3.5 Evidence Output and Regulatory Audit Layer

To ensure regulatory usability, the evidence graph output system enforces **single attribution per forensic edge**, timestamps for chain-of-custody, contract-level forensic tags, wallet-identity provenance markers, sanctioned address node labeling, and regulator-friendly audit trails that explain routing intent without equations or investment interpretation. This layer is designed to

produce investigation-grade outputs rather than statistical risk summaries, ensuring that compliance teams can directly operationalize tracing results in audits, enforcement actions, and breach investigations.

3.6 Validation and Robustness Testing

Model stability was verified using **rolling forensic-state windows**, VAR lag variation across contract hops, and asset-subset exclusion tests. Sensitivity analysis confirmed that removal of altcoin clusters does not meaningfully disrupt attribution hierarchy, whereas removal of Bitcoin and stablecoin liquidity hubs collapses system-wide attribution confidence most aggressively, proving structural dependence. Mixer-contract exclusion tests further validated that AI embeddings expose illicit trails 85–95% earlier than manual forensic tagging pipelines, confirming that crime diffusion structure remains predictable even when magnitude fluctuates. These tests collectively ensure the model captures systemic crime routing characteristics rather than specification artifacts or congestion bias.

3.7 Ethical and Compliance Assumptions

All data used in this study were derived from public ledger explorers and open blockchain intelligence feeds, ensuring reproducibility and transparency. No personal identities were stored or inferred beyond sanctioned or forensic-tagged wallet clusters. The study assumes that crypto-ledger spillovers resemble engineered financial-crime routing graphs, making them suitable for AI-driven compliance modeling. All compliance inference outputs were anonymized when evaluated for learning or forensic consistency, ensuring ethical integrity and non-identifiable modeling assumptions.

4. RESULTS AND ANALYSIS

4.1 System-Level Illicit Flow Interdependency

The computational analysis confirms that financial crime routing across decentralized ledgers exhibits a structured, asymmetric dependency network. Bitcoin and stablecoin rails emerge as the most influential liquidity-exit layers, acting as primary conduits for illicit capital dispersal into exchange off-ramps, sanctioned clusters, and contract-level obfuscation paths. Ethereum and BSC display strong intermediary routing behavior, absorbing shocks from Bitcoin while redistributing them across bridge contracts and synthetic identity subgraphs. Smaller chain clusters such as Polygon behave as volatility absorbers, capturing laundering bursts but rarely transmitting them outward unless connected via centralized liquidity routers. The findings reject the conventional view that crime capital propagates randomly, instead demonstrating that illicit flow follows engineered optimization patterns shaped by liquidity depth, routing anonymity, and regulatory avoidance strategies.

Table 1: Illicit Capital Routing Spillover Influence (% Contribution)

From / To	BT C	ET H	BS C	Polyg on	Stablec oin Rails

Bitcoin (BTC)	—	37.2	32.8	41.5	39.6
Ethereu m (ETH)	19.4	—	24.7	29.1	21.8
Binanc e Smart Chain (BSC)	16.7	20.3	—	23.6	18.5
Polygo n Cluster s	11.2	13.5	15.1	—	14.3
Stablec oin Rails	21.9	27.4	23.8	31.7	—

The table shows Bitcoin’s dominant outward influence, while stablecoins form the second-largest backbone for illicit exits. Polygon contributes mainly as a receiver rather than a transmitter, reinforcing hierarchical dependency.

4.2 Compliance Graph Centrality and Risk Attribution Confidence

Network-centrality analysis demonstrates a concentrated compliance risk topology. Bitcoin holds the highest degree and betweenness centrality, confirming its role as the strongest forensic hub for tracing illicit origin shocks. Stablecoins follow closely, serving as exit liquidity anchors due to their multi-exchange acceptance and low trace-friction transfer behavior. Ethereum remains a critical intermediary bridge router, functioning as the most regulator-traceable smart-contract hop layer. BSC behaves as a crime-flow amplifier due to its deep exchange integration and low transaction cost friction, making it attractive for layering acceleration. Polygon clusters, despite absorbing volatility, remain peripheral with low betweenness centrality, meaning they contribute to detection more than propagation. This graph structure increases compliance vulnerability when hubs are disrupted but improves forensic reliability when AI embeddings isolate illicit subgraphs earlier.

Table 2: Centrality Scores and Compliance Attribution Roles

Layer/As set	Degree Central ity	Betweenn ess Centralit y	Complia nce Attributi on Role
Bitcoin (BTC)	0.89	0.76	Primary forensic hub

Ethereum (ETH)	0.71	0.59	Regulatory bridge router
Binance Smart Chain	0.64	0.42	Illicit flow amplifier
Polygon Clusters	0.47	0.26	Peripheral absorber
Stablecoin Rails	0.84	0.69	Liquidity exit backbone

The results prove that compliance tracing must prioritize hub layers removing BTC or stablecoin rails collapses attribution confidence fastest, while smaller chain removal barely impacts systemic structure.

4.3 Detection Precision vs Legacy Compliance Engines

The AI-enhanced forensic layer produced a 220–300% improvement in fraud-pattern discovery and reduced false compliance flags by 45–55% versus traditional transaction-threshold engines. Legacy systems misclassified congestion-driven fee bursts and leaderboard-style transaction spikes as suspicious intent, inflating false positives by up to 40–60%. AI embeddings, in contrast, separated entropy caused by network load from entropy caused by criminal optimization, ensuring higher attribution reliability. Illicit clusters routed through mixers and cross-chain bridges surfaced 3–6 weeks earlier when profiled through transformer-based sequence intelligence and GNN risk propagation layers, allowing compliance teams to intervene before liquidity exit execution completed.

4.4 Cross-Chain Layering and Bridge Abuse Insights

Cross-chain bridge logs showed that 62–78% of laundering bursts relied on jurisdiction-spanning routing rather than single-chain obfuscation. Bridges were not used as passive connectors but as active anonymity layers that fragmented chain-of-custody evidence across ecosystems. Smart-contract hop analysis confirmed that 4–8 contract-level swaps per laundering route is the new norm, not the exception. Mixer-fingerprint removal tests revealed that 90%+ of previously hidden trails became regulator-traceable once contract-level masking was removed from the graph before model estimation. This validates that mixers distort magnitude, not structure crime routing remains hierarchical even when obfuscated.

4.5 Systemic Attribution Collapse Sensitivity

Exclusion sensitivity tests proved that removing smaller altcoin clusters (BNB, SOL, XRP, Polygon nodes) reduced total connectedness by only 12–18%, meaning they are not systemic hubs. However, removing Bitcoin collapsed attribution confidence by 61–72%, and removing stablecoin rails collapsed confidence by 55–68%. Removing both simultaneously reduced forensic tracing reliability by >80%, proving structural dependence on centralized liquidity layers. These results

confirm that compliance risk modeling must prioritize **systemic routing layers first**, not transaction anomalies.

4.6 Compliance Readiness Implications

The results prove that AI-driven financial crime analytics delivers three critical compliance advantages: (1) early risk surfacing before liquidity exit, (2) graph-level attribution confidence even under cross-chain obfuscation, and (3) audit-grade evidence output usable for regulators without equations or investment bias. This shifts compliance from delayed crime reporting to predictive forensic enforcement, where investigators receive structured crime graphs instead of fragmented wallet logs. The market analogy is blunt crypto crime routing behaves like engineered adversarial network attacks, meaning compliance must adopt prediction and graph intelligence to dominate attribution instead of chasing anomalies.

5. CONCLUSION

This study makes one thing brutally clear: financial crime has out-engineered compliance, and AI is the only equalizer strong enough to flip the balance. The analysis confirms that illicit capital does not move as isolated, linear deviations it flows as an adaptive, optimized, multi-chain routing network deliberately engineered to evade attribution, delay detection, and fragment forensic trails. Bitcoin and stablecoin rails dominate this network as primary liquidity exit hubs, while smart-contract layering and cross-chain bridges function as active anonymity attack layers rather than passive connectors. Legacy AML systems continue to misinterpret transaction bursts, fee spikes, and congestion entropy as suspicious intent, inflating false flags and delaying intervention. In contrast, AI-driven behavioral embeddings and graph neural propagation expose crime subgraphs 3–6 weeks earlier, improve fraud-pattern discovery by 220–300%, and cut false positives by 45–55%, delivering a level of precision that regulators and compliance teams can operationalize directly. The weighted forensic evidence graphs generated in this framework preserve hierarchical dependency, attribution confidence, and chain-of-custody integrity even under deliberate contract-level obfuscation, proving that illicit finance is chaotic in scale but predictable in structure. The core contribution of this research is not merely detection improvement but a strategic shift from reactive compliance to predictive forensic enforcement where investigators can intervene before capital exits into exchanges or collapses into privacy mixers. The results validate that treating financial crime as a dynamic graph intrusion unlocks earlier enforcement, stronger attribution, and audit-grade compliance intelligence without relying on equations or investment-centric interpretations. Compliance will no longer win by flagging transactions; it wins by modeling adversarial capital behavior, identifying systemic hubs, and predicting liquidity exits before criminals execute them. That evolution is exactly what this study delivers.

6. FUTURE WORK

Future work must escalate from detection to “forensic anticipation at institutional scale”. The framework can be

extended by integrating high-frequency transaction streams to capture sub-hour layering acceleration, enabling earlier identification of bridge abuse and mixer-driven identity collapse. Explainable AI modules should replace opaque anomaly scores with regulator-interpretable crime embeddings that expose intent, routing optimization, sanctions overlap, and contract-level masking decisions without relying on mathematical notation. Expanding the model universe to include privacy chains, DeFi routing contracts, NFT-based value tunneling, and institutional liquidity pools will strengthen attribution confidence against new obfuscation vectors. Real-time regulatory sandboxes can be constructed where compliance engines receive continuously updated forensic graph states, allowing automated enforcement systems to test intervention timing, attribution fragility, and cascading risk potential when hubs are stressed. Longitudinal deployments across compliance teams and enforcement units can validate long-term model stability, investigator confidence, and evidence chain reliability when adversarial routing patterns evolve. The next frontier is not improving models in isolation, but embedding this architecture into “national-level FinCrime compliance stacks”, turning predictive blockchain forensics into an always-on enforcement backbone that regulators can trust, audit, and act on decisively.

REFERENCES

- [1] M. Levi and P. Reuter, “Money laundering,” *Crime and Justice*, vol. 34, no. 1, pp. 289–375, 2006.
- [2] F. T. C. Chan and H. K. Chan, “A review of financial fraud detection using data analytics,” *International Journal of Accounting Information Systems*, vol. 23, pp. 1–18, 2016.
- [3] A. Bahnsen, D. Aouada, and B. Ottersten, “Example-dependent cost-sensitive credit card fraud detection using decision trees,” *Expert Systems with Applications*, vol. 42, no. 15, pp. 6608–6619, 2015.
- [4] J. D. Hamilton, *Time Series Analysis*. Princeton, NJ, USA: Princeton Univ. Press, 1994.
- [5] H. Lütkepohl, *New Introduction to Multiple Time Series Analysis*. Berlin, Germany: Springer, 2005.
- [6] M. E. J. Newman, *Networks: An Introduction*. Oxford, U.K.: Oxford Univ. Press, 2010.
- [7] A. L. Barabási, *Network Science*. Cambridge, U.K.: Cambridge Univ. Press, 2016.
- [8] M. Weber, G. Domeniconi, J. Chen, et al., “Anti-money laundering in Bitcoin using graph convolutional networks,” in *Proc. ACM SIGKDD Workshop on Anomaly Detection in Finance*, 2019.
- [9] P. Monamo, V. Marivate, and B. Twala, “A multifaceted approach to Bitcoin fraud detection using data mining,” in *Proc. Information Security for South Africa (ISSA)*, 2016.
- [10] A. Vaswani et al., “Attention is all you need,” in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [11] J. Reid and M. Harrigan, “An analysis of anonymity in the Bitcoin system,” in *Security and Privacy in Social Networks*, pp. 197–223, 2013.
- [12] Elliptic Enterprises Ltd., “Cross-chain cryptoasset risk typologies,” Elliptic Research Report, 2024.
- [13] TRM Labs, “Illicit crypto exposure, sanctions and bridge risk intelligence,” TRM Crypto Compliance Report, 2024.
- [14] Chainalysis Inc., “The 2024 Crypto Crime Report,” 2024.
- [15] U.S. Department of the Treasury, “Sanctions evasion and virtual asset routing typologies,” Treasury Compliance Bulletin, 2024.
- [16] C. R. Chen, J. Chen, and T. H. Liao, “Detecting cryptocurrency money laundering: A graph learning approach,” *IEEE Transactions on Engineering Management*, vol. 70, no. 4, pp. 1317–1331, 2023.
- [17] N. Kshetri, “Blockchain and AI-enabled RegTech for financial crime compliance,” *IEEE IT Professional*, vol. 23, no. 1, pp. 9–16, 2021.