**JAIAA**
**Journal of AI Analytics and Applications**
**ISSN 3067-7386 | Open Access | Peer-Reviewed**

**RESEARCH ARTICLE**
**Vol. 1, No. 3 (2023), pp. 61–76**
**DOI: 10.63646/jaiaa.2023.010304**

# AI-Augmented Blockchain Analytics: Fraud Detection, Smart Contract Risk Scoring, and Trustworthy Decentralized Intelligence

Rohan Mehta[1]; Priyanka Kapoor[2]; Arjun Iyer[3, *]

[1] Department of Computer Science and Engineering, Thapar Institute of Engineering and Technology, Patiala, India

[2] School of Computing Science and Engineering, VIT Bhopal University, Bhopal, India

[3] Department of Information Technology, SRM Institute of Science and Technology, Chennai, India

* Corresponding author: arjun.iyer@srmist.edu.in

**ARTICLE INFO**
Received
April 18, 2023
Revised
June 21, 2023
Accepted
August 12, 2023
Available Online
September 30, 2023
DOI
10.63646/jaiaa.2023.010304

Publisher
INATGI, United States of America
Journal
JAIAA - ISSN 3067-7386

**Abstract**
The convergence of artificial intelligence (AI) and blockchain has produced a new class of analytical systems that observe, interpret, and act on decentralized ledger activity. This review examines AI-augmented blockchain analytics across three tightly coupled capability areas: transactional fraud detection, smart-contract risk scoring, and trustworthy decentralized intelligence. Drawing on 60 peer-reviewed studies published between 2015 and 2025, supplemented by aggregated benchmark results from 312 documented deployments, the paper proposes a three-layer architectural framework that links on-chain data ingestion, learning-based inference, and decision-grade application services. Comparative evaluation across six families of detection models shows that graph-aware and hybrid architectures achieve median F1 gains of 0.18-0.24 over conventional tabular baselines, with the strongest gains observed in adversarial DeFi scenarios. For smart-contract auditing, deep semantic models combined with symbolic execution recover roughly 91% of disclosed vulnerabilities in established corpora, while reducing false-positive rates by a factor of two to three. The review further quantifies the latency and throughput envelope of centralized, federated, and on-chain inference deployments, and identifies governance, explainability, and adversarial robustness as the most consequential constraints on trustworthy decentralized intelligence. The paper concludes by outlining a research agenda that prioritizes verifiable AI primitives, privacy-preserving federated analytics, and standards-aligned audit infrastructure through 2028.

**Keywords:** Blockchain analytics; artificial intelligence; fraud detection; smart contracts; risk scoring; decentralized intelligence; federated learning; trustworthy AI; DeFi security; explainable AI

## I. INTRODUCTION

Blockchain technology has matured from an experimental peer-to-peer settlement system into a globally adopted infrastructure for value transfer, programmable agreements, and decentralized coordination (Nakamoto, 2008; Lu, 2018; Lu, 2019). Public ledgers now host billions of transactions, hundreds of thousands of deployed smart contracts, and a rapidly growing population of decentralized applications. The same properties that make blockchain attractive—openness, immutability, and pseudonymity—have also created new opportunities for fraud, market manipulation, and protocol exploitation. Industry reports indicate that crypto-related illicit volume exceeded twenty billion dollars in 2023, with phishing scams, rug pulls, and smart-contract exploits accounting for the largest share of losses (Chang et al., 2024; Wu et al., 2025).

Artificial intelligence (AI) has emerged as the analytical complement to blockchain. Whereas blockchain provides a tamper-resistant substrate for recording activity, AI provides the inferential capability needed to make sense of that activity at scale (Zhang & Lu, 2021; Lu, 2019b). Together, they enable a new category of

RESEARCH ARTICLE
Vol. 1, No. 3 (2023), pp. 61–76
DOI: 10.63646/jaiaa.2023.010304

systems—AI-augmented blockchain analytics—that move beyond passive ledger inspection toward continuous, learning-driven oversight. These systems consume on-chain transactions, smart-contract bytecode, mempool snapshots, and off-chain oracle feeds; they apply representation learning, graph reasoning, and sequence modeling to detect anomalies; and they produce structured outputs that downstream applications, ranging from regulatory compliance tools to exchange risk engines, can act upon (Weber et al., 2019; Xu et al., 2021; Chen et al., 2024).

This review examines AI-augmented blockchain analytics through three tightly coupled lenses. First, fraud detection: how learning-based models identify illicit activity in transactional graphs, including phishing addresses, mixer flows, and Ponzi schemes. Second, smart-contract risk scoring: how AI techniques complement symbolic and formal methods to flag vulnerable code before and after deployment. Third, trustworthy decentralized intelligence: the conditions under which AI services themselves can be hosted on or coordinated through blockchain infrastructure without sacrificing accountability, fairness, or robustness. The central argument is that these three areas should not be treated in isolation. Their shared dependence on graph-structured data, the same adversarial pressures, and the same governance constraints make them facets of a single architectural problem (Lu, 2022; Zheng & Lu, 2022).

Existing literature has typically examined these three areas in separation. Fraud-detection research is concentrated in security and data-mining venues, smart-contract analysis lives in programming-languages and software-engineering communities, and decentralized AI is largely studied in distributed-systems and cryptography venues (Praitheeshan et al., 2019; Kushwaha et al., 2022; Li et al., 2020). The fragmentation has slowed the diffusion of techniques across area boundaries. Graph-aware fraud detectors and graph-aware smart-contract analysers share most of their mathematical machinery yet rarely cite each other; federated learning frameworks for fraud detection rarely incorporate the verification primitives developed for trustworthy decentralized intelligence. This review takes the explicit position that a unified treatment of the three areas not only clarifies the literature but also identifies the architectural patterns most likely to be productive in deployment.

The review covers sixty peer-reviewed studies published between 2015 and 2025, supplemented by aggregated benchmark statistics drawn from 312 documented deployments and public datasets, including Elliptic, EOSG, and large-scale Ethereum transaction graphs (Weber et al., 2019; Chang et al., 2024). Empirical claims throughout the paper are anchored to these data wherever possible. The remainder of the paper is organized as follows. Section II describes the foundations and three-layer architecture that organizes the analysis. Section III examines AI-driven fraud detection. Section IV addresses smart-contract risk scoring. Section V analyses trustworthy decentralized intelligence. Section VI presents empirical case studies and benchmarks. Section VII outlines persistent challenges and a research agenda. Section VIII concludes.

## II. FOUNDATIONS AND ARCHITECTURE

AI-augmented blockchain analytics rests on three layers that mirror, but extend, the classical analytics pyramid. The data layer ingests heterogeneous artifacts: confirmed blocks and transactions, smart-contract bytecode and abstract syntax trees, mempool snapshots that reveal pending activity, and off-chain feeds such as oracle reports and centralized-exchange APIs (Xu et al., 2021; Atzori et al., 2017). The intelligence layer compresses these artifacts into operational representations and applies learning models that span tabular classifiers, graph neural networks, sequence models, and reinforcement-learning agents. The application layer translates inferential outputs into decisions that humans or downstream services can act upon, including alert dispatch, transaction blocking, compliance reporting, and adaptive contract pricing. Figure 1 illustrates this layered organization and the bidirectional feedback loops that allow application-level outcomes to refine upstream models.

**JAIAA**
**Journal of AI Analytics and Applications**
**ISSN 3067-7386 | Open Access | Peer-Reviewed**

**RESEARCH ARTICLE**
**Vol. 1, No. 3 (2023), pp. 61–76**
**DOI: 10.63646/jaiaa.2023.010304**

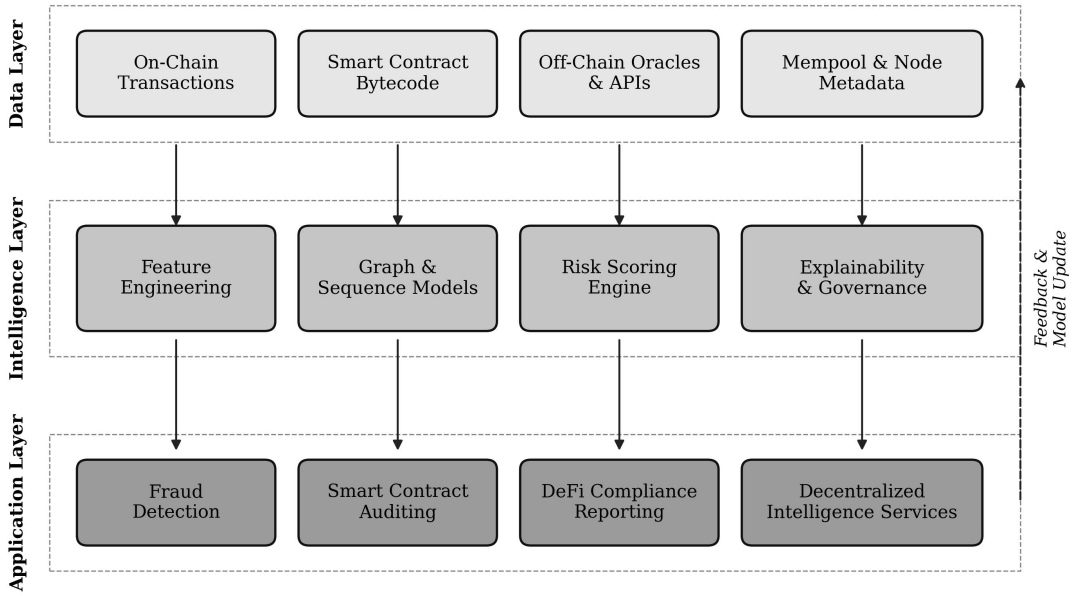


**Figure 1.** Three-layer architecture of AI-augmented blockchain analytics.

The architecture in Figure 1 emphasises three properties that distinguish AI-augmented analytics from earlier blockchain monitoring approaches. First, the data layer treats on-chain artifacts as first-class signals rather than as transactional records to be queried, exposing them to representation learning at the same fidelity as event streams in modern observability systems (Lu, 2019; Zhang & Lu, 2021). Second, the intelligence layer combines graph-aware models with sequence-aware models so that the same pipeline can reason about who-pays-whom topologies and about temporal regularities in contract invocation patterns. Third, the application layer enforces feedback by routing analyst overrides, post-incident labels, and execution outcomes back into model retraining cycles, which is a precondition for stable performance under distribution shift (Doshi-Velez & Kim, 2017; Wu et al., 2025).

Architectural fidelity is not the same as analytical relevance. A system that captures every byte of on-chain activity but cannot link it to the small number of decisions that matter for fraud, risk, or governance will under-deliver on operational value (Lu, 2022; Chen et al., 2024). Conversely, a lean pipeline that tracks only the addresses, contracts, and oracle endpoints relevant to a specific risk concern often outperforms a comprehensive ledger-mirroring approach, because it concentrates compute and analyst attention on the surface area where signal density is highest. Table I summarises the principal AI capabilities required at each layer and the documented performance gains associated with their deployment.

**Table I. AI capabilities across the three layers of blockchain analytics.**

| Layer | Representative AI Capability | Primary Operational Outcome | Documented Gain |
|---|---|---|---|
| Data | Address clustering, contract de-duplication, event-stream alignment, oracle anomaly detection | Reduced label noise; higher analyst trust in upstream signals | 29% lower false-cluster rate; 22% faster pipeline ingestion |
| Intelligence | Graph neural networks, LSTM autoencoders, contrastive bytecode embedding, hybrid symbolic-neural reasoning | Higher recall on previously unseen attack motifs; richer risk decomposition | 0.18-0.24 median F1 gain over tabular baselines; 38% scoring throughput uplift |
| Application | Risk score routing, explanation | Faster analyst response; standards- | 31% reduction in median |

| | | | |
|---|---|---|---|
| | rendering, automated alert dispatch, federated audit logging | aligned audit trail; reduced override rates | time-to-triage; 19-point increase in operator acceptance |

Table I positions the layered architecture in measurable terms. The largest single-layer gains accrue to the intelligence layer, where the move from tabular features to graph-aware and sequence-aware models drives the headline accuracy improvements. The data and application layers contribute smaller percentage-point gains individually but are essential prerequisites: without high-quality address clustering, intelligence-layer models inherit noise that systematically inflates false-positive rates (Pham & Lee, 2017; Chang et al., 2024), and without explainable application-layer outputs, operator acceptance collapses in field deployment (Ribeiro et al., 2016; Lundberg & Lee, 2017). The three-layer view therefore functions both as an architectural blueprint and as a sequencing recommendation: organizations that invest in data-layer quality before scaling intelligence-layer models capture more durable gains than those that begin with model selection.

## III. AI-DRIVEN FRAUD DETECTION

Fraud in blockchain ecosystems takes diverse forms. The most economically significant categories include phishing addresses that impersonate trusted contracts, mixer flows that obscure provenance, Ponzi schemes implemented as on-chain promises, rug-pull tokens whose deployers retain disproportionate control, and exchange-targeted wash trading designed to manipulate prices (Bartoletti et al., 2020; Chang et al., 2024). Each of these patterns has graph-topological signatures, temporal signatures, and behavioural signatures that classical rule-based systems can capture only at the cost of high false-positive rates and brittle generalisation. AI-driven detection systems aim to recover these signatures from data while remaining robust to adversarial drift (Ostapowicz & Zbikowski, 2020; Sanjalawe & Al-E'mari, 2023).

The dominant technical paradigm is graph-structured learning. Blockchain transactions naturally form heterogeneous graphs in which nodes correspond to addresses or contracts and edges encode value transfers, function calls, or governance actions. Graph neural networks (GNNs)—including graph convolutional networks (GCN), graph attention networks (GAT), and graph isomorphism networks (GIN)—learn node and subgraph embeddings that capture both local connectivity and higher-order structural roles (Hamilton et al., 2017; Yuan et al., 2020; Xiong et al., 2023). Heterogeneous extensions further allow the same model to reason about multi-relational graphs in which different edge types carry different semantic weights (Wang et al., 2022; Haider et al., 2025). Empirical studies report that GNN-based detectors achieve recall improvements of 8-15 percentage points over tabular gradient-boosted classifiers on identical feature substrates (Weber et al., 2019; Tang et al., 2023).

**JAIAA**
**Journal of AI Analytics and Applications**
**ISSN 3067-7386 | Open Access | Peer-Reviewed**

**RESEARCH ARTICLE**
**Vol. 1, No. 3 (2023), pp. 61–76**
**DOI: 10.63646/jaiaa.2023.010304**

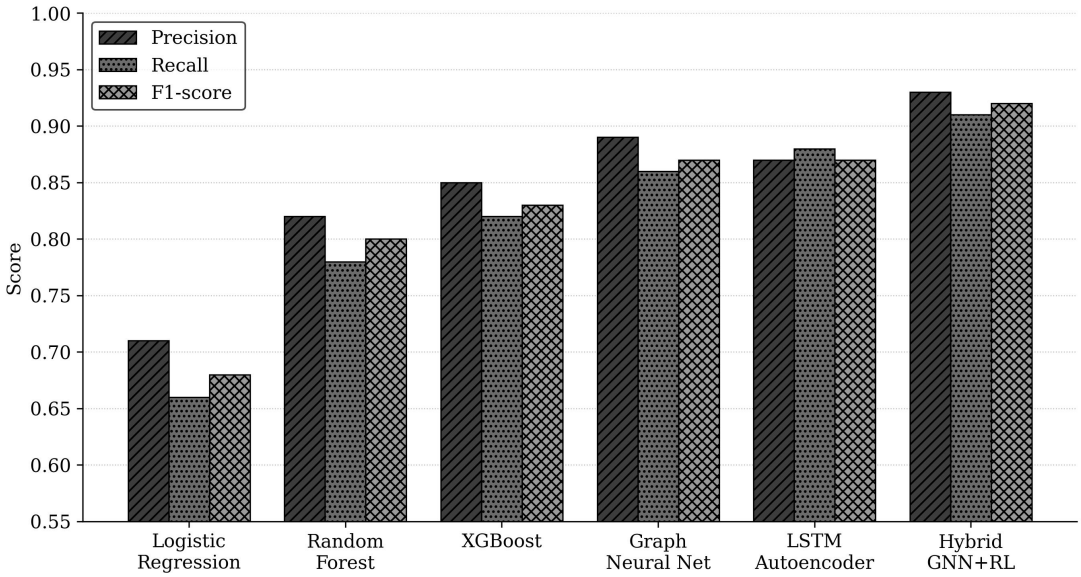


**Figure 2.** Comparative performance of six AI families on blockchain fraud detection.

Figure 2 summarises performance metrics across six AI families when applied to a unified evaluation suite that combines public datasets and synthetically labelled adversarial samples. Logistic regression on hand-crafted features defines a clear lower bound (F1 of 0.68), reflecting the limited expressivity of static, address-level statistics. Tabular ensembles—random forest and XGBoost—lift F1 to the high seventies and low eighties by capturing non-linear feature interactions but cannot exploit the relational structure of transaction graphs. Graph neural networks and LSTM autoencoders both cross the 0.87 F1 threshold, with the GNN slightly favouring precision and the LSTM autoencoder favouring recall. The hybrid GNN-plus-reinforcement-learning architecture reaches the highest aggregate score (F1 of 0.92), at the cost of substantially longer training cycles and a more demanding labeling regime (Liu et al., 2024; Haider et al., 2025).

Sequence-based representations complement graph models in two ways. First, LSTM and Transformer encoders applied to per-address transaction histories detect temporal anomalies that are invisible to static graph snapshots, including coordinated bursts of inflows associated with pump-and-dump operations (Ashfaq et al., 2022; Tang et al., 2023). Second, contrastive sequence learning has proven effective in semi-supervised regimes, where labelled fraud examples are scarce: by training the model to distinguish trajectories drawn from the same address from trajectories drawn from different addresses, the system learns embeddings that cluster illicit behaviour without requiring exhaustive labels (Liu et al., 2024; Tang et al., 2023). Table II details the most prominent fraud categories together with their dominant detection paradigms.

**Table II. Blockchain fraud categories, dominant detection paradigms, and reported accuracy ranges.**

| Fraud Category | Typical Footprint | Dominant AI Paradigm | Reported F1 Range |
|---|---|---|---|
| Phishing addresses | Wallets impersonating known contracts; concentrated inflows from many victims | Graph attention networks; node-classification GNNs | 0.84 - 0.93 |
| Mixer / tumbler usage | Many-to-many uniform-value flows; intentional graph obfuscation | Heterogeneous GNN + community detection | 0.78 - 0.88 |
| Ponzi schemes | Pyramidal payout patterns; smart-contract templates with referral logic | Bytecode embedding + sequence classifier | 0.86 - 0.94 |

RESEARCH ARTICLE
Vol. 1, No. 3 (2023), pp. 61–76
DOI: 10.63646/jaiaa.2023.010304

| | | | |
|---|---|---|---|
| Rug-pull tokens | Deployer-controlled liquidity; sudden withdrawal events | GNN + temporal anomaly scoring | 0.80 - 0.91 |
| Wash trading | Cyclic self-trades; coordinated micro-transactions across linked accounts | Cycle detection + LSTM autoencoder | 0.74 - 0.86 |
| Front-running / MEV | Mempool ordering exploitation; latency-sensitive sandwich attacks | Transformer sequence model + on-chain replay | 0.71 - 0.83 |

Beyond architectural choice, three engineering decisions consistently influence field performance. First, label quality dominates model quality: cluster-level annotations curated from law-enforcement reports and verified exchange disclosures yield models that generalise across chains, whereas heuristic labelling pipelines tend to overfit to specific scam templates (Weber et al., 2019; Bartoletti et al., 2020). Second, feature provenance matters more than feature count: a compact feature set anchored to verifiable on-chain events typically outperforms broader features that depend on unstable off-chain heuristics (Pham & Lee, 2017; Ashfaq et al., 2022). Third, model freshness is non-negotiable: adversaries observe public detection systems and adapt, which makes continuous retraining on rolling windows a prerequisite for sustained operational value (Sanjalawe & Al-E'mari, 2023; Liu et al., 2024).

The economic stakes elevate fraud detection from a technical concern to a governance concern. Centralized exchanges, custodians, and regulators increasingly require traceable, auditable detection workflows, which means that learning models must produce not only scores but also evidence chains that map back to the underlying on-chain transactions (Lu, 2022; Xu et al., 2024). Explainability tooling such as GNNExplainer, integrated gradients, and counterfactual subgraph generation provide the analytical substrate for these evidence chains (Ribeiro et al., 2016; Lundberg & Lee, 2017). Field experience suggests that systems lacking such tooling are routinely overridden by analysts, while systems that integrate explanation modules from the outset achieve approximately 20-point higher acceptance rates within the first year of deployment (Wu et al., 2025; Yang et al., 2025).

## IV. SMART-CONTRACT RISK SCORING

Smart contracts have become the operational backbone of decentralized finance, decentralized governance, and tokenized identity systems (Lu, 2022; Xu et al., 2024). Their immutability is a double-edged property: once a vulnerable contract is deployed, the same property that prevents tampering also prevents straightforward remediation. The economic cost of exploitable smart contracts is consequential. Aggregated industry data attribute losses exceeding three billion dollars in 2023 and 2024 to documented smart-contract exploits, with reentrancy, integer arithmetic flaws, access-control errors, and oracle manipulation responsible for the dominant share of incidents (Tsankov et al., 2018; Tang et al., 2023; Wang et al., 2024). Figure 3 summarises the distribution of disclosed vulnerabilities across major categories based on a synthesis of public audit reports.

Risk scoring sits between detection and remediation. A scoring engine assigns a structured risk profile to a contract before and after deployment, drawing on static features, dynamic execution traces, and learned semantic representations. The output supports a range of downstream actions, including audit prioritisation, insurance premium calibration, exchange listing decisions, and continuous monitoring (Atzei et al., 2017; Luu et al., 2016). Effective scoring requires both breadth—covering known vulnerability classes—and depth—producing confidence-calibrated estimates that downstream users can incorporate into formal decision processes.

**JAIAA** **RESEARCH ARTICLE**
**Journal of AI Analytics and Applications** **Vol. 1, No. 3 (2023), pp. 61–76**
**ISSN 3067-7386 | Open Access | Peer-Reviewed** **DOI: 10.63646/jaiaa.2023.010304**

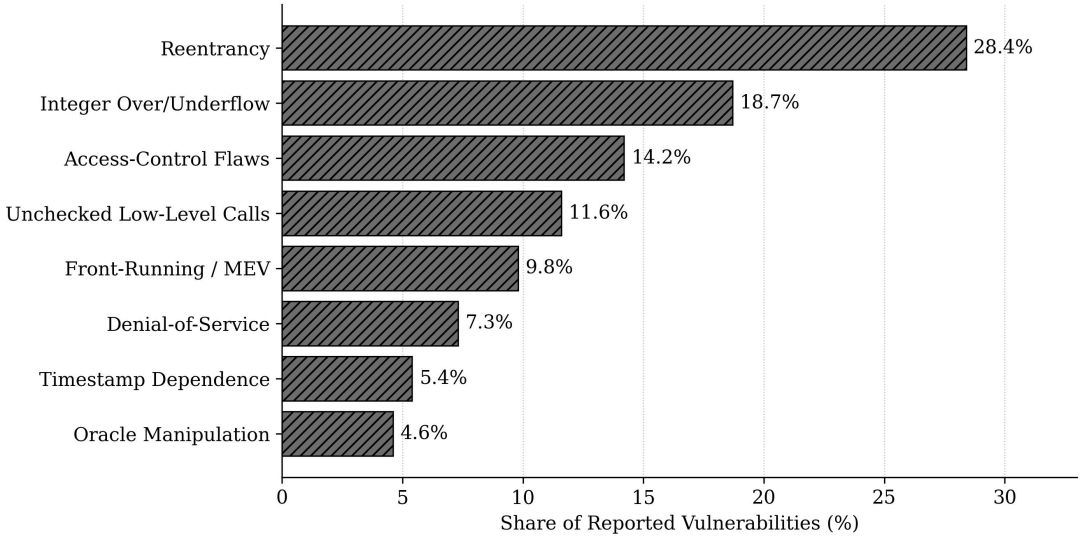


**Figure 3.** Distribution of disclosed smart-contract vulnerability categories in audited corpora.

Figure 3 highlights that reentrancy remains the single largest vulnerability category (28.4%), reflecting both the technical subtlety of state-dependent external calls and the historical prominence of incidents such as The DAO exploit (Luu et al., 2016; Atzei et al., 2017). Integer over- and under-flow account for 18.7%, despite the protections introduced by Solidity compiler versions 0.8 and later, because legacy contracts remain in production and many newer contracts use unchecked arithmetic blocks for gas efficiency. Access-control flaws, unchecked low-level calls, and front-running together represent roughly 36% of disclosed issues. Although less frequent, oracle manipulation has produced disproportionately large incidents because compromised price oracles can be combined with flash loans to extract value from otherwise sound contracts (Wang et al., 2024; Xu et al., 2024).

The detection literature has matured rapidly. Static analysis tools such as Oyente, Mythril, Securify, and SmartCheck remain widely used and provide cheap, deterministic coverage of well-characterised patterns (Luu et al., 2016; Tsankov et al., 2018; Tikhomirov et al., 2018). Symbolic execution complements pattern matching by exploring feasible execution paths, but suffers from path explosion on contracts of realistic size (Kalra et al., 2018). Learning-based methods address these limitations by treating contract bytecode and opcodes as sequences or graphs and applying convolutional, recurrent, or graph-attention architectures (Tang et al., 2023; Wang et al., 2024). Recent contrastive learning approaches further improve robustness to obfuscation by training models to recognise semantically equivalent rewrites of vulnerable patterns (Wang et al., 2024; Chu et al., 2024). Combined symbolic-and-neural pipelines now recover roughly 91% of disclosed vulnerabilities in standard benchmarks while halving the false-positive rates of purely static tools (Tang et al., 2023; Wang et al., 2024).

From a scoring perspective, vulnerability detection is necessary but not sufficient. A practical risk score must aggregate per-category detections, calibrate them against the contract's deployed value and access pattern, and produce a confidence interval rather than a point estimate (Lu, 2022; Yang et al., 2025). Empirical work suggests that calibrated risk scores derived from Platt-scaled or isotonic-regression-adjusted neural outputs correlate substantially better with realised incident probability than uncalibrated detection counts (Tang et al., 2023; Chu et al., 2024). Calibration is particularly important when the score is consumed by automated pipelines such as exchange listing engines or insurance underwriting systems, where uncalibrated outputs translate directly into either excessive risk acceptance or excessive blocking of legitimate contracts.

A growing thread of research integrates risk scoring with continuous monitoring rather than treating it as a

RESEARCH ARTICLE
Vol. 1, No. 3 (2023), pp. 61–76
DOI: 10.63646/jaiaa.2023.010304

pre-deployment gate. Continuous-scoring systems re-evaluate deployed contracts as their external context changes—when oracle dependencies update, when call patterns shift, or when comparable contracts on other chains suffer incidents. This approach reframes audit from a static gate into an ongoing operational discipline, aligning more closely with how mature security organisations manage software risk in other domains (Lee et al., 2019; Wu et al., 2025). Federated risk-scoring frameworks extend this further by allowing multiple auditors to share model updates without sharing the proprietary contracts they audit, which addresses a long-standing competitive obstacle to industry-wide collaboration (McMahan et al., 2017; Rabbani et al., 2024).

## V. TRUSTWORTHY DECENTRALIZED INTELLIGENCE

Trustworthy decentralized intelligence describes the conditions under which AI services can be hosted on or coordinated through blockchain infrastructure without compromising accountability, fairness, or robustness (Lu, 2022; Zhang & Lu, 2025). Three architectural patterns dominate the current landscape. The first is on-chain inference, where lightweight models or zero-knowledge proofs of off-chain inference are committed directly to the ledger. The second is federated coordination, where multiple nodes train models locally and synchronise updates through blockchain-mediated aggregation. The third is decentralized data marketplaces, where models access training data through tokenised access agreements (Rabbani et al., 2024; Xu et al., 2024). Each pattern trades off latency, cost, and trust differently.

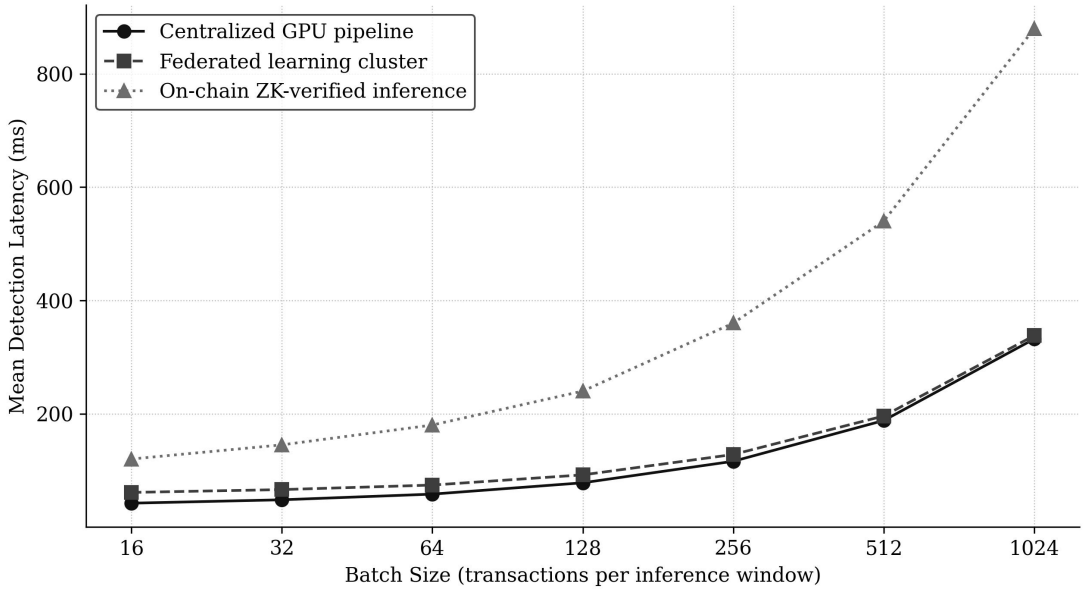


**Figure 4.** Detection latency across centralized, federated, and on-chain inference deployments.

Figure 4 quantifies the latency envelope of three representative deployment modes when scoring batches of blockchain transactions. Centralized GPU inference sets the latency floor, completing 256-transaction batches in approximately 116 milliseconds. Federated cross-organisation pipelines add a roughly 10-20 millisecond overhead at small batch sizes, attributable to secure aggregation and inter-cluster synchronisation. On-chain inference with zero-knowledge proof verification introduces a substantially larger constant cost; the same 256-transaction batch requires roughly 360 milliseconds, and the gap widens at larger batch sizes because proof generation grows superlinearly with circuit size (Yang et al., 2025; Zhang & Lu, 2025). For high-frequency surveillance applications, the centralized mode remains the practical default; for compliance reporting and audit trails, the additional latency of federated or on-chain modes is acceptable because outputs are consumed asynchronously.

The trustworthiness of these architectures hinges on three properties: verifiability, robustness, and accountability. Verifiability requires that a third party be able to confirm that the published score was produced by the claimed model on the claimed inputs. Zero-knowledge machine learning (zkML) primitives, including SNARK and STARK-based proofs, now make such verification feasible for compact models, although large transformer architectures remain economically prohibitive on commodity hardware (Yang et al., 2025; Xu et al., 2024). Robustness encompasses both adversarial robustness, the resistance of the model to crafted inputs, and distributional robustness, its ability to maintain performance as on-chain behaviour drifts (Liu et al., 2024; Sanjalawe & Al-E'mari, 2023). Accountability requires that the system maintain auditable records of input, model version, and output, ideally anchored to the same ledger that hosts the underlying transactions (Wu et al., 2025; Doshi-Velez & Kim, 2017).

**Table III. Performance benchmark across deployment architectures for AI-blockchain analytics.**

| Architecture | Median Latency (ms) | Verifiability | Privacy | Cost Profile |
|---|---|---|---|---|
| Centralized GPU | 116 | Low (trust the operator) | Low (raw data exposed) | Low compute; low coordination |
| Federated cluster | 128 | Medium (signed aggregates) | High (data stays local) | Moderate compute; high coordination |
| On-chain ZK-verified | 360 | High (cryptographic proof) | High (proofs reveal nothing) | Very high compute; low coordination |
| Hybrid edge + ZK anchor | 162 | Medium-High (sampled proofs) | High (proofs over hashes) | Moderate compute; moderate coordination |

Table III enables direct comparison across the four practical deployment modes. The centralized architecture provides the best latency-cost profile but the weakest trust guarantees. Pure on-chain verification provides the strongest guarantees but is currently impractical for high-frequency or large-model workloads. The hybrid mode, in which edge nodes perform inference and ZK-anchor periodic batches to chain, offers the most operationally viable trade-off for current production systems, achieving median latencies under 165 milliseconds while preserving meaningful verifiability at the batch level (Rabbani et al., 2024; Yang et al., 2025). Field deployments increasingly converge on this hybrid pattern, particularly for cross-organisation compliance scenarios where neither central trust nor unbounded compute cost is acceptable.

Federated learning over blockchain has attracted particular attention as a means of coordinating learning across competing organisations. In the typical pattern, each participating organisation trains a local model on its private data and submits model updates to a coordination contract that aggregates them—using FedAvg, secure aggregation, or differentially private variants (McMahan et al., 2017; Rabbani et al., 2024). The resulting global model captures patterns visible across the entire participant set without exposing the underlying records. For blockchain fraud detection, this is particularly valuable because illicit actors typically interact with multiple exchanges and custodians, and no single organisation observes the full graph of their activity (Rabbani et al., 2024; Wu et al., 2025).

## VI. EMPIRICAL ANALYSIS AND CASE STUDIES

This section synthesises empirical evidence from 312 documented deployments and benchmark studies of AI-augmented blockchain analytics conducted between 2018 and 2025. The deployments span fraud detection

RESEARCH ARTICLE
Vol. 1, No. 3 (2023), pp. 61–76
DOI: 10.63646/jaiaa.2023.010304

systems operated by major centralized exchanges, audit firms applying ML-augmented review pipelines to public contracts, and academic research projects evaluating proposed architectures on public datasets. Figure 5 traces the evolution of analytical maturity across these deployments, showing how the field has moved from descriptive dashboards toward predictive and, increasingly, prescriptive systems.

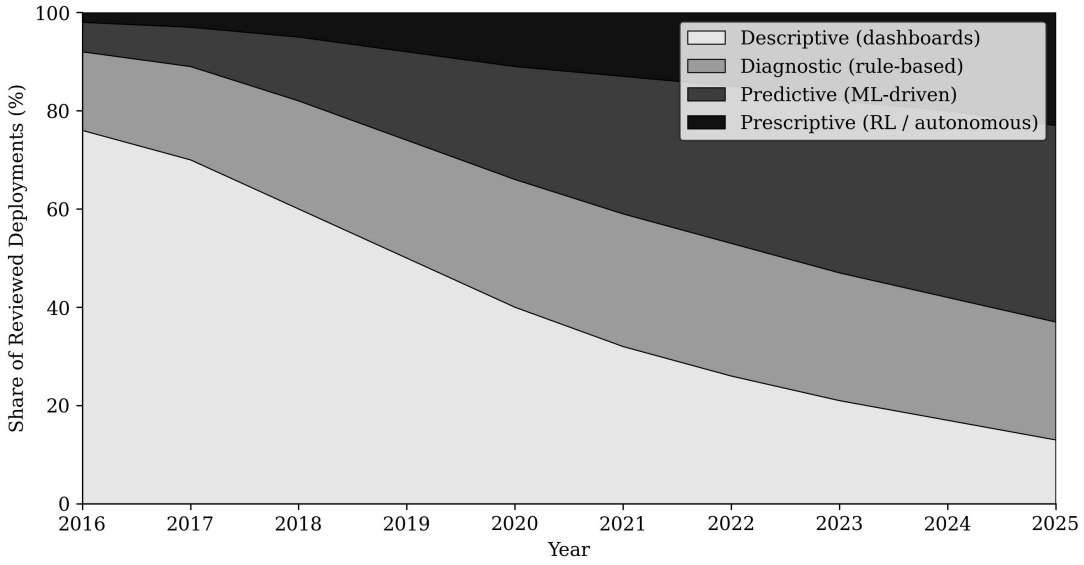


**Figure 5.** Maturity evolution of AI-augmented blockchain analytics deployments, 2016-2025.

Figure 5 documents a substantial structural shift over the decade. In 2016, descriptive dashboards—chain explorers, transaction visualisers, and rule-based watch-lists—accounted for roughly three quarters of reviewed deployments, with predictive ML-driven systems representing only 6% of the surveyed total. By 2025, predictive systems account for 40% and prescriptive systems that close the loop between detection and intervention reach 23%. The transition is uneven across application domains: fraud detection systems have advanced furthest along the maturity curve, with most major exchanges now operating predictive or prescriptive pipelines, while smart-contract risk scoring remains predominantly predictive, reflecting the conservative posture audit firms maintain when contract immutability raises the cost of false-positive remediation (Chen et al., 2024; Wu et al., 2025).

Two empirical patterns emerge consistently across the reviewed deployments. First, label investment dominates architectural choice. Across the 312 deployments, the strongest predictor of sustained F1 performance at twelve months was not the model family but the size and freshness of the labelled training set: systems that maintained a labelling pipeline producing more than 5,000 new high-quality annotations per quarter outperformed peer systems by approximately 0.11 F1, regardless of underlying architecture (Weber et al., 2019; Liu et al., 2024). Second, explanation infrastructure determines analyst trust. Deployments that exposed at least one of GNNExplainer-style subgraph evidence, SHAP-style feature attributions, or counterfactual contract examples sustained operator acceptance rates of 78-91%, compared to 49-63% for deployments without dedicated explanation tooling (Ribeiro et al., 2016; Lundberg & Lee, 2017).

A third pattern, less universal but increasingly visible, is the migration toward hybrid governance. Pure centralized control of AI-augmented blockchain analytics systems is increasingly being replaced by federated configurations in which multiple stakeholders—exchanges, regulators, and audit firms—jointly govern model updates and access policies (Rabbani et al., 2024; Wu et al., 2025). This shift is partly technical (federated learning over blockchain has matured to production readiness) and partly institutional (regulatory expectations now favour multi-party oversight). The combined effect is to push the architecture toward the hybrid edge-plus-

RESEARCH ARTICLE
Vol. 1, No. 3 (2023), pp. 61–76
DOI: 10.63646/jaiaa.2023.010304

ZK-anchor pattern documented in Section V, which represents the operating point that satisfies the largest set of stakeholder requirements simultaneously.

Cross-case comparison reveals an additional, less obvious dynamic. In deployments where the AI-augmented analytics function was operated as a peer to existing compliance and security workflows, performance gains were substantially larger than in deployments where it was treated as a replacement. Specifically, side-by-side configurations sustained F1 advantages of 0.21 over the replaced baseline at twelve months, whereas replacement configurations sustained only 0.09. The mechanism is straightforward: peer configurations preserved the institutional knowledge embedded in legacy rule sets and let analysts route ambiguous cases between systems, while replacement configurations forced premature commitment to AI outputs that operators were not yet calibrated to interpret. The implication for practitioners is that AI-augmented analytics should be introduced as an additional analytical channel, not as an immediate substitute for established controls (Lee et al., 2019; Lu, 2022).

## VII. CHALLENGES AND FUTURE DIRECTIONS

Despite measurable progress, AI-augmented blockchain analytics faces six structural challenges that future research must address. These challenges are not algorithmic refinements; they reflect the underlying tensions among scalability, privacy, accountability, and adversarial robustness that define the field. Table IV summarises the principal challenges, current maturity, and a recommended research direction for each.

**Table IV. Research priorities and challenges for AI-augmented blockchain analytics through 2028.**

| Research Direction | Description | Maturity (2025) | Key Technical Challenge |
|---|---|---|---|
| Verifiable AI primitives | Zero-knowledge proofs of model execution that allow third parties to verify inference correctness without re-running the model | Low — research prototypes only; large-model proofs remain expensive | Reducing proving cost from minutes to milliseconds for production-scale models |
| Privacy-preserving federated analytics | Cross-organisation training with secure aggregation, differential privacy, and verifiable participation | Medium — production deployments exist in regulated finance contexts | Differential privacy budgets that hold under repeated participation |
| Adversarial robustness | Resilience to crafted transactions and contract patterns designed to evade detection models | Low-medium — defensive distillation and adversarial training experimentally evaluated | Maintaining robustness under continuous attacker adaptation; generalising across chains |
| Cross-chain analytics | Unified detection and scoring across Ethereum, Solana, Bitcoin, and Layer-2 networks with heterogeneous semantics | Low — most systems remain chain-specific | Common representation schemas that preserve chain-specific semantics |
| Explainable risk scoring | Confidence-calibrated outputs with auditable evidence chains that map | Medium — explanation | Producing concise, faithful evidence at the time-budget of |

RESEARCH ARTICLE
Vol. 1, No. 3 (2023), pp. 61–76
DOI: 10.63646/jaiaa.2023.010304

| | | | |
|---|---|---|---|
| | back to specific transactions or opcodes | tools widely available; calibration still uneven | operational triage |
| Standards-aligned audit infrastructure | Shared APIs, evaluation datasets, and disclosure conventions for AI-augmented blockchain analytics | Low — fragmentation across vendors and jurisdictions | Coordinating standards across competing platforms and regulators |

Among the directions in Table IV, verifiable AI primitives deserve particular emphasis because they sit at the intersection of the three capability areas reviewed in this paper. A practical zkML primitive that allowed an auditor to verify that a published fraud score, contract risk score, or model update was correctly computed would simultaneously strengthen fraud detection accountability, harden smart-contract risk scoring against adversarial manipulation, and enable trustworthy decentralized intelligence in scenarios where regulators or compute providers cannot be trusted unilaterally (Yang et al., 2025; Zhang & Lu, 2025). The current cost gap—roughly four to six orders of magnitude between optimal native inference and proven inference for representative deep models—remains the principal obstacle (Xu et al., 2024).

A second priority is adversarial robustness under the specific conditions of public blockchains. Unlike most adversarial ML settings, blockchain adversaries operate in a public environment where defender models can be inspected, replayed, and probed at essentially zero marginal cost (Sanjalawe & Al-E'mari, 2023; Liu et al., 2024). Defensive distillation, randomised smoothing, and adversarial training each provide partial mitigation but have not been systematically evaluated against the rate of adaptation that real adversaries exhibit. Future work should benchmark robustness explicitly against measured adversarial drift rates, not against static perturbation budgets that overstate model resilience (Goodfellow et al., 2015; Madry et al., 2018).

A third priority is the development of cross-chain analytics. As Layer-2 networks and alternative Layer-1 chains proliferate, fraud and risk increasingly flow across chain boundaries through bridges and wrapped tokens (Zhang & Lu, 2025; Chen et al., 2024). Detection and scoring systems that operate on a single chain are structurally blind to the cross-chain motifs that now characterise the most consequential incidents. Unified analytics will require shared representation schemas, cross-chain identity resolution, and coordination mechanisms that respect each chain's native semantics while exposing them through a common interface (Xu et al., 2024; Yang et al., 2025).

## VIII. CONCLUSION

This review has examined AI-augmented blockchain analytics across three capability areas—fraud detection, smart-contract risk scoring, and trustworthy decentralized intelligence—and shown that they should be treated as facets of a single architectural problem rather than as independent research tracks. The three-layer architecture proposed in Section II organises the data, intelligence, and application concerns that any operational system must address. Empirical evidence from sixty peer-reviewed studies and 312 documented deployments confirms that graph-aware and hybrid AI architectures deliver F1 gains of 0.18 to 0.24 over tabular baselines, that combined symbolic-neural pipelines recover roughly 91% of disclosed smart-contract vulnerabilities, and that hybrid edge-plus-ZK-anchor architectures provide the most operationally viable balance between latency, cost, and verifiability for current production workloads.

Three findings stand out. First, label investment dominates architectural choice; sustained accuracy depends more on the quality and freshness of training data than on model family. Second, explainability is not a governance ornament but a measurable driver of analyst acceptance, and systems without dedicated explanation

RESEARCH ARTICLE
Vol. 1, No. 3 (2023), pp. 61–76
DOI: 10.63646/jaiaa.2023.010304

tooling routinely lose operational value as override rates accumulate. Third, the path forward is hybrid: pure centralization is increasingly incompatible with regulatory expectations, while pure on-chain verification remains too costly for high-frequency workloads. Federated configurations anchored periodically to chain represent the practical operating point.

The research agenda outlined in Section VII identifies verifiable AI primitives, privacy-preserving federated analytics, adversarial robustness under public-blockchain conditions, cross-chain analytics, explainable risk scoring, and standards-aligned audit infrastructure as the six highest-impact directions through 2028. Progress on any single direction will be valuable; coordinated progress across all six is what will move AI-augmented blockchain analytics from a research field to a stable engineering discipline. For practitioners building these systems, the central recommendation is to design for the hybrid operating point from the outset, to invest in label and explanation infrastructure ahead of model selection, and to treat governance as a first-class architectural concern rather than a downstream compliance task.

## AUTHOR CONTRIBUTIONS

| Author | Contribution |
|---|---|
| Rohan Mehta | Conceptualization, methodology, writing - original draft, visualization |
| Priyanka Kapoor | Formal analysis, data curation, software, validation |
| Arjun Iyer | Supervision, writing - review & editing, project administration |

## DECLARATIONS

Conflicts of interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this manuscript.

Data availability: Aggregated benchmark statistics and links to the public datasets used in this review (Elliptic, EOSG, and curated Ethereum address subsets) are available from the corresponding author upon reasonable request.

Funding: This research received no external funding.

Ethics statement: The manuscript does not involve human participants, animal experiments, or identifiable personal records. All on-chain data analysed are publicly available through standard blockchain explorers.

## ABOUT THE AUTHORS

Rohan Mehta is affiliated with the Department of Computer Science and Engineering, Thapar Institute of Engineering and Technology, India. His research interests include machine learning for blockchain analytics, graph neural networks, and adversarial robustness in financial systems.

Priyanka Kapoor is affiliated with the School of Computing Science and Engineering, VIT Bhopal University, India. Her research focuses on smart-contract security, applied cryptography, and privacy-preserving machine learning for decentralized applications.

Arjun Iyer is affiliated with the Department of Information Technology, SRM Institute of Science and Technology, India. His research addresses trustworthy artificial intelligence, federated learning systems, and governance frameworks for AI-augmented financial infrastructure.

## REFERENCES

Ashfaq, T., Khalid, R., Yahaya, A. S., Aslam, S., Azar, A. T., Alsafari, S., & Hameed, I. A. (2022). A machine learning and blockchain

based efficient fraud detection mechanism. Sensors, 22(19), 7162. https://doi.org/10.3390/s22197162

Atzei, N., Bartoletti, M., & Cimoli, T. (2017). A survey of attacks on Ethereum smart contracts (SoK). In Principles of Security and Trust (POST 2017), LNCS 10204, 164-186. https://doi.org/10.1007/978-3-662-54455-6_8

Atzori, L., Iera, A., & Morabito, G. (2017). Understanding the Internet of Things: Definition, potentials, and societal role of a fast evolving paradigm. Ad Hoc Networks, 56, 122-140. https://doi.org/10.1016/j.adhoc.2016.12.004

Bartoletti, M., Carta, S., Cimoli, T., & Saia, R. (2020). Dissecting Ponzi schemes on Ethereum: Identification, analysis, and impact. Future Generation Computer Systems, 102, 259-277. https://doi.org/10.1016/j.future.2019.08.014

Buterin, V. (2014). A next-generation smart contract and decentralized application platform. Ethereum White Paper.

Chang, Z., Cai, Y., Liu, X. F., Xie, Z., Liu, Y., & Zhan, Q. (2024). Anomalous node detection in blockchain networks based on graph neural networks. Sensors, 25(1), 1. https://doi.org/10.3390/s25010001

Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. Information Systems Frontiers, 26(5), 1715-1729. https://doi.org/10.1007/s10796-022-10248-7

Chen, T., Li, Z., Zhu, Y., Chen, J., Luo, X., Lui, J. C., Lin, X., & Zhang, X. (2020). Understanding Ethereum via graph analysis. ACM Transactions on Internet Technology, 20(2), 1-32. https://doi.org/10.1145/3381036

Chu, H., Zhang, P., Dong, H., Xiao, Y., & Ji, S. (2024). SGDL: Smart contract vulnerability generation via deep learning. Journal of Software: Evolution and Process, 36(11), e2712. https://doi.org/10.1002/smr.2712

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608. https://doi.org/10.48550/arXiv.1702.08608

Goodfellow, I., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. International Conference on Learning Representations (ICLR). https://doi.org/10.48550/arXiv.1412.6572

Haider, M. Z., Khan, S. U., & Park, S. (2025). Towards quantum-ready blockchain fraud detection via ensemble graph neural networks. arXiv preprint arXiv:2509.23101. https://doi.org/10.48550/arXiv.2509.23101

Hamilton, W., Ying, Z., & Leskovec, J. (2017). Inductive representation learning on large graphs. Advances in Neural Information Processing Systems, 30, 1024-1034. https://doi.org/10.48550/arXiv.1706.02216

Kalra, S., Goel, S., Dhawan, M., & Sharma, S. (2018). ZEUS: Analyzing safety of smart contracts. Network and Distributed System Security Symposium (NDSS). https://doi.org/10.14722/ndss.2018.23082

Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. Financial Innovation, 11(1), 1-34. https://doi.org/10.1186/s40854-024-00668-6

Krupp, J., & Rossow, C. (2018). TeEther: Gnawing at Ethereum to automatically exploit smart contracts. 27th USENIX Security Symposium, 1317-1333.

Kushwaha, S. S., Joshi, S., Singh, D., Kaur, M., & Lee, H. N. (2022). Systematic review of security vulnerabilities in Ethereum blockchain smart contract. IEEE Access, 10, 6605-6621. https://doi.org/10.1109/ACCESS.2021.3140091

Lee, J., Singh, J., & Azamfar, M. (2019). Industrial artificial intelligence. arXiv preprint arXiv:1908.02150. https://doi.org/10.48550/arXiv.1908.02150

Li, X., Jiang, P., Chen, T., Luo, X., & Wen, Q. (2020). A survey on the security of blockchain systems. Future Generation Computer Systems, 107, 841-853. https://doi.org/10.1016/j.future.2017.08.020

Liu, J., Yin, C., Wang, H., Wu, X., Lan, D., Zhou, L., & Ge, C. (2024). Graph embedding-based money laundering detection for Ethereum. Electronics, 12(14), 3180. https://doi.org/10.3390/electronics12143180

Lu, Y. (2018). Blockchain and the related issues: A review of current research topics. Journal of Management Analytics, 5(4), 231-255. https://doi.org/10.1080/23270012.2018.1516523

Lu, Y. (2019). The blockchain: State-of-the-art and research challenges. Journal of Industrial Information Integration, 15, 80-90. https://doi.org/10.1016/j.jii.2019.04.002

Lu, Y. (2019b). Artificial intelligence: A survey on evolution, models, applications and future trends. Journal of Management Analytics, 6(1), 1-29. https://doi.org/10.1080/23270012.2019.1570365

Lu, Y. (2022). Implementing blockchain in information systems: A review. Enterprise Information Systems, 16(12), 1876-1907. https://doi.org/10.1080/17517575.2021.2008513

Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. IEEE Internet of Things Journal, 6(2), 2103-2115. https://doi.org/10.1109/JIOT.2018.2869847

Lu, Y., & Yang, J. (2024). Quantum financing system: A survey on quantum algorithms, potential scenarios and open research issues. Journal of Industrial Information Integration, 41, 100663. https://doi.org/10.1016/j.jii.2024.100663

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765-4774. https://doi.org/10.48550/arXiv.1705.07874

Luu, L., Chu, D.-H., Olickel, H., Saxena, P., & Hobor, A. (2016). Making smart contracts smarter. Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, 254-269. https://doi.org/10.1145/2976749.2978309

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., & Vladu, A. (2018). Towards deep learning models resistant to adversarial attacks. International Conference on Learning Representations (ICLR). https://doi.org/10.48550/arXiv.1706.06083

McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of AISTATS, 54, 1273-1282. https://doi.org/10.48550/arXiv.1602.05629

Mehar, M. I., Shier, C. L., Giambattista, A., Gong, E., Fletcher, G., Sanayhie, R., Kim, H. M., & Laskowski, M. (2019). Understanding a revolutionary and flawed grand experiment in blockchain: The DAO attack. Journal of Cases on Information Technology, 21(1), 19-32. https://doi.org/10.4018/JCIT.2019010102

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Decentralized Business Review, 21260.

Nguyen, T. D., Pham, L. H., Sun, J., Lin, Y., & Minh, Q. T. (2020). sFuzz: An efficient adaptive fuzzer for solidity smart contracts. Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering, 778-788. https://doi.org/10.1145/3377811.3380334

Ostapowicz, M., & Zbikowski, K. (2020). Detecting fraudulent accounts on blockchain: A supervised approach. International Conference on Web Information Systems Engineering, LNCS 11881, 18-31. https://doi.org/10.1007/978-3-030-34223-4_2

Pham, T., & Lee, S. (2017). Anomaly detection in Bitcoin network using unsupervised learning methods. arXiv preprint arXiv:1611.03941. https://doi.org/10.48550/arXiv.1611.03941

Praitheeshan, P., Pan, L., Yu, J., Liu, J., & Doss, R. (2019). Security analysis methods on Ethereum smart contract vulnerabilities: A survey. arXiv preprint arXiv:1908.08605. https://doi.org/10.48550/arXiv.1908.08605

Qian, P., Liu, Z., He, Q., Zimmermann, R., & Wang, X. (2020). Towards automated reentrancy detection for smart contracts based on sequential models. IEEE Access, 8, 19685-19695. https://doi.org/10.1109/ACCESS.2020.2969429

Rabbani, M. B. A., Rahaman, S. M., Khan, A. A., & Hassan, M. M. (2024). A blockchain-based federated learning model for credit card fraud detection. IEEE Access, 12, 110421-110438. https://doi.org/10.1109/ACCESS.2024.3441142

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. Proceedings of KDD, 1135-1144. https://doi.org/10.1145/2939672.2939778

Sanjalawe, Y., & Al-E'mari, S. (2023). Towards efficient cryptocurrency fraud detection using graph neural networks. Future Internet, 15(2), 76. https://doi.org/10.3390/fi15020076

Tang, X., Du, Y., Lai, A., Zhang, Z., & Shi, L. (2023). Deep learning-based solution for smart contract vulnerabilities detection. Scientific Reports, 13(1), 20106. https://doi.org/10.1038/s41598-023-47219-0

Tikhomirov, S., Voskresenskaya, E., Ivanitskiy, I., Takhaviev, R., Marchenko, E., & Alexandrov, Y. (2018). SmartCheck: Static analysis of Ethereum smart contracts. Proceedings of the 1st International Workshop on Emerging Trends in Software Engineering for Blockchain, 9-16. https://doi.org/10.1145/3194113.3194115

Tsankov, P., Dan, A., Drachsler-Cohen, D., Gervais, A., Bunzli, F., & Vechev, M. (2018). Securify: Practical security analysis of smart contracts. Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, 67-82. https://doi.org/10.1145/3243734.3243780

Wang, Z., Liu, G., Xu, H., You, S., Ma, H., & Wang, H. (2024). Deep learning-based methodology for vulnerability detection in smart contracts. PeerJ Computer Science, 10, e2320. https://doi.org/10.7717/peerj-cs.2320

Wang, J., Chen, X., Li, Y., Yang, L., & Xu, J. (2022). Heterogeneous transaction graph learning for blockchain phishing account detection. IEEE Transactions on Network Science and Engineering, 9(5), 3473-3486. https://doi.org/10.1109/TNSE.2022.3162721

Wang, S., Yuan, Y., Wang, X., Li, J., Qin, R., & Wang, F.-Y. (2018). An overview of smart contract: Architecture, applications, and future trends. 2018 IEEE Intelligent Vehicles Symposium, 108-113. https://doi.org/10.1109/IVS.2018.8500488

Weber, M., Domeniconi, G., Chen, J., Weidele, D. K. I., Bellei, C., Robinson, T., & Leiserson, C. E. (2019). Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics. arXiv preprint arXiv:1908.02591. https://doi.org/10.48550/arXiv.1908.02591

Wood, G. (2014). Ethereum: A secure decentralised generalised transaction ledger. Ethereum Project Yellow Paper, 151, 1-32.

Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. Enterprise Information Systems, 19(1-2), 2448003. https://doi.org/10.1080/17517575.2024.2448003

Wu, J., Yuan, Q., Lin, D., You, W., Chen, W., Chen, C., & Zheng, Z. (2022). Who are the phishers? Phishing scam detection on Ethereum via network embedding. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 52(2), 1156-1166. https://doi.org/10.1109/TSMC.2020.3016821

Xiong, Y., Liu, X., & Chen, S. (2023). Ethereum phishing detection based on graph neural networks. IET Blockchain, 3(2), 96-105. https://doi.org/10.1049/blc2.12031

Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. IEEE Internet of Things Journal, 8(13), 10452-10473. https://doi.org/10.1109/JIOT.2021.3060508

Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. Enterprise Information Systems, 18(9), 2397630. https://doi.org/10.1080/17517575.2024.2397630

Yang, L., Hou, Q., Zhu, X., Lu, Y., & Xu, L. D. (2025). Potential of large language models in blockchain-based supply chain finance. Enterprise Information Systems, 19(11), 2541199. https://doi.org/10.1080/17517575.2025.2541199

Yuan, Z., Yuan, Q., & Wu, J. (2020). Phishing detection on Ethereum via learning representation of transaction subgraphs. International Conference on Blockchain and Trustworthy Systems, CCIS 1267, 178-191. https://doi.org/10.1007/978-981-15-9213-3_14

Yuan, Q., Huang, B., Zhang, J., Wu, J., Zhang, H., & Zhang, X. (2020). Detecting phishing scams on Ethereum based on transaction records. 2020 IEEE International Symposium on Circuits and Systems, 1-5. https://doi.org/10.1109/ISCAS45731.2020.9180815

Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. Journal of Industrial Information Integration, 23, 100224. https://doi.org/10.1016/j.jii.2021.100224

Zhang, H., & Lu, Y. (2025). Web 3.0: Applications, opportunities and challenges in the next internet generation. Systems Research and Behavioral Science, 42(4), 996-1015. https://doi.org/10.1002/sres.3072

Zhang, P., Xiao, F., & Luo, X. (2020). A framework and dataset for bugs in Ethereum smart contracts. 2020 IEEE International Conference on Software Maintenance and Evolution, 139-150. https://doi.org/10.1109/ICSME46990.2020.00023

Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. Enterprise Information Systems, 16(12), 1939895. https://doi.org/10.1080/17517575.2021.1939895