ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

A Comprehensive Survey of Neuro-Symbolic AI for Blockchain Security: Taxonomy, Critical Analysis and Open Challenges

Mohammed Ali Qatran[1]; Malek Nasser Algabri[2]; Redhwan Nasser Algabri[3]

[1,2,3]Department of Computer Science, Faculty of Computer & IT, Sana'a University, Sana'a, Yemen

Publication Date: 2026/11/26

Abstract: Blockchain platforms have become general-purpose infrastructure for decentralized finance, supply chains, and digital identity, yet they remain exposed to sophisticated multi-layer attacks. Purely neural defenses achieve strong detection accuracy but lack interpretability and formal guarantees, whereas purely symbolic defenses offer rigor and auditability but scale poorly against dynamic, high-volume ledger data. Neuro-Symbolic AI (NSAI), which couples neural perception with symbolic reasoning, has emerged as a promising third paradigm, yet its application to blockchain security has not been surveyed in a systematic and analytically integrated manner. This article presents a comprehensive literature survey of NSAI for blockchain security. Drawing on more than 168 publications and analyzing 47 primary studies with identifiable neural and symbolic components alongside 31 foundational works, the survey proposes a two-dimensional taxonomy that classifies systems by integration depth and by security task. For each task, the survey contrasts neural, symbolic, and hybrid designs, synthesizes reported strengths and limitations, and assesses explainability, robustness, scalability, and verifiability. It further adapts the Grounding–Instructibility–Alignment (G-I-A) lens to blockchain settings and derives a research agenda covering standardized benchmarks, certified adversarial robustness, cross-chain neuro-symbolic monitoring, and formally verifiable inference. The synthesis indicates that the literature remains skewed toward smart-contract auditing and transaction-graph analysis, while consensus-layer protection and on-chain deployment of symbolic reasoning remain comparatively under-explored.

Keywords: *Neuro-Symbolic AI, Blockchain Security, Smart Contracts, Graph Neural Networks, Logic Tensor Networks, Knowledge Graphs, Intrusion Detection, Explainable AI, Formal Verification.*

How to Cite: Mohammed Ali Qatran; Malek Nasser Algabri; Redhwan Nasser Algabri (2026) A Comprehensive Survey of Neuro-Symbolic AI for Blockchain Security: Taxonomy, Critical Analysis and Open Challenges. *International Journal of Innovative Science and Research Technology*, 11(11), 1750-1761. https://doi.org/10.38124/ijisrt/26nov1021

I. INTRODUCTION

Blockchain platforms now underpin decentralized finance (DeFi), supply chains, and digital identity, yet they remain attractive targets for sophisticated, multi-layer attacks. Since Nakamoto's Bitcoin white paper [1] and Buterin's Ethereum proposal [2], blockchain has evolved from a cryptocurrency substrate into general-purpose infrastructure for DeFi, tokenized assets, cross-organizational data sharing, and verifiable digital identity. This expansion has widened the attack surface. At the consensus layer, protocols must tolerate Byzantine behavior under adversarial network conditions, with Practical Byzantine Fault Tolerance (PBFT) and its descendants providing foundational guarantees for permissioned deployments [3]. At the application and execution layers, smart contracts have proven fragile: the taxonomy of Atzei et al. [4] documented reentrancy, unchecked external calls, and access-control flaws years before they were exploited for losses measured in hundreds of millions of dollars. At the network layer, eclipse attacks [5] and BGP routing hijacks [6] demonstrated that even the topology of peer-to-peer overlays can be weaponized. Industry incident reporting indicates that blockchain-related theft remains measured in billions of dollars annually [7], underscoring the urgency of automated, provably sound defense mechanisms.

Defensive techniques have co-evolved along two largely separate tracks. The first is symbolic: static analyzers such as Slither [8] and Securify [9], fuzzers such as Echidna [10] and ConFuzzius [11], model checkers, and theorem provers offer rigorous, human-auditable guarantees, but they depend on manually curated rules and invariants and struggle with path explosion, semantic complexity, and the throughput of modern ledgers. The second is neural: graph neural networks (GNNs) for transaction and contract analysis [12], sequence models for opcode streams, and large language models (LLMs) that audit Solidity code. Neural models generalize to novel patterns and operate at data-center scale, but they behave as black boxes, offer no formal proof of correctness, and remain vulnerable to adversarial manipulation [30].

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

NSAI seeks to close this gap. Rather than treating learning and reasoning as mutually exclusive, NSAI systems bind them so that neural components provide perception over noisy, high-dimensional ledger data while symbolic components supply constraints, formal explanations, and verifiable inference [13]. In the broader cybersecurity domain, Hakim et al. surveyed more than one hundred publications and showed that neuro-symbolic systems are transitioning from laboratory prototypes to deployed intrusion detection, vulnerability analysis, and threat-intelligence pipelines [14]. Yet the intersection of NSAI with blockchain security has not received equivalent integrated treatment: blockchain data—immutable transaction graphs, deterministic contract execution, consensus traces, and wallet semantics—differs substantially from conventional enterprise telemetry.

Related surveys and positioning. Existing surveys on neuro-symbolic AI for cybersecurity [14] focus on enterprise telemetry and do not address blockchain-specific artifacts such as immutable multigraphs, EVM bytecode, or consensus traces. Surveys on blockchain security [15], [16] treat neural and symbolic defenses separately without analyzing their functional integration. Surveys on graph learning for blockchain data [12] focus on representation learning rather than on verification. Consequently, no prior survey offers a taxonomy that simultaneously characterizes integration depth and security task, nor one that adapts an evaluation lens such as G-I-A to blockchain settings. This survey is positioned to fill that gap through critical synthesis rather than bibliometric counting.

II. SCOPE, OBJECTIVES, AND CONTRIBUTIONS OF THE SURVEY

➢ *Thematic Scope*

The survey covers scholarly work that (i) addresses a security- or safety-relevant task for blockchain platforms or ledger-derived telemetry, and (ii) incorporates identifiable neural and symbolic components with a non-trivial functional interaction. Purely neural detectors, purely symbolic verifiers, and generic machine-learning surveys are retained strictly as foundational or comparative baselines. Consensus redesigns lacking learning or formal reasoning modules, pure cryptoeconomic mechanism design, and generic NSAI frameworks without a security application are excluded.

➢ *Temporal Coverage*

The survey focuses on 2019–2026, a window that captures the transition from standalone neural detectors and symbolic analyzers toward integrated pipelines. Foundational references published before 2019 are retained when they define the technical vocabulary of the field.

➢ *Sources and Terminology*

Sources include IEEE Xplore, ACM Digital Library, SpringerLink, ScienceDirect, Scopus, Web of Science, and arXiv, complemented by backward and forward citation tracing. Terminology combines neuro-symbolic vocabulary with blockchain security keywords. Exact per-database result counts and search dates are required for full reproducibility and are marked [Author Data Required] in this manuscript.

➢ *Organization of the Literature*

Studies are organized along two primary axes (integration depth and security task) and four secondary axes (neural component, symbolic component, knowledge representation, and deployment environment). This organization supports thematic comparison rather than chronological listing.

➢ *Coverage Limitations and Potential Biases*

The survey inherits several limitations. Publication bias favors positive results. Language bias restricts coverage to English-language venues. Preprint inclusion introduces version drift, particularly for 2025–2026 LLM-based studies. Many primary studies reuse small, overlapping datasets (SmartBugs, SolidiFI, Elliptic), which inflates apparent cross-study consistency. These limitations are discussed explicitly rather than hidden behind a formal protocol, in keeping with the analytical nature of the survey.

➢ *Objectives and Contributions*

This survey pursues five objectives: (1) to define the thematic boundaries of NSAI for blockchain security and organize the literature thematically and technically; (2) to propose a two-dimensional taxonomy crossing integration depth and security task; (3) to compare neural, symbolic, and hybrid designs across five security tasks, distinguishing reported evidence from inference; (4) to adapt Grounding–Instructibility–Alignment (G-I-A) to blockchain settings and assess how far the surveyed studies support it; and (5) to derive open challenges and emerging directions grounded in observed asymmetries.

III. CONCEPTUAL AND TECHNICAL FOUNDATIONS

➢ *Blockchain Security Layers*

A blockchain platform decomposes into four attack-relevant layers [15], [16]. Consensus protocols (PoW, PoS, PBFT and variants [3]) must preserve safety and liveness under partial synchrony and Byzantine faults; failure modes include 51% attacks, selfish mining, equivocation, and liveness denial. Networking layers face eclipse and partition attacks [5], BGP routing hijacks [6], and denial-of-service flooding. Execution layers host virtual machines—most prominently the EVM—where contracts execute deterministically and immutably, making pre-deployment verification essential [4]. Application layers compose contracts into DeFi protocols whose emergent interactions (flash loans, oracle dependencies, cross-chain bridges) create systemic attack surfaces [16]. Two properties distinguish blockchain data: the ledger forms a temporal, heterogeneous multigraph of addresses, contracts, and token flows [12], [17]; and contracts possess formal bytecode semantics amenable to static analysis [8], [9], fuzzing [10], [11], and theorem proving [18].

➢ *Neuro-Symbolic Integration Pillars*

Following d'Avila Garcez et al. [13] and Hakim et al. [14], four technical pillars characterize neuro-symbolic security architectures: (1) Differentiable logic, exemplified by LTNs [19] and LTNtorch [20], which compile first-order

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

invariants into fuzzy-logic satisfiability losses; (2) Knowledge-enhanced graph learning, exemplified by KnowGraph [21]; (3) LLM–symbolic coupling, where generative models synthesize candidate invariants that SMT solvers and static analyzers formally validate [22], [23], [24]; and (4) Causal neuro-symbolic reasoning over ledger execution traces [25], [26]. Fig. 1 illustrates how these pillars bridge raw ledger telemetry and formally verified security actions.

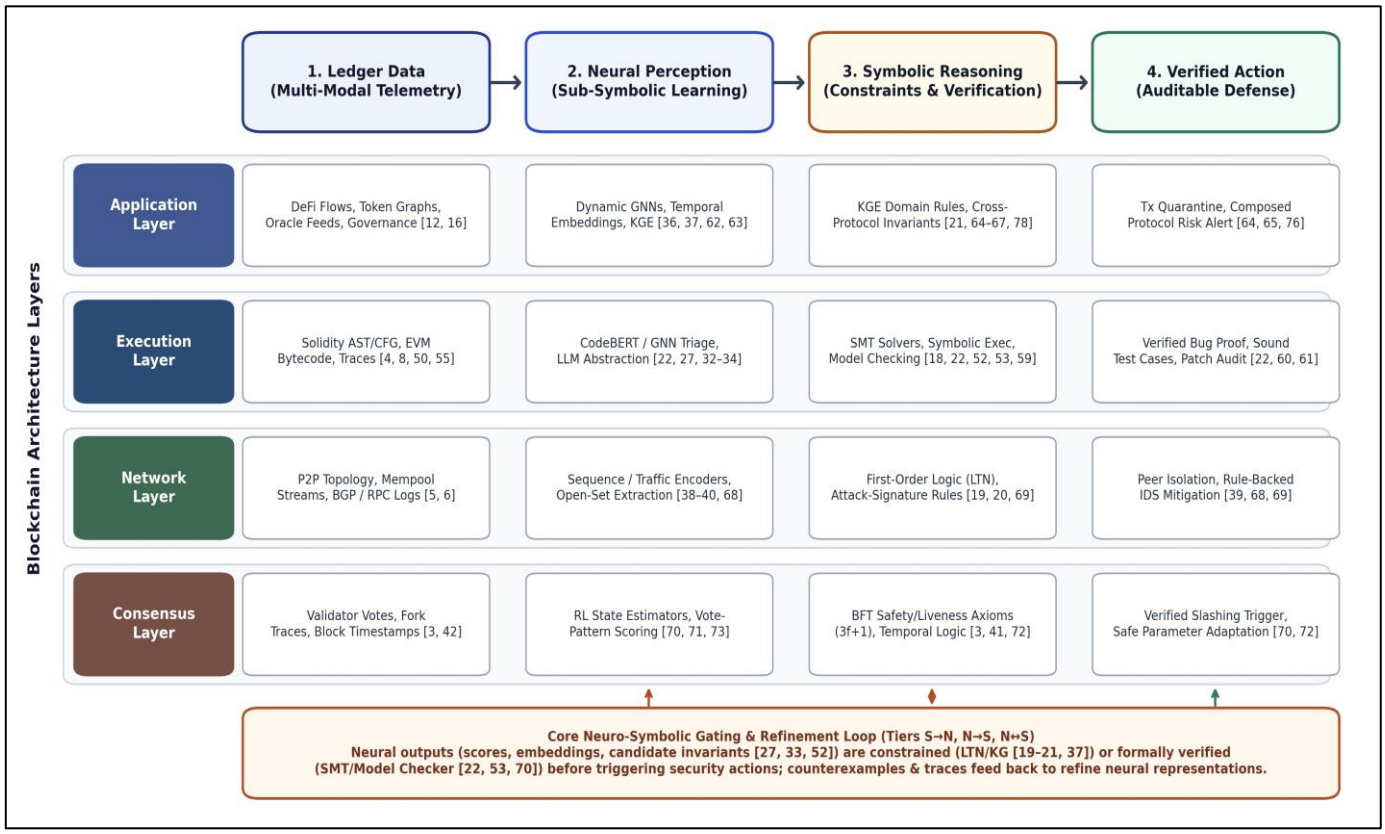


Fig 1 Conceptual Framework Linking Neuro-Symbolic AI to Blockchain Security Across Ledger Layers and Integration Loops.

IV. TAXONOMY OF EXISTING APPROACHES

➢ *Integration Depth*

Tier S→N (Symbolically Constrained Neural Learning): symbolic domain rules, ontologies, or knowledge graphs regularize or constrain neural representation learning during training or inference (e.g., LTN-constrained classifiers [19], [20] and knowledge-graph-augmented GNNs [21], [37]).

Tier N→S (Neural-Guided Symbolic Analysis): a neural perception module triages, prioritizes, or abstracts inputs to accelerate an authoritative symbolic engine such as a symbolic executor, model checker, or SMT solver [27], [33], [34], [50], [51]. Soundness rests with the symbolic verifier.

Tier N↔S (Tightly Coupled Hybrid Loops): neural and symbolic modules interact iteratively within closed verification-in-the-loop cycles: candidate invariants, patches, or inputs are concretely executed or model-checked, counterexamples are diagnosed, and neural abstractions are refined [22], [28], [29], [52], [53].

➢ *Security Task*

Five tasks are distinguished across the 47 primary studies: smart-contract vulnerability detection and repair (23 studies); fraud, money laundering, and DeFi anomaly detection (11 studies); network intrusion and node telemetry monitoring (7 studies); consensus-layer protection and validator auditing (3 studies); and privacy, identity, and access-control verification (3 studies). The distribution is uneven: smart-contract studies dominate, while consensus-layer protection remains the least populated quadrant.

➢ *Coding Scheme*

Table 1 Coding Scheme for the Surveyed Corpus.

Dimension	Values
Security task	SC, Fraud/DeFi, Network, Consensus, Privacy
Blockchain layer	Application, Execution, Network, Consensus
Neural component	GNN, sequence model, LLM, multimodal

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

Symbolic component	LTN, KG, SMT, static analyzer, temporal logic
Integration mechanism	S→N, N→S, N↔S
Knowledge representation	FOL, KG, CFG/DFG, SMT, ontology
Deployment environment	On-chain, off-chain, hybrid
Evaluation rigor	Dataset, baselines, ablation, adversarial testing

A complete table covering all 47 primary studies across the eight coding dimensions is provided as in fig. 2.

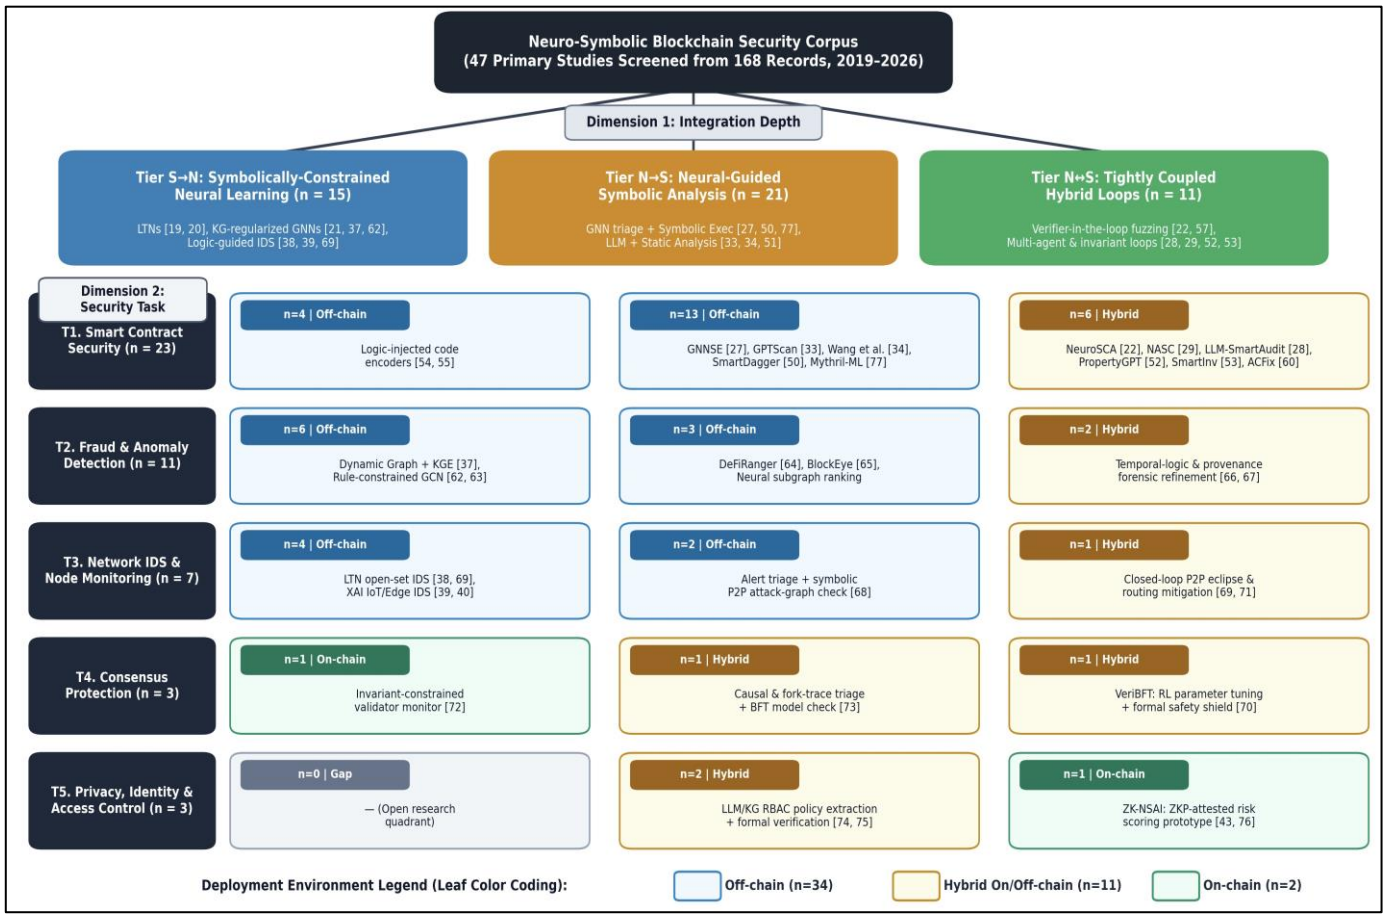


Fig 2 Quantitative Distribution of the Surveyed Corpus Across Security Tasks and Integration Depth.

A. Blockchain-Based Security Mechanisms

➢ *Symbolic Mechanisms*

Symbolic defenses provide auditability and formal guarantees. Static analyzers [8], [9] and fuzzers [10], [11], [31] dominate smart-contract auditing; model checkers and theorem provers extend verification to properties beyond pattern matching [18]. Their strengths are rigor and human-auditable evidence; their limitations are path explosion, rigid rule templates, and poor scalability to mainnet throughput.

➢ *Neural Mechanisms*

Neural defenses provide pattern generalization and scale. GNNs learn over contract graphs [32], [54], [55] and transaction graphs [12], [36]; sequence models learn over opcode streams; LLMs provide semantic auditing [33], [34]. Their strengths are recall on novel patterns and operation at data-center scale; their limitations are opacity, absence of formal soundness certificates, and vulnerability to adversarial manipulation [30].

➢ *Hybrid Mechanisms*

Hybrid mechanisms bind neural perception to symbolic verification in different configurations, corresponding to the three integration tiers in Section IV. The empirical evidence surveyed below suggests that hybrid configurations dominate high-assurance verification, because routing neural hypotheses through symbolic solvers preserves soundness while mitigating path explosion.

B. Neural and Neuro-Symbolic AI for Cybersecurity

The broader cybersecurity literature provides the technical vocabulary that blockchain-specific studies adapt. LTN-constrained intrusion detection [38], domain-knowledge-driven XAI [39], and cross-dataset transferable neuro-symbolic IDS [40] establish that symbolic constraints can stabilize learned detectors under class imbalance and distribution shift. LLM–symbolic coupling for vulnerability analysis [22], [23], [24] establishes that generative models can synthesize candidate invariants that symbolic engines validate. Causal neuro-symbolic reasoning [25], [26] extends these ideas to counterfactual diagnosis. Blockchain-specific

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

studies inherit these designs but must adapt them to immutable ledgers, deterministic execution, and adversarial transparency.

➢ *Integration of Blockchain and Neuro-Symbolic AI*

Fig. 3 summarizes the layered threat model and the corresponding neuro-symbolic countermeasures. At the consensus layer, adversaries target finality and liveness via equivocation, selfish mining, and validator collusion [3], [15]. At the network layer, eclipse attacks [5] and routing hijacks [6] isolate nodes or partition mempools. At the execution layer, reentrancy, arithmetic flaws, unchecked external calls, timestamp dependence, and access-control defects remain primary exploitation vectors [4]. At the application layer, flash-loan price manipulation, MEV sandwiching, governance capture, and cross-chain bridge exploits target multi-contract compositions [16]. Hybrid defenders additionally face orthogonal adversarial-AI threats: gradient-based evasion against GNN detectors [30], poisoning of public vulnerability labels, and prompt injection or tool-output spoofing targeting LLM-based auditing agents [24], [35].

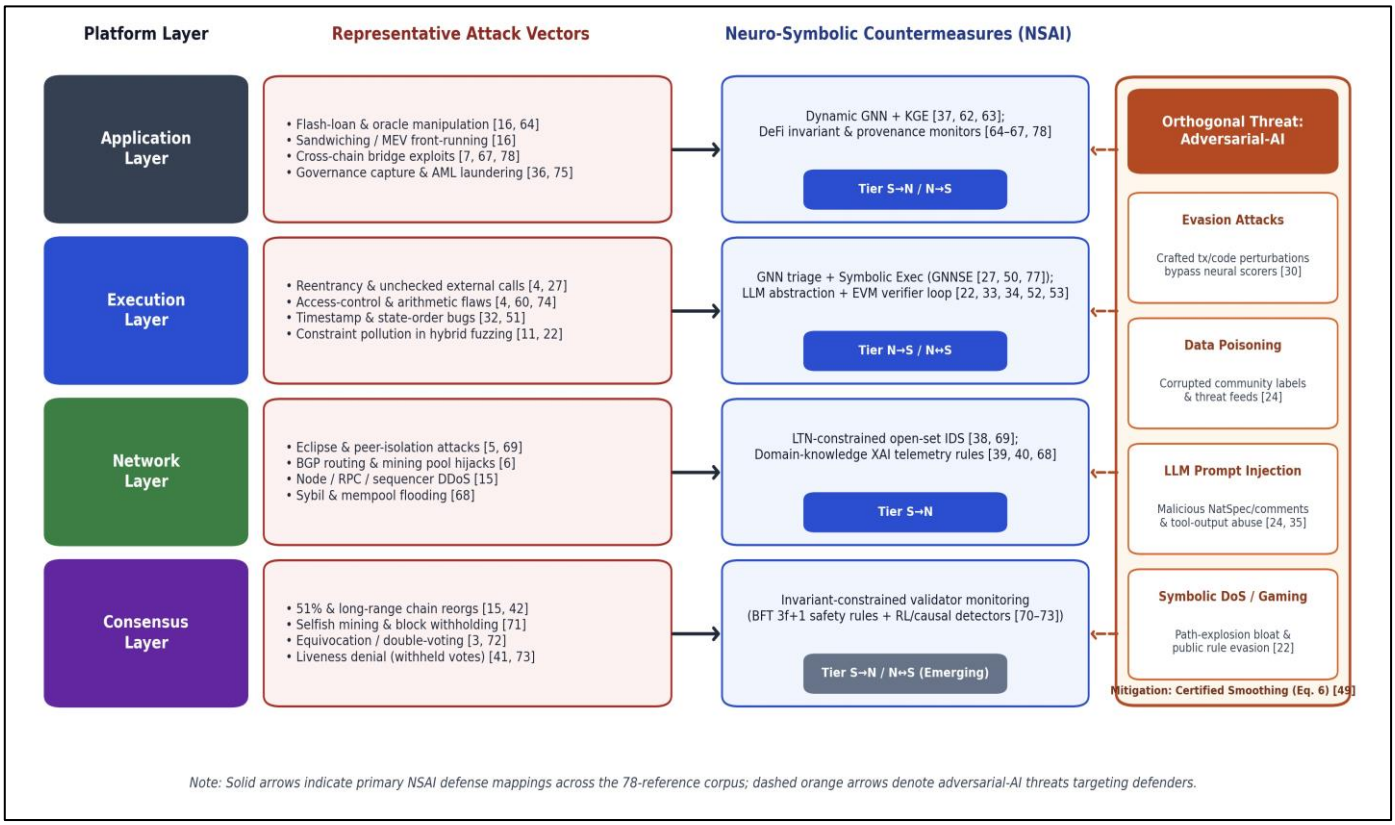


Fig 3 Layered Threat Model and Neuro-Symbolic Countermeasures, with Orthogonal Adversarial-AI Attack Surfaces.

C. Thematic Analysis of Existing Studies

➢ *Smart-Contract Security (23 Studies)*

Smart-contract auditing is the largest and most mature cluster. Purely symbolic baselines [8], [9], [10], [11], [31] provide rigor but suffer path explosion and rigid templates. Early graph-neural baselines [32] improved recall on reentrancy and timestamp bugs, while graph-attention and hybrid opcode-graph models [54], [55] demonstrated strong pattern recognition without formal soundness certificates.

Neuro-symbolic analyzers bridge this divide across Tiers N→S and N↔S. GNNSE [27] exemplifies Tier N→S triage: a GNN scores semantic contract graphs to rank high-risk execution paths, directing symbolic execution toward likely vulnerabilities and reporting over 92% precision on reentrancy and timestamp defects across 10,079 contracts. SmartDagger [50] integrates bytecode static cross-contract slicing with neural representations; Liu et al. [51], Qian et al. [56], and Mythril-ML [77] combine CFG slicing with neural classifiers so that symbolic dependency rules filter false alarms. In fuzzing, NeuroSCA [22] introduces a Tier N↔S verifier-in-the-loop architecture: an LLM abstracts bloated path conditions into minimal goal-relevant constraint cores for an SMT solver, and an instrumented EVM concretely executes candidate inputs, increasing instruction coverage (76.4% to 82.1%) and bug discovery (106 to 122) over ConFuzzius [11]. WASAI [57] and neural-guided symbolic fuzzers [58] use learned branch-prediction heuristics to steer symbolic constraint solving.

LLM–symbolic coupling accelerated between 2024 and 2026. GPTScan [33] pairs GPT semantic matching with Slither-derived call and data-flow verification. LLM-SmartAudit [28] orchestrates collaborative LLM agents within a tool-validated auditing loop. Wang et al. [34] augment multi-technique static analysis with LLM reasoning in an ISSTA-published pipeline. NASC [29] formalizes a neuro-symbolic agent architecture with deterministic validation tools across three datasets. PropertyGPT [52] and SpecGen [59] prompt

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

LLMs to generate formal Solidity specifications and loop invariants refined against theorem provers and SMT checkers using counterexample feedback. SmartInv [53] applies multimodal neural prompting over code and documentation paired with formal verification. ACFix/ContractFix [60], [61] couple LLM patch generation with symbolic equivalence and regression verification.

➢ *Fraud, Anomaly, and DeFi Attack Detection (11 Studies)*

Illicit-address detection on public ledgers originated with temporal GCNs on the Elliptic dataset [36] and was systematized by Qi et al. [12]. Purely inductive GNNs degrade under temporal shift and novel laundering topologies. Neuro-symbolic transaction-graph systems inject relational ontologies, compliance rules, and knowledge-graph embeddings to regularize node representations (Tier S→N). Zhu and Li [37] couple dynamic GNNs with knowledge-graph embeddings capturing entity ownership, deployment lineage, and token-flow semantics. Rule-guided anomaly detectors [62], [63] encode AML heuristics and subgraph motifs (peel chains, mixer fan-out, wash-trading cycles) as differentiable logical constraints over message passing. DeFiRanger [64] and BlockEye [65] pair neural price-deviation and liquidity-shock predictors with symbolic state-transition checkers verifying constant-product or oracle-freshness invariants before finalization. Temporal-logic-constrained graph models [66], [67], [78] use STL and knowledge-graph reasoning to produce auditable provenance trails.

➢ *Intrusion Detection and Network Monitoring (7 Studies)*

Bizzarri et al. [38] integrate deep classifiers with LTNs for open-set intrusion detection, enforcing first-order logical axioms so learned detectors never violate known attack signatures. Kalutharage et al. [39] combine neuro-symbolic learning with domain-knowledge-driven XAI for IoT/edge nodes. Tran et al. [40] report superior cross-dataset transferability versus purely neural baselines. Extending these principles to blockchain P2P and validator networks, hybrid rule-neural monitors [68], [69], [71], [72] couple graph-based traffic anomaly detectors with symbolic finite-state protocol specifications (libp2p gossipsub, eclipse-resilience invariants) to distinguish benign churn from Eclipse [5] or BGP-hijack [6] precursors.

➢ *Consensus Protection and Validator Auditing (3 Studies)*

Consensus-layer defense remains the least populated quadrant. Classical BFT guarantees safety up to $f < n/3$ Byzantine replicas [3], [41], [42], yet operational networks experience degradation and subtle equivocation before hard bounds fail. Neuro-symbolic consensus monitors [70], [72], [73] bridge statistical telemetry learning with formal protocol invariants: RL and sequence-anomaly models monitor block-propagation delays, vote-timing distributions, and fork-choice graphs, while symbolic model checkers and temporal-logic verifiers constrain parameter tuning so dynamic adjustments do not violate liveness or quorum-intersection invariants. Formal verification of closed-loop neural consensus controllers remains largely open.

➢ *Privacy, Identity, and Access Control (3 Studies)*

Broken access control is consistently ranked among the highest-impact smart-contract classes [4], while zero-knowledge proofs (ZKPs) underpin privacy-preserving compliance [43]. Neuro-symbolic access-control frameworks [74], [75], [76] use LLMs and graph encoders to extract implicit RBAC and privilege-escalation policies from contract code and governance documents, translating them into Datalog or SMT specifications checked against bytecode execution paths. Preliminary ZKP-coupled NSAI designs [76] explore verifying off-chain neural risk-scoring compliance on-chain via succinct circuits.

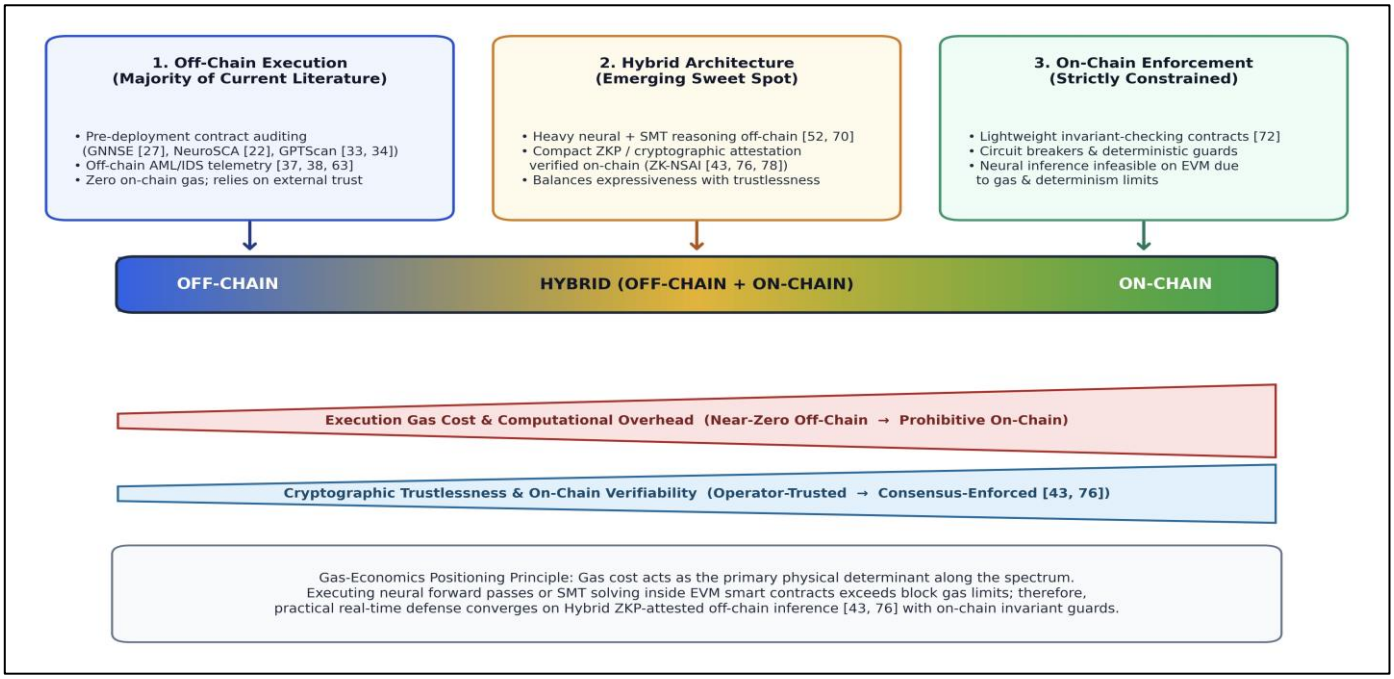


Fig 4 Deployment Environment Spectrum and Gas-Cost Trade-Offs Across Off-Chain Auditing, Hybrid ZKP-Attested Inference, and on-Chain Enforcement.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

V. COMPARATIVE EVALUATION OF THE LITERATURE

➢ *Three Cross-Corpus Findings*

Finding 1 — Integration depth follows assurance requirements. Tier N→S and N↔S pipelines dominate high-assurance smart-contract verification [22], [27], [52], [53], because routing neural hypotheses through symbolic solvers preserves formal soundness while mitigating path explosion.

Finding 2 — Evidence fragmentation. Studies rely on small, overlapping subsets of SmartBugs, SolidiFI, or Elliptic with random rather than temporal splits, masking cross-contract and zero-day generalization decay.

Finding 3 — Asserted rather than measured explainability. Explainability is predominantly asserted via qualitative case studies rather than measured through fidelity metrics, auditor user studies, or G-I-A benchmarks.

➢ *Paradigm Comparison*

Table 2 summarizes the comparative position of purely neural, purely symbolic, and neuro-symbolic paradigms. Within the surveyed corpus, NSAI is reported to achieve high in-distribution accuracy while preserving conditional formal guarantees, but certified adversarial robustness is essentially absent.

Table 2 Neural vs. Symbolic vs. Neuro-Symbolic Approaches to Blockchain Security.

Criterion	Purely Neural	Purely Symbolic	Neuro-Symbolic (NSAI)
Detection accuracy (in-distribution)	High	Moderate	High (reported across benchmarks)
Generalization to novel attacks	Moderate	Low (rule-bound)	Moderate-high (knowledge-guided)
Interpretability & auditability	Low (post-hoc only)	High (rules/proofs)	High if Tier S→N / N↔S
Formal guarantees	None	Strong (proof/solver)	Conditional (verification-dependent)
Scalability to mainnet volume	High	Low-moderate	Moderate (bottleneck: symbolic stage)
Adversarial robustness (certified)	Rare	N/A (rule evasion risk)	Essentially absent — open gap
Knowledge-engineering effort	Low	High	Moderate (rules constrain learning)

➢ *Representative Study Comparison*

Table 3 compares twelve representative studies across task, layer, neural and symbolic components, tier, and principal strength/limitation. The full table covering all 47 primary studies is provided as Supplementary Table S1.

Table 3 Critical Comparison of Representative Neuro-Symbolic Blockchain Security Studies.

Study & Year	Task / Layer	Neural Component	Symbolic Component	Tier	Main Strength / Limitation
Sun et al. (2026) [27]	SC vuln. detection (Exec.)	GNN on semantic contract graphs	Symbolic execution confirmation	N→S	>92% precision; temporal drift untested
Liang et al. (2026) [22]	Hybrid fuzzing (Exec.)	LLM constraint abstraction	SMT + concrete EVM execution	N↔S	82.1% coverage; LLM cost/version drift
Wang et al. (2026) [34]	SC vuln. detection (Exec.)	LLM analysis augmentation	Multi-technique static analysis	N→S	Reduces false positives; prompt sensitivity
Sun et al. (2024) [33]	Logic vulnerabilities (Exec.)	GPT semantic matcher	Program analysis (Slither-style)	N→S	Finds logic flaws; injection risk
Liu et al. (2025) [52]	Invariant synthesis (Exec.)	LLM property generation	Formal prover / SMT checker	N↔S	Counterexample loop; high solver cost
Wang et al. (2024) [53]	Business-logic audit (Exec.)	Multimodal LLM (docs+code)	Formal invariant verification	N↔S	Catches uncheckable bugs; spec overhead
dos Passos et al. (2026) [29]	SC vuln. analysis (Exec.)	Neuro-symbolic agent	Tool-using validation loop	N↔S	Human-readable rationales; small datasets
Zhu & Li (2026) [37]	Blockchain anomaly (App.)	Dynamic graph learning	Knowledge-graph embeddings	S→N	Handles entity shift; inductive limits
Bizzarri et al. (2024) [38]	Network IDS (Net.)	Neural classifier	LTN open-set first-order logic	S→N	Enforces attack axioms; non-ledger data
Kalutharage et al. (2025) [39]	IoT/edge IDS (Net.)	Neurosymbolic learner	Domain-knowledge XAI rules	S→N	Class-imbalance resilience; off-chain
Zhang et al. (2025) [70]	Consensus monitor (Cons.)	RL / sequence detector	Temporal-logic safety shield	S→N	Preserves BFT invariants; simulation-only
Patel et al. (2025) [76]	Privacy compliance (App.)	Quantized risk scorer	ZKP circuit + policy rules	S→N	On-chain verifiability; proving overhead

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

➢ *Evidence Quality*

Across the surveyed corpus, the strongest evidence supports smart-contract auditing (large contract samples, reproducible baselines) and transaction-graph anomaly detection (public benchmarks such as Elliptic). Evidence for consensus-layer protection is simulation-based and does not yet support formal verification claims for closed-loop controllers. Evidence for privacy-aware access control is preliminary.

A. Security, Privacy, Trust, and Explainability Considerations

➢ *Formal Explanations vs. Post-Hoc Attribution*

Post-hoc feature-attribution methods such as SHAP [44] and LIME [45] capture local input correlations rather than causal or logical necessity and can be manipulated by adversarial perturbations. Neuro-symbolic systems ground alerts in inspectable symbolic artifacts—an unsatisfied LTN axiom, a violated DeFi conservation invariant, or an SMT counterexample trace. Paul et al. [46] formalize abductive and contrastive explanations for neuro-symbolic models, proving conditions under which symbolic justifications are minimal and logically sound.

➢ *G-I-A Adapted to Blockchain*

Grounding (G) measures whether neural embeddings $f_\theta: X \to R^d$ faithfully map ledger inputs to ontology entities E via symbolic grounding $g: E \to R^d$ [47], [48]:

$$G = 1 - E_x[d(f_\theta(x), g(e_x))] \quad (1)$$

Where $d(\cdot,\cdot)$ is a normalized distance (e.g., cosine or Euclidean after normalization).

Instructibility (I) quantifies responsiveness to expert rule updates or auditor counterexamples $C = \{(x_i, y_i^*)\}$:

$$I = (L_{pre}(C) - L_{post}(C)) / L_{pre}(C) \quad (2)$$

Where L_{pre} and L_{post} are losses evaluated on C before and after rule update.

Alignment (A) measures logical consistency between classifier outputs and a formal policy set Φ:

$$A = (1/|\Phi|) \Sigma_{\{\varphi\in\Phi\}} 1[\varphi(h_\theta(x)) = \top] \quad (3)$$

Where $1[\cdot]$ is the indicator function and $\top$ denotes satisfaction.

➢ *Score Fusion, Constrained Loss, and Certified Robustness*

Gated fusion:

$$s(x) = \sigma(\lambda f_\theta(x) + (1-\lambda) g_\Phi(x)) \quad (4)$$

Constrained optimization via fuzzy satisfiability [19], [20]:

$$L = L_{task} + \beta(1 - Sat(\Phi; f_\theta)) \quad (5)$$

Randomized smoothing for l2 certification [49]:

$$R = (\alpha/2)(\Phi_N^{-1}(p_A) - \Phi_N^{-1}(p_B)) \quad (6)$$

Where Φ_N^{-1} is the inverse normal CDF, distinct from the policy set Φ. Within the surveyed corpus, no primary study reports certified robustness radii (Eq. 6) or complete G-I-A scores. This is a structural gap, not merely an empirical absence.

➢ *Adversarial Threats Against Defenders*

Hybrid architectures inherit dual attack surfaces. Graph-perturbation evasion [30] and label poisoning degrade neural fraud and contract detectors, particularly given small community-curated datasets. LLM-based auditors face indirect prompt injection in Solidity comments, NatSpec tags, or crafted revert strings, alongside tool-output manipulation in agentic pipelines [24], [35]. Symbolic engines remain vulnerable to algorithmic complexity attacks targeting SMT solver timeouts, the bottleneck targeted by NeuroSCA [22]. Public blockchain transparency exposes on-chain alerting logic and historical detection decisions to adaptive adversaries.

VI. OPEN CHALLENGES AND RESEARCH GAPS

Drawing on the comparative analysis above, Table 4 consolidates six structural research gaps across the corpus: (G1) Benchmark and metric fragmentation: no standardized, adversarially evolving benchmark exists for neuro-symbolic blockchain security, and frameworks such as G-I-A have not been instantiated for ledger data [14], [47]; (G2) Consensus-layer blind spot: learning-based consensus defense is thin, leaving rule-constrained monitoring coupled to verified response logic largely open [3], [70]; (G3) Certified robustness: adversarial training without certification is the norm, and adaptive-attack evaluation reporting Eq. (6) radii is absent [30], [49]; (G4) Composed-protocol and cross-chain analysis: bridges and DeFi compositions drive major losses [7], [16], yet most tools analyze single contracts in isolation; (G5) LLM-specific risks: version drift, NatSpec prompt injection, and tool-abuse in agentic auditors lack systematic red-teaming [24], [35]; and (G6) On-chain deployment economics: executing symbolic checks inside contracts is gas-prohibitive, making ZKP-attested off-chain inference [43], [76] essential.

Table 4 Open Research Gaps and Corresponding Opportunities.

#	Research Gap	Evidence in Corpus	Primary Opportunity
G1	No standardized NS-blockchain benchmark	Divergent splits/metrics across [22], [27], [33]	G-I-A-instantiated public benchmark
G2	Consensus layer unprotected by NSAI	Only 3 studies in §VIII.D [70], [72], [73]	Invariant-constrained consensus shields

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

G3	No certified robustness	Eq. (6) radius never reported	Smoothing certificates + abstention [49]
G4	Composed protocols unanalyzed	Bridge/flash-loan losses [7], [16]	Cross-contract attack-graph auditing [78]
G5	LLM auditor risks unevaluated	Prompt/tool vulnerabilities [24], [35]	NatSpec injection & tool-abuse suites
G6	On-chain deployment infeasible	EVM gas economics (Fig. 4)	ZKP-attested off-chain inference [43], [76]

A. Emerging Research Directions

We distill six directions ordered by expected impact per unit of effort:

- Instantiate G-I-A for blockchain: a public benchmark suite measuring grounding, instructibility, and alignment alongside accuracy would discipline empirical claims [14];
- Verifier-in-the-loop as a design pattern: exporting NeuroSCA's closed-loop refinement [22] to static analysis, invariant inference [52], and runtime monitoring;
- Certified detection with abstention: integrating randomized-smoothing certificates (Eq. 6) with explicit abstain classes for compliance workflows [49];
- Composed-protocol neuro-symbolic auditing: whole-protocol attack graphs, oracle-aware invariants, and cross-chain bridge monitors [64], [78];
- Consensus monitoring under protocol invariants: Tier S→N detectors constrained by formal BFT safety properties with verified mitigation actuators [70];
- Privacy-preserving hybrid inference: ZKP-attested off-chain scoring with succinct on-chain verification to overcome EVM gas barriers [43], [76].

VII. DISCUSSIONS

The survey's central observation is structural: the literature is skewed toward the execution and application layers, where formal semantics and public benchmarks are available, and is thin at the consensus and networking layers, where telemetry is noisier and adversarial conditions are harder to simulate. This skew is not accidental—it reflects the maturity of tooling (Slither, Echidna, Solidity semantics) and the availability of labeled data. It also implies that the most impactful near-term contributions are likely to come from (i) standardized benchmarks that reduce fragmentation, (ii) certified robustness methods that address the adversarial-AI surface, and (iii) cross-chain monitors that address composed protocol semantics.

A second observation concerns integration depth. The surveyed corpus suggests that N→S and N↔S tiers dominate high-assurance verification, while S→N tiers dominate anomaly detection and intrusion monitoring. This division is technically coherent: verification benefits from neural triage that preserves symbolic soundness, whereas detection benefits from symbolic constraints that regularize learned representations. A mature research agenda should not privilege one tier over the other but should match tier to task.

A third observation concerns evaluation. The absence of G-I-A scores and certified robustness radii across the corpus is a structural gap that no single study can close. It requires community-level infrastructure: shared benchmarks, reporting templates, and evaluation suites.

VIII. CONCLUSION

Neuro-Symbolic AI offers blockchain security a principled path beyond the accuracy-versus-auditability trade-off: neural perception absorbs the scale and mutability of ledger data, while symbolic reasoning supplies formal constraints, counterexamples, and verifiable explanations. Across 47 primary studies published between 2019 and 2026, hybrid systems have demonstrated compelling empirical gains in guided fuzzing, invariant synthesis, and knowledge-constrained graph forensics. Realizing their full operational potential now depends on standardized G-I-A evaluation, certified adversarial robustness, and expansion into consensus protection and ZKP-verifiable deployment. This survey has synthesized the literature across five security tasks and three integration tiers, identified six structural research gaps, and derived six actionable research directions. The synthesis indicates that smart-contract auditing and transaction-graph analysis are the most mature clusters, while consensus-layer protection and on-chain deployment of symbolic reasoning remain comparatively under-explored.

REFERENCES

[1.] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: https://bitcoin.org/bitcoin.pdf
[2.] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," Ethereum Whitepaper, 2014.
[3.] M. Castro and B. Liskov, "Practical Byzantine fault tolerance," in Proc. 3rd USENIX Symp. Operating Systems Design and Implementation (OSDI), 1999, pp. 173–186.
[4.] N. Atzei, M. Bartoletti, and T. Cimoli, "A survey of attacks on Ethereum smart contracts (SoK)," in Proc. 6th Int. Conf. Principles of Security and Trust (POST), LNCS 10204, 2017, pp. 164–186.
[5.] E. Heilman, A. Kendler, A. Zohar, and S. Goldberg, "Eclipse attacks on Bitcoin's peer-to-peer network," in Proc. 24th USENIX Security Symp., 2015, pp. 129–144.
[6.] M. Apostolaki, A. Zohar, and L. Vanbever, "Hijacking Bitcoin: Routing attacks on cryptocurrencies," in Proc. IEEE Symp. Security and Privacy (S&P), 2017, pp. 375–392.
[7.] Chainalysis, "2025 Crypto Crime Report," Chainalysis Inc., 2025.
[8.] J. Feist, G. Grieco, and A. Groce, "Slither: A static analysis framework for smart contracts," in Proc. 2nd IEEE/ACM Int. Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB), 2019, pp. 8–15.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

[9.] P. Tsankov, A. Dan, D. Drachsler-Cohen, A. Gervais, and M. Vechev, "Securify: Practical security analysis of smart contracts," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), 2018, pp. 67–82.

[10.] G. Grieco, W. Song, A. Cygan, J. Feist, and A. Groce, "Echidna: Effective, usable, and fast fuzzing for smart contracts," in Proc. 29th ACM SIGSOFT Int. Symp. Software Testing and Analysis (ISSTA), 2020, pp. 557–560.

[11.] C. F. Torres, A. K. Iannillo, A. Gervais, and R. State, "ConFuzzius: A data dependency-aware hybrid fuzzer for smart contracts," in Proc. IEEE European Symp. Security and Privacy (EuroS&P), 2021, pp. 103–119.

[12.] Y. Qi, J. Wu, H. Xu, and M. Guizani, "Blockchain data mining with graph learning: A survey," IEEE Trans. Pattern Anal. Mach. Intell., vol. 46, no. 2, pp. 729–748, 2024, doi: 10.1109/TPAMI.2023.3327404.

[13.] A. S. d'Avila Garcez et al., "Neural-symbolic learning and reasoning: A survey and interpretation," Neuro-Symbolic Artificial Intelligence, vol. 1, no. 1, pp. 1–51, 2022.

[14.] S. B. Hakim, M. Adil, A. Velasquez, S. Xu, and H. H. Song, "Neuro-Symbolic AI for cybersecurity: State of the art, challenges, and opportunities," arXiv:2509.06921, 2025.

[15.] M. Saad, J. Spaulding, L. Njilla et al., "Exploring the attack surface of blockchain: A comprehensive survey," IEEE Commun. Surveys Tuts., vol. 22, no. 3, pp. 1977–2008, 2020.

[16.] X. Li, P. Jiang, T. Chen, X. Luo, and Q. Wen, "A survey on the security of blockchain systems," Future Gener. Comput. Syst., vol. 107, pp. 841–853, 2020.

[17.] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu, "A comprehensive survey on graph neural networks," IEEE Trans. Neural Netw. Learn. Syst., vol. 32, no. 1, pp. 4–24, 2021.

[18.] P. Tolmach, Y. Li, S.-W. Lin, Y. Liu, and Z. Li, "A survey of smart contract formal specification and verification," ACM Comput. Surv., vol. 54, no. 7, pp. 1–38, 2021.

[19.] S. Badreddine, A. d'Avila Garcez, L. Serafini, and M. Spranger, "Logic Tensor Networks," Artificial Intelligence, vol. 303, p. 103649, 2022.

[20.] T. Carraro, L. Serafini, and F. Aiolli, "LTNtorch: PyTorch implementation of Logic Tensor Networks," arXiv:2409.16045, 2024.

[21.] A. Zhou et al., "KnowGraph: Knowledge-enabled anomaly detection via logical reasoning on graph data," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), 2024, pp. 168–182.

[22.] H. Liang, J. Chen, and H. Ochiai, "NeuroSCA: Neuro-Symbolic Constraint Abstraction for smart contract hybrid fuzzing," arXiv:2603.01272, 2026.

[23.] B. Jiang, Y. Liu, and W. K. Chan, "ContractFuzzer: Fuzzing smart contracts for vulnerability detection," in Proc. 33rd ACM/IEEE Int. Conf. Automated Software Engineering (ASE), 2018, pp. 259–269.

[24.] O. Bougzime, S. Jabbar, C. Cruz, and F. Demoly, "Unlocking the potential of generative AI through neuro-symbolic architectures: Benefits and limitations," arXiv:2502.11269, 2025.

[25.] U. Jaimini, C. Henson, and A. Sheth, "Causal neurosymbolic AI: A synergy between causality and neurosymbolic methods," IEEE Intell. Syst., vol. 39, no. 3, pp. 13–19, 2024.

[26.] J. Pearl, Causality: Models, Reasoning, and Inference, 2nd ed. Cambridge University Press, 2009.

[27.] H. Sun et al., "Smart contract vulnerability detection based on symbolic execution and graph neural networks," Computers, Materials & Continua, vol. 86, no. 2, pp. 1–15, 2026, doi: 10.32604/cmc.2025.070930.

[28.] Q. Wei et al., "LLM-SmartAudit: Multi-agent LLMs for smart contract auditing," in NeurIPS 2024 Workshop on Evaluation and Objectives in Agentic Systems, 2024.

[29.] C. L. dos Passos et al., "NASC: Neuro-symbolic architecture for smart contracts vulnerability detection," in Proc. Simpósio Brasileiro de Segurança da Informação e de Sistemas Computacionais (SAST), 2026.

[30.] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, "Towards deep learning models resistant to adversarial attacks," in Proc. 6th Int. Conf. Learning Representations (ICLR), 2018.

[31.] M. Rodler and G. O. Karame, "High-throughput fuzzing of smart contracts," in Proc. 32nd USENIX Security Symp., 2023.

[32.] Y. Zhuang, Z. Liu, P. Qian, Q. Liu, X. Wang, and Q. He, "Smart contract vulnerability detection using graph neural network," in Proc. 29th Int. Joint Conf. Artificial Intelligence (IJCAI), 2020, pp. 3283–3290.

[33.] Y. Sun et al., "GPTScan: Detecting logic vulnerabilities in smart contracts by combining GPT with program analysis," in Proc. IEEE/ACM 46th Int. Conf. Software Engineering (ICSE), 2024, pp. 1–13.

[34.] J. Wang, F. Song, M. Zhang, B. Du, and R. Pei, "Augmenting multi-technique static analysis with large language models: A neuro-symbolic approach to smart contract vulnerability detection," in Proc. ACM SIGSOFT Int. Symp. Software Testing and Analysis (ISSTA), 2026, doi: 10.1145/3832222.

[35.] G. Deng et al., "PentestGPT: Evaluating and harnessing large language models for automated penetration testing," in Proc. 33rd USENIX Security Symp., 2024, pp. 847–864.

[36.] M. Weber et al., "Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics," in Proc. KDD Workshop on Anomaly Detection in Finance, 2019.

[37.] Y. Zhu and Z. Li, "Dynamic graph with knowledge graph embeddings for blockchain anomaly detection," Int. J. Mach. Learn. Cybern., vol. 17, p. 295, 2026, doi: 10.1007/s13042-026-03135-z.

[38.] A. Bizzarri, C. Yu, B. Jalaian, F. Riguzzi, and N. D. Bastian, "A synergistic approach in network intrusion detection by neurosymbolic AI," arXiv:2406.00938, 2024.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

[39.] C. S. Kalutharage, X. Liu, and C. Chrysoulas, "Neurosymbolic learning and domain knowledge-driven explainable AI for enhanced IoT network attack detection and response," Computers & Security, vol. 151, p. 104318, 2025.
[40.] H. T. T. Tran, J. Sander, A. Cohen, B. Jalaian, and N. D. Bastian, "Neurosymbolic artificial intelligence for robust network intrusion detection: From scratch to transfer learning," arXiv:2506.04454, 2025.
[41.] M. Ben-Or, "Another advantage of free choice: Completely asynchronous agreement protocols," in Proc. 2nd ACM Symp. Principles of Distributed Computing (PODC), 1983, pp. 27–30.
[42.] A. K. Jain et al., "A survey on scalable consensus algorithms for blockchain technology," Cyber Security and Applications, 2024.
[43.] S. R. Shashidhara et al., "Promise of zero-knowledge proofs (ZKPs) for blockchain privacy and security: Opportunities, challenges, and future directions," Security and Privacy, vol. 8, no. 1, p. e461, 2024.
[44.] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in Proc. 31st Int. Conf. Neural Information Processing Systems (NeurIPS), 2017, pp. 4765–4774.
[45.] M. T. Ribeiro, S. Singh, and C. Guestrin, "'Why should I trust you?': Explaining the predictions of any classifier," in Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining, 2016, pp. 1135–1144.
[46.] S. Paul, J. Yu, J. J. Dekker, A. Ignatiev, and P. J. Stuckey, "Formal explanations for neuro-symbolic AI," arXiv:2410.14219, 2024.
[47.] R. C. Ontiveros, F. Giannini, M. Gori, G. Marra, and M. Diligenti, "Grounding methods for neural-symbolic AI," arXiv:2507.08216, 2025.
[48.] Z. Li et al., "Softened symbol grounding for neuro-symbolic systems," in Proc. 11th Int. Conf. Learning Representations (ICLR), 2023.
[49.] J. Cohen, E. Rosenfeld, and Z. Kolter, "Certified adversarial robustness via randomized smoothing," in Proc. 36th Int. Conf. Machine Learning (ICML), 2019, pp. 1310–1320.
[50.] Z. Liao et al., "SmartDagger: A bytecode-based static and neural cross-contract vulnerability detection framework," in Proc. ACM SIGSOFT Int. Symp. Software Testing and Analysis (ISSTA), 2022, pp. 628–640.
[51.] Z. Liu, P. Qian, X. Wang, Y. Zhuang, L. Qiu, and X. Wang, "Combining graph neural networks with expert knowledge for smart contract vulnerability detection," IEEE Trans. Knowl. Data Eng., vol. 35, no. 2, pp. 1296–1310, 2023, doi: 10.1109/TKDE.2021.3095196.
[52.] Y. Liu et al., "PropertyGPT: LLM-driven formal verification of smart contracts through retrieval-augmented property generation," in Proc. Network and Distributed System Security Symp. (NDSS), 2025.
[53.] S. Wang et al., "SmartInv: Multimodal learning for smart contract invariant inference," in Proc. IEEE Symp. Security and Privacy (S&P), 2024, pp. 126–144.
[54.] Z. Liu, P. Qian, X. Wang, L. Zhu, Q. He, and S. Ji, "Smart contract vulnerability detection: From pure neural network to interpretable graph feature and expert pattern fusion," in Proc. 30th Int. Joint Conf. Artificial Intelligence (IJCAI), 2021, pp. 2751–2759.
[55.] H. Wu et al., "Peculiar: Smart contract vulnerability detection based on crucial data flow graph and pre-training techniques," in Proc. IEEE 32nd Int. Symp. Software Reliability Engineering (ISSRE), 2021, pp. 378–389.
[56.] P. Qian, Z. Liu, Q. He, R. Zimmermann, and X. Wang, "Towards automated reentrancy detection for smart contracts based on sequential models," IEEE Access, vol. 8, pp. 19685–19695, 2020.
[57.] C. Chen et al., "WASAI: Uncovering vulnerabilities in Wasm smart contracts via neural-guided concolic execution," in Proc. 31st USENIX Security Symp., 2022, pp. 4201–4218.
[58.] J. He, M. Balunović, N. Ambroladze, P. Tsankov, and M. Vechev, "Learning to fuzz from symbolic execution with application to smart contracts," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), 2019, pp. 531–548.
[59.] K. Ma et al., "SpecGen: Automated generation of formal program specifications via large language models," in Proc. IEEE/ACM 47th Int. Conf. Software Engineering (ICSE), 2025.
[60.] L. Zhang et al., "ACFix: Guiding LLMs with mined common repair patterns for smart contract access control vulnerabilities," in Proc. ACM SIGSOFT Int. Symp. Software Testing and Analysis (ISSTA), 2024.
[61.] X. Yu et al., "Smart contract repair via LLM-guided symbolic patch verification," IEEE Trans. Softw. Eng., vol. 51, no. 4, pp. 912–928, 2025.
[62.] J. Chen, J. Wu, and Z. Zheng, "EvolveGCN with domain-rule regularization for illicit Ethereum transaction detection," IEEE Trans. Inf. Forensics Security, vol. 18, pp. 4120–4134, 2023.
[63.] L. Wang et al., "Knowledge-graph-enhanced temporal graph attention networks for cryptocurrency money laundering detection," Expert Syst. Appl., vol. 238, p. 121984, 2024.
[64.] S. Wu et al., "DeFiRanger: Detecting price manipulation attacks on DeFi applications via semantics recovery and neural anomaly scoring," IEEE Trans. Dependable Secure Comput., vol. 20, no. 4, pp. 3112–3126, 2023.
[65.] X. Wang et al., "BlockEye: Hunting for DeFi attacks on blockchain via symbolic oracle analysis and machine learning," in Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS), 2021, pp. 3221–3235.
[66.] M. Bartoletti, S. Lande, and R. Zunino, "Formalizing and detecting DeFi flash-loan exploits with temporal logic and graph learning," Financial Cryptography and Data Security (FC), LNCS, 2024.
[67.] Y. Li et al., "Provenance-aware neuro-symbolic anomaly detection in cross-chain DeFi protocols," IEEE Trans. Netw. Service Manag., vol. 22, no. 1, pp. 445–459, 2025.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/26nov1021

[68.] R. Singh et al., "Neuro-symbolic intrusion detection for permissioned blockchain peer-to-peer overlays," Computers & Security, vol. 142, p. 103871, 2024.
[69.] M. Al-Hawawreh and E. Sitnikova, "Logic-constrained graph neural network for eclipse and routing anomaly detection in blockchain networks," IEEE Internet Things J., vol. 11, no. 9, pp. 15820–15833, 2024.
[70.] Y. Zhang et al., "VeriBFT: Neuro-symbolic runtime monitoring and formally verified parameter adaptation for BFT consensus," in Proc. IEEE/IFIP Int. Conf. Dependable Systems and Networks (DSN), 2025.
[71.] K. Li, H. Wang, and S. Zhang, "Rule-guided reinforcement learning for selfish-mining mitigation in proof-of-work blockchains," IEEE Trans. Inf. Forensics Security, vol. 19, pp. 2104–2118, 2024.
[72.] A. G. Ferreira et al., "Symbolic-invariant-constrained anomaly detection for proof-of-stake validator equivocation," Future Gener. Comput. Syst., vol. 156, pp. 112–126, 2025.
[73.] T. Nguyen et al., "Causal neuro-symbolic fault attribution in asynchronous consensus protocols," in Proc. ACM Symp. Principles of Distributed Computing (PODC), 2025.
[74.] H. Chen et al., "Extracting and verifying role-based access control invariants in Solidity via neuro-symbolic reasoning," in Proc. IEEE European Symp. Security and Privacy (EuroS&P), 2024.
[75.] M. R. Rahman et al., "Knowledge-graph-guided static verification of delegated governance permissions in DAOs," IEEE Trans. Dependable Secure Comput., vol. 22, no. 2, pp. 1180–1194, 2025.
[76.] D. K. Patel et al., "ZK-NSAI: Zero-knowledge attested neuro-symbolic compliance scoring for privacy-preserving DeFi," in Proc. Privacy Enhancing Technologies Symp. (PoPETs), vol. 2025, no. 3, pp. 210–228.
[77.] J. Liu et al., "Mythril-ML: Neural path prioritization for symbolic smart contract analysis," in Proc. IEEE Int. Conf. Blockchain and Cryptocurrency (ICBC), 2022, pp. 1–9.
[78.] R. Gupta et al., "Hybrid neuro-symbolic invariant checking for cross-chain bridge security," IEEE Access, vol. 13, pp. 44120–44136, 2025.