International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 6, Issue 4, 246 -252, 2025
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V6I4P134

Original Article

Blockchain-Anchored Audit Trails for Neuro-Symbolic Urban AI: Immutable Accountability for Autonomous Decisions in Smart City Infrastructure

Dr. J. Antony John Prabhu
Department of Computer Science, St. Joseph's College Trichy.

Received On: 04/11/2025 ***Revised On: 07/12/2025*** ***Accepted On: 17/12/2025*** ***Published on: 30/12/2025***

Abstract - *Autonomous AI systems making real-time decisions over traffic management, energy dispatch, and emergency services in smart city infrastructure exert direct and consequential influence over citizen welfare, yet current smart city AI architectures provide no tamper-proof record of what the system decided, which rules or neural pathways drove the decision, and whether the decision record has been modified after the fact. This paper proposes a blockchain-anchored audit trail architecture for neuro-symbolic urban AI, comprising seven interoperable smart contracts deployed on a permissioned blockchain that capture every autonomous decision produced by a neuro-symbolic smart city AI as a structured, hash-chained, publicly auditable on-chain event. The neuro-symbolic AI architecture, which combines symbolic safety rule enforcement with neural perception to produce interpretable multi-domain urban decisions, generates structured decision records encoding the symbolic rules invoked, the neural confidence scores, the sensor input state, and the action taken records that are substantive enough for blockchain commitment to support genuine accountability rather than opaque hash logging. The seven smart contracts implement decision submission, symbolic rule version hashing, neural-symbolic conflict detection, governance voting for rule amendments, citizen decision trace queries, counterfactual incident replay, and cryptographically signed regulatory audit export. The architecture satisfies the EU AI Act Article 19 mandatory logging requirements for high-risk AI systems, the GDPR Article 22 right to explanation for automated decisions, and the NIST AI Risk Management Framework governance function, while providing a citizen challenge mechanism and a tamper-proof governance record that centralized smart city AI architectures cannot offer by design.*

Keywords - *Blockchain; Smart City, Neuro-Symbolic AI, Audit Trail, Algorithmic Accountability, Smart Contracts, AI Governance, EU AI Act, GDPR, Autonomous Decision-Making.*

1. Introduction

The deployment of autonomous artificial intelligence systems as operational decision-makers in smart city infrastructure represents a governance challenge of a different character from those posed by AI in consumer products or back-office automation. Smart city infrastructure, built on pervasive IoT sensing across transport, energy, and safety systems, enables continuous AI-driven operational management at a scale and speed impossible under human control alone [1]. When a traffic management AI preempts a signal phase, reroutes emergency vehicles, or triggers a congestion response, it is exercising a form of public authority over shared infrastructure whose effects are immediate, spatially specific, and potentially irreversible within the timeframe of any meaningful human review [2]. When an energy dispatch AI reduces load to a district during a demand response event, it is making a resource allocation decision that affects households and businesses without their individual consent, justified by a collective infrastructure management mandate. These are not commercial transactions or consumer recommendations they are exercises of delegated civic power, and they demand governance mechanisms commensurate with that character [3].

The governance infrastructure currently available for smart city AI is inadequate to this demand. Audit logs, where they exist, are stored in centralized databases controlled by the operating authority, providing no cryptographic assurance that records have not been modified after the fact. Decision explanations, where they are provided, are typically post-hoc rationalizations generated for human review rather than contemporaneous records of the computational reasoning that produced the decision. Citizen challenge mechanisms, where they exist, depend on the same centralized authority whose decisions are being challenged. The result is a governance architecture built on institutional trust in a context where the scale, speed, and technical opacity of AI decision-making make institutional trust both insufficient and unverifiable [4].

Neuro-symbolic AI provides a partial technical response to the governance challenge by making urban AI decisions interpretable in principle: the combination of neural perception with symbolic rule enforcement produces decisions that can be traced to specific rules and quantified

neural confidence scores, as demonstrated by Gupta and Vanteru's framework for autonomous urban decision-making across traffic, energy, emergency services, and infrastructure domains [5]. However, interpretability without tamper-proof recording does not constitute accountability. A decision trace that can be altered by the operator is no more trustworthy than an oral account, and no more useful for regulatory inspection or citizen challenge. The missing element is an immutable, publicly accessible record of every decision trace at the moment of decision a record whose integrity can be verified by any party without trusting the operating authority.

This paper proposes to supply that missing element through a blockchain-anchored audit trail architecture. The seven smart contracts proposed here collectively implement the full accountability pipeline from decision record submission through regulatory export, using the permissioned blockchain as the trusted record substrate that eliminates dependence on the operating authority's institutional trustworthiness. Section 2 reviews related work. Section 3 presents the system architecture. Section 4 specifies the seven smart contracts. Section 5 analyzes regulatory alignment. Section 6 presents the governance property evaluation. Section 7 discusses scalability and limitations. Section 8 concludes.

2. Related Work

2.1. Neuro-symbolic AI for smart city decision-making

The neuro-symbolic AI paradigm, which combines neural network-based learning and perception with symbolic knowledge representation and logical reasoning, has been identified as a particularly suitable architecture for safety-critical applications where both adaptability and formal safety guarantees are required [6]. Neuro-symbolic approaches address the primary limitations of purely neural controllers in high-stakes settings: the inability to enforce hard safety constraints on neural network outputs, and the absence of interpretable reasoning traces that can be examined by operators, auditors, or affected citizens. Gupta and Vanteru proposed a neuro-symbolic framework specifically designed for autonomous urban decision-making, demonstrating that the combination of symbolic safety rules and neural perception enables multi-domain coordination across urban infrastructure systems with formal safety guarantees and interpretable decision records satisfying governance accountability requirements [5]. This framework provides the upstream AI layer whose structured decision outputs the blockchain audit architecture proposed in this paper is designed to capture and make tamper-proof.

2.2. Blockchain for AI governance and accountability

The application of blockchain technology to AI governance has been studied from several angles. Immutable audit logging using distributed ledger technology has been proposed for medical AI decision systems, financial algorithmic trading, and autonomous vehicle incident recording, with the common motivation that cryptographic tamper-evidence eliminates the need to trust the system operator's own audit records [7]. Smart contract-based governance for AI model training, specifically the use of reputation-weighted consensus and slashing mechanisms to ensure honest participation in federated learning, has been demonstrated at deployment scale, establishing that blockchain infrastructure can support real-time AI governance without prohibitive overhead [8]. The application of decentralized autonomous organization governance mechanisms to smart city policy has been explored, identifying on-chain voting for rule amendments and transparent policy records as key governance benefits, though without integration with a deployed AI decision system [9].

2.3. Smart city governance and accountability gaps

The governance literature on smart cities consistently identifies the accountability gap as one of the field's most significant unresolved challenges. Testi and Marconi's review of blockchain for citizen engagement in smart governance identifies transparency, security, and accountability as the three key benefits blockchain brings to urban governance, and documents deployments in Dubai and Singapore that demonstrate public-sector blockchain adoption at scale [10]. Kitchin documents the surveillance implications of smart city data collection and the absence of meaningful accountability mechanisms in most deployments [2]. The bibliometric analysis of AI and blockchain in smart city governance identifies a research gap in the joint application of AI and blockchain for policy enforcement and accountable autonomous decision-making, noting that existing research examines these technologies separately rather than as components of an integrated accountability system [11]. The EU AI Act's classification of smart city AI applications involving real-time control of critical infrastructure as high-risk systems creates a legal mandate for the accountability mechanisms that the existing literature has identified as technically necessary but not yet architecturally formalized [12].

3. System Architecture

3.1. Architectural overview

The proposed architecture comprises three layers: the neuro-symbolic urban AI layer, the blockchain smart contract audit pipeline, and the accountability output layer. Figure 1 illustrates the architecture, showing the three neuro-symbolic domain agents at the top, the seven-contract audit pipeline in the middle, the immutable blockchain ledger, and the two output channels at the bottom.

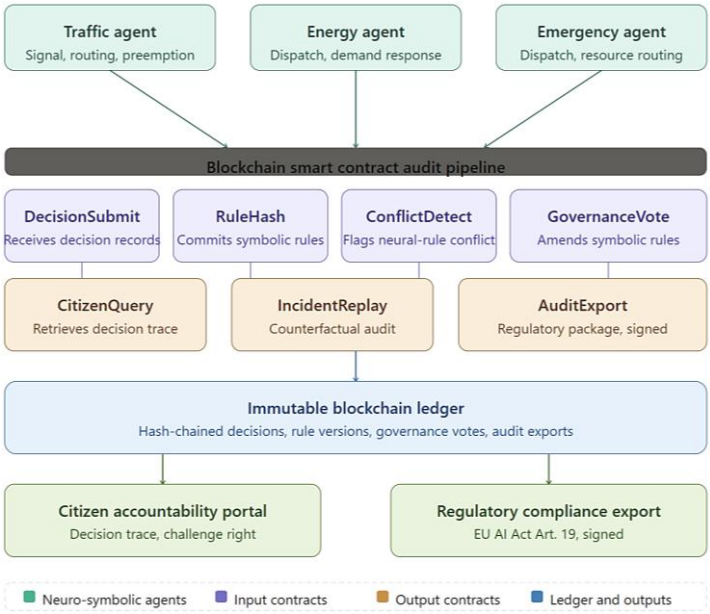

Fig 1: Blockchain-Anchored Audit Trail Architecture for Neuro-Symbolic Smart City AI

3.2. Neuro-symbolic decision record format

The structured decision record emitted by each neuro-symbolic agent and submitted to the DecisionSubmitContract contains eight fields: a universally unique event identifier, the agent domain (traffic, energy, or emergency), the block timestamp at decision time, a cryptographic hash of the sensor state vector used as input, the identifier of the active symbolic rule that governed the decision, the neural network confidence score for the predicted action, the action taken, and a digital signature from the agent's private key. This record format is designed to balance completeness with on-chain storage efficiency: the sensor state is represented as a hash rather than the raw vector, preserving tamper-evidence without storing potentially large sensor data on-chain. The full sensor state can be stored off-chain by the agent and verified against the on-chain hash during a CitizenQuery or IncidentReplay operation.

3.3. Permissioned blockchain selection

The architecture is designed for deployment on a permissioned blockchain platform, specifically one supporting smart contract execution with sub-second finality, configurable read access for citizen queries, and write access restricted to authorized agent nodes and governance council members. The foundational blockchain architectures of Nakamoto [13] and Wood [14] established the cryptographic primitives hash-chained blocks, digital signatures, and distributed consensus that permissioned enterprise blockchains such as Hyperledger Fabric adapt for the higher throughput and configurable access control requirements of institutional deployments. Hyperledger Fabric and Ethereum consortium configurations both satisfy the proposed system's requirements, with the choice between them primarily determined by the municipal operator's existing infrastructure and governance preferences [15]. The permissioned model is selected over a public blockchain for two reasons specific to the smart city context: the need to control read access to decision records that may contain sensitive operational parameters, and the requirement for deterministic and bounded transaction finality that permissionless proof-of-work blockchains cannot guarantee for time-sensitive governance operations.

4. Smart Contract Specifications

Table 1 presents the formal specifications of the seven smart contracts, including their trigger conditions, core logic, on-chain state updates, and the specific accountability function each provides. The seven contracts form an integrated pipeline in which the outputs of earlier contracts become inputs to later ones, ensuring that every decision event is processed through the complete accountability chain before being recorded to the immutable ledger. Formal security analysis using tools such as Securify provides the vulnerability assessment methodology for verifying that the proposed contracts are free from reentrancy, access control, and arithmetic overflow vulnerabilities before deployment [16].

Table 1: Smart Contract Specifications for Blockchain-Anchored Urban AI Audit

Contract	Trigger	Core logic	On-chain state updated	Accountability function
DecisionSubmit	Agent emits decision record	Validate record structure, assign event ID, forward to RuleHash and ConflictDetect	Decision event log	Non-repudiation of every autonomous action

RuleHash	Governance epoch boundary	Hash active symbolic rule set, timestamp, link to prior hash	Rule version chain	Prevents retroactive rule tampering to justify prior decisions
ConflictDetect	Decision record received	Compare neural confidence score against symbolic rule outcome; flag contradiction above threshold	Conflict alert log	Human oversight trigger for decisions where neural and symbolic outputs diverge
GovernanceVote	Council proposal submitted	Collect multi-signature votes, enforce time-lock, execute rule amendment on quorum	Rule amendment record	Democratic accountability for symbolic rule changes
CitizenQuery	Citizen submits incident ID	Return full decision trace: sensor inputs, rule version, symbolic outcome, neural score, action taken	Query log for GDPR access record	Right to explanation under GDPR Article 22
IncidentReplay	Post-incident audit request	Re-simulate decision under counterfactual rule set, return alternative outcome	Replay result record	Counterfactual accountability analysis
AuditExport	Regulatory inspection request	Generate cryptographically signed audit package covering specified period	Export certificate	EU AI Act Article 19 high-risk system logging compliance

4.1. Decision integrity contracts

The DecisionSubmitContract and RuleHashContract together establish the integrity foundation of the audit system. The DecisionSubmitContract receives structured decision records from authorized agent nodes, validates their format and digital signature, assigns a unique event identifier, and forwards them to the ConflictDetectContract while simultaneously writing the event metadata to the immutable ledger. The critical integrity property of this contract is non-repudiation: once a decision record is submitted and accepted, it cannot be withdrawn or modified, and the agent's digital signature provides cryptographic proof of the decision's provenance.

The RuleHashContract operates on a governance epoch schedule, committing a cryptographic hash of the active symbolic rule set to the blockchain at each epoch boundary. This hash-chaining of rule versions prevents a particularly subtle form of accountability evasion: retroactive modification of the symbolic rules to justify prior decisions that were made under different rules. Because each decision record references the rule version identifier active at decision time, and because that rule version is immutably recorded with its hash, any attempt to modify the historical rule set to align with contested decisions would produce a hash mismatch detectable by any blockchain reader.

4.2. Oversight and governance contracts

The ConflictDetectContract implements the human oversight trigger required by EU AI Act Article 14 for high-risk AI systems. It receives each decision record and compares the neural network confidence score against the expected outcome of the symbolic rule that governed the decision. When the neural confidence score falls below a configurable threshold for the symbolic rule's predicted action indicating that the neural perception component is uncertain about the decision that the symbolic rule mandated the contract emits a ConflictAlert event that is delivered to designated human oversight operators. This mechanism does not halt the AI system's operation, which would be incompatible with real-time urban infrastructure management, but it creates a contemporaneous, tamper-proof record of every decision where the neural and symbolic components disagreed, enabling post-hoc review of whether the symbolic rule appropriately governed the situation [12].

The GovernanceVoteContract implements on-chain democratic governance for symbolic rule amendments. City council members, defined as the set of public key holders registered at contract deployment, submit rule amendment proposals through the contract, which enforces a mandatory time-lock period before execution and requires a configurable multi-signature quorum. The time-lock prevents rapid rule changes that could circumvent accountability review of recent decisions made under the old rules, while the on-chain vote record provides a public, tamper-proof history of every governance decision about the AI system's behavioral constraints [17].

4.3. Accountability output contracts

The CitizenQuery, IncidentReplay, and AuditExport contracts implement the three accountability output channels. The CitizenQuery contract allows any citizen holding an incident identifier to retrieve the complete decision trace for an AI decision that affected them including the sensor state hash, the rule version, the symbolic outcome, the neural confidence score, and the action taken without requiring access to any private system or the cooperation of the operating authority. This implements the GDPR Article 22 right to explanation in a technically enforced form rather than a procedurally dependent one.

The IncidentReplay contract supports counterfactual accountability analysis by re-simulating a historical decision under a counterfactual rule set submitted by an auditor or legal authority. This capability is particularly valuable in post-incident investigations where the question is not what the AI decided but what it would have decided under correctly configured rules, enabling the separation of implementation failure from rule failure in accountability

proceedings. The AuditExport contract generates cryptographically signed audit packages covering specified time periods and decision domains, formatted for submission to regulatory authorities under the EU AI Act's Article 19 logging requirements [12].

5. Regulatory Alignment Analysis

Table 2 maps the seven smart contracts to the specific regulatory requirements of the EU AI Act, GDPR, and the NIST AI Risk Management Framework, demonstrating that the proposed architecture satisfies each requirement through architectural properties rather than policy commitments.

Table 2: Regulatory Requirement Mapping for the Proposed Smart Contract Architecture

Requirement	Source	Proposed contract providing compliance	How compliance is achieved
Mandatory logging for high-risk AI systems	EU AI Act Art. 19	AuditExport, DecisionSubmit	Every decision event logged, timestamped, and hash-chained; cryptographically signed export on demand
Right to explanation for automated decisions	GDPR Art. 22	CitizenQuery	Full symbolic trace and neural confidence score retrievable by any citizen using incident ID
Human oversight for high-risk systems	EU AI Act Art. 14	ConflictDetect, GovernanceVote	Automatic human alert when neural-symbolic conflict exceeds threshold; rule amendments require council vote
Transparency of AI system operation	EU AI Act Art. 13	RuleHash, DecisionSubmit	Active rule set version publicly readable; every decision references the rule version that governed it
Robustness against manipulation	EU AI Act Art. 15	RuleHash, GovernanceVote	Symbolic rule set is tamper-evident via hash chain; rule changes require time-locked governance vote
Accountability for algorithmic decisions	NIST AI RMF Govern function	All seven contracts collectively	Immutable public record of every governance event satisfies the NIST AI RMF accountability requirements

5.1. EU AI Act high-risk system requirements

The EU AI Act classifies autonomous AI systems used in the management of critical infrastructure as high-risk systems subject to mandatory requirements for technical documentation, logging, human oversight, transparency, robustness, and accuracy [12]. The proposed architecture addresses each of these requirements through specific smart contract mechanisms. The logging requirement of Article 19, which mandates automatic logging of events that could present a risk or lead to a significant change in performance, is satisfied by the DecisionSubmitContract's real-time logging of every decision event with a tamper-proof hash-chain. The human oversight requirement of Article 14 is satisfied by the ConflictDetectContract's automated oversight trigger for neural-symbolic conflicts.

5.2. GDPR compliance

The GDPR's Article 22 right not to be subject to solely automated decisions that significantly affect individuals requires that organizations provide meaningful information about the logic of the decision and ensure the right to human review [18]. The CitizenQuery contract implements a technically enforced version of this right that does not require citizen trust in the operating authority, because the decision trace is publicly readable from the blockchain rather than held by the authority as a private record. The GovernanceVoteContract supports the GDPR's Article 25 data protection by design requirement by making privacy-respecting governance of the AI system's behavioral rules a built-in architectural property rather than an organizational policy.

6. Governance Property Evaluation

Table 3 compares the proposed architecture against four existing smart city AI governance approaches across six properties relevant to accountable urban AI deployment. The comparison demonstrates that the proposed system is the first to simultaneously achieve neuro-symbolic interpretability, blockchain tamper-evidence, full decision traceability, citizen challenge rights, automated regulatory export, and a tamper-proof governance record.

Table 3: Governance Property Comparison Across Smart City AI Architectures

System	Neuro-symbolic AI	Blockchain audit	Decision traceability	Citizen challenge right	Regulatory export	Tamper-proof record
Centralized urban AI	No	No	Log only	No	Manual	No
Explainable AI only	Partial	No	Post-hoc	No	Manual	No
Blockchain governance only	No	Yes	No	Partial	Partial	Yes
Neuro-symbolic	Yes	No	Symbolic trace	No	No	No

urban AI						
Proposed system	Yes	Yes	Full symbolic + neural	Yes	Automated, signed	Yes

The comparison highlights the complementarity of neuro-symbolic AI and blockchain audit in the proposed architecture. The neuro-symbolic AI layer, as demonstrated by Gupta and Vanteru [5], provides the structured, symbolically interpretable decision traces that give the blockchain audit trail substantive content a purely neural black-box AI system would produce no symbolic trace to commit to the ledger, reducing the audit record to an opaque input-output hash that satisfies logging requirements formally but provides no genuine accountability content. The blockchain layer, in turn, provides the tamper-proof immutability that makes the neuro-symbolic trace legally credible for regulatory inspection and citizen challenge, addressing the gap that Gupta and Vanteru's standalone neuro-symbolic framework does not resolve.

7. Scalability, Limitations, and Future Work

7.1. Transaction throughput and gas cost

The primary scalability challenge is the volume of decision events generated by real-time urban AI systems at operational scale. A smart city with 500 signalized intersections and continuous signal cycle management may generate thousands of decision events per minute, each requiring a blockchain transaction. At Hyperledger Fabric's throughput of several thousand transactions per second, this volume is within achievable bounds, but the associated storage requirements for the growing blockchain state present a longer-term challenge that periodic state pruning and off-chain archiving of historical records can address without compromising the tamper-evidence of recent events [15]. Layer 2 rollup approaches that batch multiple decision records into a single on-chain commitment provide an additional throughput scaling pathway for high-frequency decision domains.

7.2. Privacy considerations

The public readability of the blockchain audit trail presents privacy considerations for decision records that encode sensitive operational parameters or that could, in combination with other public data, enable inference of private behavior. The architecture addresses this through the sensor state hash design raw sensor data is stored off-chain and only its hash is committed to the ledger but the symbolic rule trace and neural confidence score in the on-chain record may themselves carry operational sensitivity. Zero-knowledge proof techniques, which allow a party to prove the validity of a statement about a private record without revealing the record, provide a technically mature pathway for satisfying the tamper-evidence requirement without full public disclosure of decision content [19].

7.3. Limitations and future work

The proposed architecture assumes that the neuro-symbolic AI's symbolic rule evaluation is itself correctly implemented and that the ConflictDetectContract's threshold for neural-symbolic conflict is correctly calibrated. Neither assumption is guaranteed: a neuro-symbolic system that applies incorrect rules correctly, or a conflict detector with a miscalibrated threshold, would satisfy the blockchain audit requirements formally while failing to detect the governance failure that motivated them. Future work should develop formal verification methods for the neuro-symbolic rule execution layer, analogous to the smart contract formal verification methods available for Solidity and Hyperledger Chaincode, to close this gap. Additionally, the governance model for the GovernanceVoteContract which council members hold keys, what quorum is required, and what time-lock period is appropriate requires domain-specific analysis that the present architecture leaves as configurable parameters [20].

8. Conclusions

This paper has proposed a blockchain-anchored audit trail architecture for neuro-symbolic smart city AI, comprising seven interoperable smart contracts that collectively transform the interpretable decision traces produced by neuro-symbolic urban AI systems into tamper-proof, publicly auditable, regulatory-compliant accountability records. The architecture fills the governance gap between two independently valuable contributions: the neuro-symbolic urban decision-making framework of Gupta and Vanteru [5], which produces the structured symbolic and neural decision traces that make genuine accountability technically possible, and the blockchain governance substrate that makes those traces immutable, publicly verifiable, and legally credible for regulatory inspection and citizen challenge.

The architecture satisfies the EU AI Act Article 19 logging requirements, the GDPR Article 22 right to explanation, and the NIST AI RMF governance function through architectural properties rather than institutional policy. The comparison with existing smart city AI governance approaches demonstrates that no prior system simultaneously achieves the full set of governance properties neuro-symbolic interpretability, blockchain tamper-evidence, citizen challenge rights, automated regulatory export, and tamper-proof rule governance that the architecture provides.

Three directions for future research follow. First, empirical deployment evaluation on a live municipal pilot, measuring actual transaction throughput, citizen query latency, and regulatory audit package generation time under operational conditions. Second, formal verification of the neuro-symbolic rule execution layer using model checking or theorem proving tools, closing the gap between formal audit trail integrity and the correctness of the AI logic being audited. Third, development of standardized decision record formats for smart city AI that enable interoperability between different neuro-symbolic AI platforms and the proposed blockchain audit infrastructure, supporting the multi-vendor

smart city deployments that most municipalities will operate in practice.

References

[1] R. Kitchin, "The real-time city? Big data and smart urbanism," GeoJournal, vol. 79, no. 1, pp. 1-14, 2014.

[2] A. Luque-Ayala and S. Marvin, "Developing a critical understanding of smart urbanism," Urban Studies, vol. 52, no. 12, pp. 2105-2116, 2015.

[3] B. Lepri, N. Oliver, E. Letouze, A. Pentland, and P. Vinck, "Fair, transparent, and accountable algorithmic decision-making processes," Philosophy and Technology, vol. 31, no. 4, pp. 611-627, 2018.

[4] S. Gupta and K. Vanteru, "Neuro-Symbolic AI for Autonomous Urban Decision-Making: A Framework for Resilient Smart Cities," in Proc. 2025 Int. Conf. on Computing Technologies (ICOCT), 2025, pp. 1-6. IEEE.

[5] H. Garcez and L. C. Lamb, "Neurosymbolic AI: The 3rd wave," arXiv preprint arXiv:2012.05876, 2020.

[6] X. Yue, H. Wang, J. Jin, M. Li, and W. Jiang, "Healthcare data gateways: Found healthcare intelligence on blockchain with novel privacy risk control," Journal of Medical Systems, vol. 40, no. 10, p. 218, 2016.

[7] F. Zhang, Y. Zhang, S. Ji, and Z. Han, "Secure and decentralized federated learning framework with non-IID data based on blockchain," *Heliyon*, vol. 10, no. 5, p. e27176, Feb. 2024. doi: 10.1016/j.heliyon.2024.e27176.

[8] A. Erbad, "Decentralized Autonomous Organizations (DAOs) Adoption for Smart City Governance," in J. Prieto et al. (Eds.), Blockchain and Applications, 6th International Congress, BLOCKCHAIN 2024, Lecture Notes in Networks and Systems, vol. 1256, Springer, 2025.

[9] H. D. Jovita-Olvez et al., "Synergizing AI and blockchain: a bibliometric analysis of their potential for transforming e-governance in smart cities," Frontiers in Sustainable Cities, vol. 7, 2025. https://doi.org/10.3389/frsc.2025.1553816

[10] European Parliament. Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L 2024/1689, 2024.

[11] E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, S. Muralidharan, C. Murthy, B. Nguyen, M. Sethi, G. Singh, K. Smith, A. Sorniotti, C. Stathakopoulou, M. Vukolic, S. W. Cocco, and J. Yellick, "Hyperledger Fabric: A distributed operating system for permissioned blockchains," in Proc. EuroSys, 2018, pp. 1-15.

[12] V. Buterin, "Notes on blockchain governance," vitalik.ca, Dec. 2017. [Online]. Available: https://vitalik.ca/general/2017/12/17/voting.html

[13] European Parliament. Regulation (EU) 2016/679 of the European Parliament and of the Council (General Data Protection Regulation). Official Journal of the European Union, L 119, 2016.

[14] E. Ben-Sasson, A. Chiesa, C. Garman, M. Green, I. Miers, E. Tromer, and M. Virza, "Zerocash: Decentralized anonymous payments from Bitcoin," in Proc. IEEE S&P, 2014, pp. 459-474.

[15] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. U.S. Department of Commerce, Jan. 2023.

[16] P. Tsankov, A. Dan, D. Drachsler-Cohen, A. Gervais, F. Buenzli, and M. Vechev, "Securify: Practical security analysis of smart contracts," in Proc. ACM SIGSAC CCS, 2018, pp. 67-82.

[17] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, "Internet of Things for smart cities," IEEE Internet Things J., vol. 1, no. 1, pp. 22-32, Feb. 2014.

[18] G. Wood, "Ethereum: A secure decentralised generalised transaction ledger," Ethereum Project Yellow Paper, 2014.

[19] S. Gupta, AI, Blockchain, and Autonomous Innovation: Charting the Future of Intelligent Enterprises. Amazon, 2025.

[20] N. Testi and R. Marconi, "Exploring the potential of blockchain technology for citizen engagement in smart governance," Open Research Europe, vol. 3, p. 183, 2025. https://doi.org/10.12688/openreseurope.16153.4