Article

Not peer-reviewed version

Using Blockchain Ledgers to Record the AI Decisions in IoT

Vikram Kulothungan *

Posted Date: 23 April 2025

doi: 10.20944/preprints202504.1789.v1

Keywords: IoT; blockchain; auditing

 doi:10.20944/preprints202504.1789.v1

Article,

Using Blockchain Ledgers to Record the AI Decisions in IoT

Vikram Kulothungan

Capitol Technology University; vikramk1986@gmail.com

Abstract: The integration of artificial intelligence (AI) into Internet of Things (IoT) systems has outpaced the development of mechanisms to explain and audit automated decisions, creating a transparency gap. This paper addresses the research problem of establishing immutable audit trails for AI-driven IoT decisions to enhance trust, accountability, and regulatory compliance. We propose a blockchain-based framework that logs each AI inference and its provenance data (inputs, model parameters, and outputs) on a tamper-proof distributed ledger, ensuring every decision is traceable and auditable. The technical method- ology centers on a permissioned blockchain ledger deployed alongside IoT infrastructure. IoT devices and edge nodes commit decision records via smart contracts, producing an im- mutable, timestamped log resistant to manipulation. This approach leverages blockchain's decentralization and cryptographic integrity to guarantee non-repudiation and data integrity. We detail how the system design balances transparency with privacy (e.g. hashing personal data) to remain compliant with data protection norms. The solution aligns closely with emerging regulatory frameworks such as the EU AI Act's mandate for automated decision logs and traceability, and GDPR's accountability and transparency requirements (e.g. maintaining audit logs of AI decisions for explainability). We demonstrate the frame- work's applicability across domains: healthcare IoT, to log diagnostic AI recommendations for accountability; and industrial IoT, to track autonomous control actions - showing that our approach generalizes to diverse high-stakes environments. The paper's contributions include a novel architecture for AI decision provenance in IoT, a detailed implementation on a blockchain ledger to securely record AI decision-making processes, and an evaluation of its performance and compliance benefits. By providing a reliable, immutable audit trail for AI in IoT, this work enhances transparency and trust in autonomous systems and offers a timely solution for auditable AI in an era of increasing regulatory scrutiny.

Keywords: IoT; blockchain; auditing

1. Introduction

The convergence of Internet of Things (IoT) and Artificial Intelligence (AI) is trans-forming modern cyber-physical systems, from smart homes and grids to autonomous healthcare and industrial automation. IoT networks now involve billions of connected devices generating vast data streams, and projections anticipate hundreds of billions more devices in the near future [1]. AI techniques (especially machine learning) are increasingly deployed on this IoT data to enable autonomous decision-making and intelligent control in real-time. However, alongside these benefits arise significant challenges around trust, transparency, and accountability in AI-driven IoT systems. Stakeholders often lack visibility into how an AI arrived at a decision, making it difficult to trust automated actions in safety-critical or sensitive domains. Regulators and users are demanding greater explain-ability and traceability for AI decisions [2]. For instance, the European Union's General.

Blockchain technology has emerged as a promising tool to address these transparency and trust gaps. Blockchains are decentralized ledgers that could provide immutable and tamper-evident records of transactions or data, maintained by a distributed network of nodes. Prior work suggests that recording AI decision processes on a blockchain can enhance transparency and user

 doi:10.20944/preprints202504.1789.v1

trust. By leveraging cryptography and consensus mechanisms, blockchains ensure that once data (such as an AI decision or input) is recorded, it cannot be altered retroactively, thereby creating an immutable audit trail. Such blockchain-based audit trails can improve traceability of AI decision-making and make compliance auditing much simpler. Furthermore, blockchain's distributed nature can remove single points of failure in IoT data management [5] and enable multiple stakeholders (e.g., device owners, service providers, regulators) to share a common trusted record of AI-driven actions without relying on a central authority. This inherently supports a higher level of trust in autonomous IoT systems [6].

Despite growing interest in combining AI and blockchain, the specific idea of using blockchain ledgers to log the provenance of AI decisions in IoT remains an evolving and relatively under-explored area. Many existing IoT deployments lack comprehensive logging of AI behavior, and traditional logs (if they exist) are centralized and prone to tampering or loss. At the same time, much of the blockchain-for-IoT research has focused on device authentication, data integrity, or cryptocurrency applications, rather than detailed audit trails of algorithmic decision-making. Thus, there is a need to chart new approaches that integrate AI decision provenance with blockchain-based immutability to fulfill emerging transparency requirements. This paper aims to fill this gap by proposing a novel framework in which every significant AI inference or decision in an IoT system is recorded to a blockchain ledger, along with key metadata to explain and reproduce that decision. By doing so, we seek to ensure that a permanent, verifiable record exists for later accountability- whether for debugging model errors, investigating accidents, demonstrating compliance, or providing explanations to users.

The paper is organized as follows: Section 2 reviews related work and situates our research within the existing literature on blockchain, AI, and IoT convergence, identifying the open gaps. Section 3 details our proposed methodology and system architecture for blockchain-based AI decision provenance in IoT. Section 4 discusses the relevant regulatory and governance landscape and how our approach addresses key requirements for trust and transparency. Section 5 presents the conceptual framework and applies it to representative case studies to illustrate its usage. Section 6 provides a critical discussion of the approach, including potential challenges (scalability, privacy, ethics, security) and corresponding mitigation strategies. Finally, Section 7 concludes the paper, summarizing the contributions and highlighting directions for future research and development.

2. Literature Review

2.1. Blockchain and IoT Convergence

Blockchain technology has been widely explored as a means to strengthen security and trust in IoT systems. Traditional IoT architectures rely on centralized brokers or cloud servers to aggregate device data, which creates single points of failure and raises concerns about data integrity and control [5]. Blockchain, by contrast, offers a distributed ledger where data from IoT devices can be recorded in an immutable, append-only chain of blocks secured by cryptographic hashes. Researchers have shown that this approach can mitigate many IoT vulnerabilities by removing centralized control and making data tamper- evident [7]. For example, Dorri et al. (2017) proposed an architecture for smart home IoT where blockchain maintains an access control list and transaction log for device interactions, improving security and privacy [8]. In a similar vein, Banafa (2017) outlined the benefits and challenges of IoT-blockchain convergence, highlighting how decentralization can increase resilience and trust in IoT data exchanges [9]. A recent survey by Fotia et al. (2022) further discusses decentralized trust management for IoT using blockchain, emphasizing that smart contracts can authenticate devices and verify data integrity in IoT's multi-layered architecture [10]. Overall, the literature suggests that blockchain can provide security, transparency, and traceability to IoT networks, with promising applications in supply chain monitoring, smart cities, and sensor networks [2]. However, much of this work has focused on

 doi:10.20944/preprints202504.1789.v1

general data or transaction logs (e.g., logging sensor readings or device access events on-chain), rather than the specific context of AI or machine learning processes running on IoT data.

2.2. AI Decision-Making and the Need for Audit Trails

As AI algorithms are increasingly deployed in IoT environments (for tasks such as anomaly detection, predictive maintenance, or autonomous control), concerns have grown regarding the opacity of AI decisions. Literature on explainable AI (XAI) and algorithmic accountability has repeatedly pointed out that AI systems require traceability and auditability to be trustworthy [11]. Users and regulators are wary of "black-box" models whose inner workings are not transparent, especially in high-stakes applications like healthcare or autonomous vehicles. Several works have proposed frameworks for algorithmic auditing and logging. For instance, Ananny and Crawford (2018) discuss "algorithmic accountability" and the importance of audit mechanisms to understand AI outputs [12]. Sokol et al. (2020) present methods for generating human-understandable explanations from AI systems to accompany decisions [13]. However, ensuring these explanations and decisions are stored reliably over time remains a challenge. Traditional databases can log decisions, but those logs could be modified or deleted (intentionally or accidentally), undermining their usefulness in forensic analysis or compliance scenarios. This is where blockchain's properties become attractive – by creating immutable records of AI system behavior, one can guarantee an audit trail that is secure from tampering. A blog by IBM described how blockchain could maintain an auditable log of the data and evidence that led an AI model to a particular prediction (illustrated with a simple fruit classification example) [1]. Academic works such as "Blockchain as a platform for Artificial Intelligence Transparency" also suggest that transparent ledgers can help verify that AI models and their training data have not been tampered with [14]. These discussions provide a foundation, but concrete implementations in IoT settings are still limited. Our work builds on this nascent area by focusing specifically on AI decision provenance in IoT -capturing not only that a decision was made, but contextual data about how it was made (inputs, model ID, etc.) in a way that is resilient and sharable.

2.3. Data Provenance and Audit Trails via Blockchain

The concept of using blockchains for data provenance (i.e., tracking the origin and evolution of data) has been explored in various domains, which we draw inspiration from. Liang et al. (2017) introduced ProvChain, a blockchain-based data provenance architecture in cloud environments that embeds provenance metadata (who accessed/modified data and when) into blockchain transactions [15]. This allows cloud users to trace how their data is used in a multi-user environment with strong guarantees of integrity. ProvChain demonstrated that blockchain can serve as a robust audit log, resilient to deletion or forgery, and this concept can be analogously applied to IoT sensor data and AI outputs. In the IoT realm, researchers have proposed tailored provenance schemes: for example, Sharma (2023) describes a hierarchical blockchain-based data provenance system for IoT to address scalability, where multiple interconnected ledgers track provenance at different layers or domains of an IoT network [16]. Also related is the work by Zyskind et al. (2015) on using blockchain for personal data management, where individuals could log data access events on a ledger for transparency [17]. These efforts, while not AI-focused, show that blockchain is effective for audit trails in distributed systems. We leverage similar techniques (smart contracts, hashing of data, event logs on chain) but target them to recording AI decision events and their provenance. Another relevant area is blockchain in accountable logging systems. Regueiro et al. (2021) present a blockchain-based audit trail mechanism for general IT systems, implementing a prototype where an enterprise's system events are recorded on a private Ethereum-based chain [18]. Their approach uses smart contracts to log events and a blockchain monitor for auditors to query the log. This shows the feasibility of integrating blockchain logging into existing software workflows. Our literature review indicates a clear research gap: while blockchain-based IoT data integrity and

 doi:10.20944/preprints202504.1789.v1

blockchain-based AI governance have been studied separately, integrating blockchain to provide an immutable audit trail for AI-driven IoT decisions remains largely unexplored and warrants original research.

In summary, the literature confirms that: (a) IoT systems benefit from blockchain for secure data sharing and logging; (b) AI systems need better transparency and audit mechanisms; and (c) marrying blockchain with AI in IoT is a promising approach to achieve trustworthy and accountable autonomy. However, prior works have not fully realized an end-to-end solution that ties together IoT device data, AI decision-making, and blockchain logging. The novel perspective we pursue is to treat AI decisions in IoT as first-class transactions to be immutably recorded, akin to financial transactions in cryptocurrency but for "AI logic transactions". By doing so, our approach goes beyond existing discussions in originality - introducing a ledger of AI decisions as an infrastructural component of IoT systems. The next sections detail this approach and how it addresses the identified gaps.

3. Proposed Methodology

To realize immutable audit trails for AI decisions in IoT, we propose a comprehensive methodology that combines a carefully designed system architecture with blockchain mechanisms tailored to IoT constraints. Figure 1 provides an overview of the envisioned framework, depicting how various components (IoT devices, AI decision engine, blockchain network, and users/auditors) interact to create a trusted ecosystem. At a high level, our approach works as follows: as IoT devices generate data and AI algorithms produce decisions or predictions based on that data, every decision event (along with essential provenance information) is encapsulated into a transaction and written to a blockchain ledger. Smart contracts on the blockchain define the structure of these decision records and enforce rules (e.g., who can write or read the log). The blockchain network – composed of distributed nodes operated by stakeholders (manufacturers, service providers, perhaps even regulators) – reaches consensus on each new decision record, timestamping and immutably linking it in the chain. This produces a chronological, tamper-proof log of AI operations that can be queried or audited by authorized parties at any time.

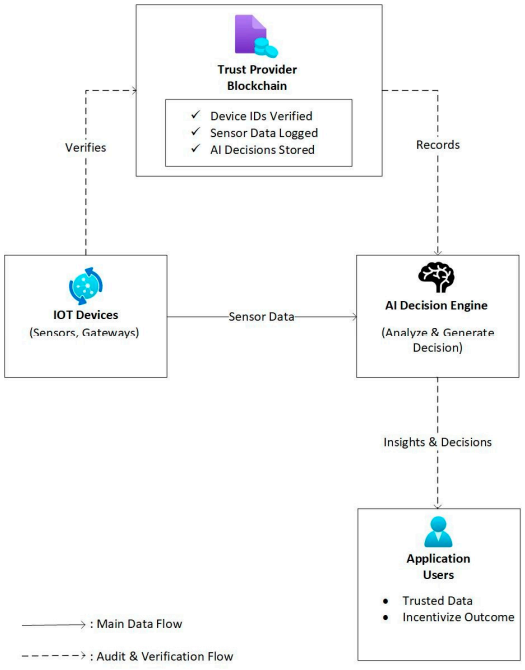

Figure 1. Conceptual framework for a blockchain-based trust and audit system.

 doi:10.20944/preprints202504.1789.v1

4. System Architecture

Our proposed architecture is structured in layers, aligning with typical IoT deployments, and augmented with blockchain and AI components at appropriate layers. We delineate four main layers: (L1) Data Generation, (L2) Data Aggregation, (L3) Data Services and Blockchain, and (L4) Application and Audit Interface. Figure 2 illustrates these layers and their key functionalities in the context of our framework. Figure 2: Layered architecture of the proposed blockchain-based AI audit trail system for IoT (inspired by the LoRaTRUST architecture [7])

4.1. L1: Data Generation (IoT Devices)

This bottom layer consists of IoT sensors, actuators, and embedded devices that generate raw data and also often execute local AI models (especially with the rise of edge AI). Each device in L1 is provisioned with a unique cryptographic identity (public-private key pair or blockchain address) which it uses to sign the data or decisions it produces. Lightweight software on the device (or associated gateway) handles sensor data retrieval (component C1.1) and may perform on-device AI inference if capable. Before sending data upward, the device can encrypt and/or sign the data (C1.4 Data Encrypter, C1.3 Data Signer) to ensure confidentiality and authenticity. For example, an IoT camera with an AI model for object detection would capture an image, run the detection locally, and then package the detected objects (decision) along with a hash of the image and a timestamp, sign it with its key, and send it as an authenticated record. The device registration with the blockchain occurs at setup: smart contracts on the blockchain maintain a verifiable registry of valid IoT nodes, recording each device's public key, type, and metadata [7]. Only registered, authenticated devices are allowed to submit decisions to the audit trail, preventing illegitimate data injection.

4.2. L2: Data Aggregation (Edge/Fog Gateways)

This layer includes IoT gateways, edge servers, or fog nodes that collect data/decisions from L1 devices (often over local networks like LoRaWAN, WiFi, etc.). Components in L2 aggregate and preprocess incoming information (C2.1 Data Router) and can enforce access control policies (C2.2 ACL Policy) to filter or control which data moves onward. In our framework, L2 also serves as the liaison to the blockchain network. Resource-constrained sensors may not run a blockchain client themselves, so a gateway can batch multiple device decisions into blockchain transactions or serve as a proxy node that uploads device-signed records to the blockchain. Importantly, the gateway also attaches its own identity and possibly additional context (e.g., location, network status) to the record, which is useful for provenance. We note that gateways too are registered on-chain (similar to devices) and their contributions can be audited [7]. By the time data reaches the end of L2, it has been authenticated, optionally encrypted, and is ready to be committed to the ledger.

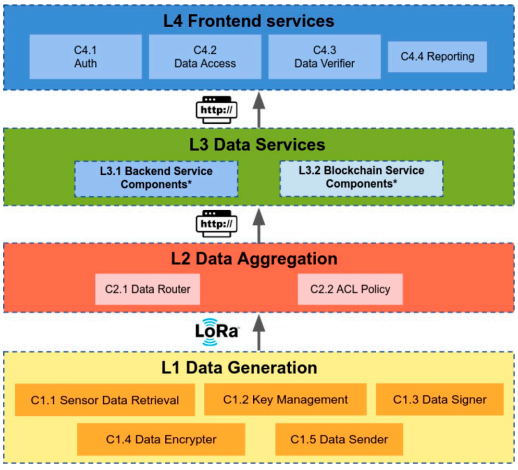

 doi:10.20944/preprints202504.1789.v1

Figure 2. Layered architecture overview. The diagram highlights the components from the IoT device layer to the Frontend services of LoRaTRUST.

4.3. L3: Data Services and Blockchain

This is the core layer where the blockchain ledger and related services reside. It encompasses two sub-components: (L3.1) Off-chain backend services and (L3.2) Blockchain services. The off-chain services (if used) might include databases or storage for large data (like raw sensor readings or images) that are referenced from the blockchain to save space. For instance, if an AI decision involves large input data, the raw data might be stored in secure cloud storage or IPFS, and only its hash and a link are put on-chain. The blockchain service components are the smart contracts and network nodes that form the distributed ledger. We implement a smart contract (AuditTrailContract) that defines a structure for AI decision records: e.g., deviceID, gatewayID, modelID, inputDataHash, outputDecision, confidence, timestamp. When an IoT gateway calls this contract to log a new decision, the contract validates the identities (only registered device/gateway combos can log; this is enforced using the registry contract) and then emits an event and stores the record on-chain [18]. We choose an appropriate blockchain platform here – for example, a permissioned blockchain (like Hyperledger Fabric or Quorum) is suitable to meet IoT performance needs and privacy, since participants are known entities. Using a permissioned chain means we avoid the expensive proof-of-work mining of public blockchains; instead, a lightweight consensus protocol (such as Practical Byzantine Fault Tolerance (PBFT) or Proof of Authority) is used among a set of validator nodes. This is crucial for IoT, as traditional proof-of-work is computationally and energetically prohibitive in this context. Our design aligns with prior research that calls for lightweight blockchain consensus for IoT to achieve low latency and low energy consumption [19]. In fact, one could implement a specialized consensus like Hierarchical Proof-of-Contribution (HPoC) which was designed for IoT settings to reduce node load, or simply use a fixed set of authority nodes (e.g., the IoT platform operator and a few independent auditors). The outcome in L3 is that all submitted AI decisions are validated and recorded on the blockchain ledger, each block containing one or more decision records. The blockchain ensures ordering (time-sequenced log) and immutability (earlier blocks cannot be altered without breaking the chain's cryptographic links). Additionally, we integrate a Blockchain Monitor API in this layer – a RESTful service that external applications or auditors in L4 can use to query the stored records without needing direct blockchain access (this concept is similar to the blockchain monitor in Regueiro et al.'s system [18]). The monitor listens to events from the AuditTrailContract and caches relevant info for user-friendly retrieval.

4.4. L4: Application and Audit Interface

The top layer represents the end-user applications, analytics dashboards, or audit tools that make use of the recorded AI decision provenance. This layer includes components like user authentication (C4.1 Auth), secure data access interfaces (C4.2 Data Access) for pulling relevant logs, data verifiers or analytics engines (C4.3 Data Verifier) that can analyze the blockchain data (for example, to detect patterns of bias or errors in decisions), and reporting tools (C4.4 Reporting) to present audit trail information to human auditors or system managers. A variety of stakeholders interact at L4: IoT application users (such as a utility company in a smart grid scenario, or a doctor in healthcare) can retrieve the sequence of decisions affecting them to understand system behavior. Regulators or third-party auditors can be granted access to inspect the logs (likely via the monitor API or by running their own blockchain node) to ensure compliance and accountability. We enforce that only authorized parties can read potentially sensitive data - while the blockchain ensures integrity, it doesn't inherently make all data public if we use a permissioned chain. Thus, access control is layered on via the Auth and Data Access services (for instance, using role-based permissions where a patient can see their health device's decision log, a doctor can see their patients' logs, an auditor can see anonymized aggregate logs, etc.). The

 doi:10.20944/preprints202504.1789.v1

Data Verifier component can cross-check the provenance: for example, it might fetch a decision's input hash from the blockchain and compare it to a recomputed hash of the input stored off-chain to ensure consistency (detecting any tampering in storage). It could also verify digital signatures on the decision to confirm it was indeed produced by the claimed device and model. These verification steps add an extra layer of trustworthiness and are automated via smart contracts and cryptographic protocols.

The synergy of these layers yields a system where AI decisions in IoT are born verifiable. From the instant a decision is made at the edge, it is signed and eventually anchored in an immutable ledger, along with information linking it to the raw data and the algorithm that produced it. This design addresses key requirements: Tamper-evidence (blockchain guarantees immutability), Real-time logging (edge gateways push decisions to the chain as they occur, with minimal delay), Provenance richness (smart contracts ensure metadata like device and model IDs are included, enabling later explanation), and Availability (the distributed ledger means even if some nodes fail, the audit trail is preserved across others).

5. Data Provenance and Smart Contract Design

A critical aspect of the methodology is how we represent and store the provenance of AI decisions on the blockchain. We define the provenance as all information necessary to trace the decision's origin and context. Concretely, for each AI inference made in the IoT system, we record: the unique ID of the device or sensor that provided input, the ID or version of the AI model/algorithm that made the decision (this could be a model hash or an index referencing a model registry), a summary or hash of the input data on which the decision was based, the decision output itself (e.g., a predicted label or control action), any confidence score or explanatory pointers (if available, such as an explanation vector or a link to an explanation artifact), the timestamp of the decision, and the identity of the entity that logged it (which might be the gateway or edge node). This information is structured as a tuple or object that the smart contract on the blockchain can store.

We implement a smart contract (in Ethereum solidity or an equivalent chain code in other platforms) called AuditTrailContract. Pseudocode for key parts of this contract:

```
    struct DecisionRecord {
    bytes32 deviceID;
    bytes32 modelID;
    bytes32 inputDataHash;
    string decisionOutput;
    uint256 timestamp;
    bytes32 loggedBy;  // gateway or node ID
}

mapping(uint256 -> DecisionRecord) public decisions;  // log storage
uint256 public decisionCount;

function logDecision(bytes32 deviceID, bytes32 modelID, bytes32 inputDataHash
                     string memory decisionOutput) public {
    require(isRegisteredDevice(deviceID), "Device not registered");
    require(isAuthorizedLogger(msg.sender, deviceID), "Not authorized");
    decisionCount += 1;
    decisions[decisionCount] = DecisionRecord(deviceID, modelID,
                                        inputDataHash, decisionOutput,
                                        block.timestamp, msg.sender);
    emit DecisionLogged(decisionCount, deviceID, modelID, inputDataHash,
                                        decisionOutput, block.timestamp,
                                        msg.sender);
}
```

 doi:10.20944/preprints202504.1789.v1

In the above, isRegisteredDevice checks the device registry contract to ensure the device is known (preventing fake device entries) [7], and isAuthorizedLogger ensures that the msg.sender (the blockchain account calling this function) is allowed to log for that device (for example, the gateway or the device itself if it has direct chain access). Each new decision gets an index decisionCount and is stored in the contract state as well as emitted as an event. Emitting an event (DecisionLogged) is useful for off-chain monitoring tools to catch the log without constantly querying the chain.

We assume that each IoT device or gateway holds credentials to authenticate to the blockchain (like an Ethereum account key). In a permissioned network, these might be tied to X.509 certificates or other identities. The registration of devices/gateways happens via an administrative process (either a privileged contract function called by the IoT platform admin at deployment time, or automatically when a device first joins using a secure join protocol). That registry contract essentially acts as a permission layer on who can write to the audit trail.

A design decision here is data storage vs. cost. Blockchain storage is typically expensive (in terms of throughput and, on public chains, fees). In permissioned chains, we don't have fees, but performance could still be an issue if we log very high-frequency events (like sensor readings every second). Our focus is on AI decisions, which usually occur less frequently than raw data generation (e.g., an anomaly detection model might raise an alert a few times a day, or a predictive model might make hourly predictions). Thus, the volume of transactions should be manageable. If needed, we can aggregate multiple decisions into one transaction (logging an array of DecisionRecords in one call) to reduce overhead. Alternatively, for bulk data provenance (e.g., every single sensor reading), it might be more efficient to store only hashes on chain (for integrity) and keep the detailed log off-chain. In our methodology, we are flexible: high-value events (like a critical decision or a violation) can be fully logged on-chain, whereas routine data can be logged as hashed references. This approach aligns with the idea of using off-chain storage + on-chain hash to balance load, a pattern used in systems like Tierion or Factom in data auditing.

Another element is the consensus protocol used by the blockchain network. As mentioned, IoT settings benefit from permissioned consensus. We might configure the blockchain nodes (which could be running on edge servers, cloud servers of the IoT service provider, and perhaps nodes run by third-party auditors) to use PBFT or a variant. PBFT can achieve fast finality (confirmation of transactions) in small networks with known nodes, which is ideal for near-real-time logging. There is precedent in literature for using PBFT-based private Ethereum networks for audit logging [18]. An alternative is Proof- of-Authority (PoA), where a few designated nodes validate all blocks. This is simpler to implement (e.g., Ethereum's Clique PoA or Polygon Edge for IoT) and would ensure low latency and energy usage – a non-financial, lightweight consensus as suggested by Nie et al. (2022) [19]. We do not use Proof-of-Work (to avoid heavy computation) and even Proof-of-Stake may be unnecessary overhead if the participants are pre-vetted; however, if a consortium of organizations runs the audit trail, a stake-based mechanism could be introduced to decentralize trust among them.

To make the methodology concrete, consider an IoT-based smart irrigation system [20], detailed in Table 1. Soil moisture sensors (devices) send readings to an edge AI model that decides whether water valves should open or remain closed. Each decision made by the AI triggers a logging event captured immutably on a blockchain ledger, enhancing transparency and auditability. As illustrated in Table 1, Sensor 123 detects moisture at 10 percent and transmits this measurement to Gateway A. Gateway A forwards the data to the AI model, which, due to a predefined threshold of 15 percent (specified by the IrrigationModelv2), determines the appropriate action: "valve=OPEN" for field X. Gateway A then prepares a transaction containing critical metadata—including the deviceID (123), modelID (IrrigationModelv2), a hash of the input data (0xabc123), and the decision outcome (OPEN valve X) - signs it, and submits it to the blockchain. Blockchain nodes verify the authenticity and permissions of Sensor 123 and Gateway A, confirm they are properly registered to field X, and subsequently store the record permanently. An event is

 doi:10.20944/preprints202504.1789.v1

emitted to indicate successful logging. Later, an auditor or the farmer can query, "Why was water released at time T?", retrieving the detailed transaction record from the blockchain. The hash (0xabc123) from the blockchain can then be cross-referenced with off-chain stored sensor data (stored on IPFS or Gateway A's database) to confirm the moisture reading was indeed 10 percent. Additionally, referencing the model registry, the farmer can validate that the irrigation threshold was set at 15 percent moisture, confirming the AI's decision was correct and justified. This structured process significantly enhances transparency: every automated action is verifiable and fully explainable through logged metadata, and blockchain's immutability ensures that neither farm operators nor device manufacturers can secretly modify decision logs.

Table 1. Steps in Soil Moisture Monitoring and Decision Process.

Step	Actor/Component	Action/Event	Data Involved
1	Soil Moisture Sensor	Measures soil moisture	Moisture = 10%
2	Soil Moisture Sensor	Sends measurement to Gateway A	Device ID (D123), Moisture = 10%
3	Gateway A	Forwards sensor data to AI model	Device ID (D123), Moisture = 10%, Field ID
4	AI Decision Model	Evaluates data against threshold	Threshold = 15%, Measured Moisture = 10%, Decision = "OPEN valve X"
5	Gateway A	Logs decision event to Blockchain	deviceID=d123, modelID=IrrigationModelž, inputDataHash=0xabc123, decisionOutput="OPEN valve X"
6	Blockchain Ledger	Validates device and gateway, stores decision event	Immutable decision record stored on-chain
7	Auditor or Farmer	Queries blockchain "Why valve opened at time T?"	Retrieves decision metadata
8	Off-chain Data Storage	Verifies stored moisture data matches input hash	Moisture=10%, inputDataHash=0xabc123
9	Auditor or Farmer	Confirms model details	IrrigationModelž, threshold=15%
10	Auditor or Farmer	Confirms transparency and correctness of decision	Decision validated

In implementing this methodology, we also consider performance optimizations. IoT systems may produce bursts of events. The smart contract could be optimized (e.g., pre- allocate array sizes, avoid writing unnecessary data, etc.). We can also utilize layer-2 scaling solutions if needed: for example, a local off-chain channel could collect decisions and periodically anchor a summary on the main chain (but since we aim for auditability, we tread carefully with off-chain channels as they introduce delay in finalizing logs). Given that our chosen examples (smart grid, healthcare, industrial logs) typically have manageable event rates (compared to, say, a high-frequency trading AI which might produce thousands of decisions a second), our direct on-chain logging approach is practical. To ensure data integrity, each record's critical fields (like input hash and output) are part of the transaction's hash that gets into the blockchain block. Thus, any attempt to change an output later would change the hash and break the chain, making tampering evident. Additionally, device signatures (where used) mean that even if a malicious gateway attempted to fabricate a decision from a device, it would not have the device's private key to sign it, and the

 doi:10.20944/preprints202504.1789.v1

contract could detect that. This cryptographic chain-of-custody from data generation to decision logging is at the heart of our provenance solution.

We have mentioned logging modelIDs - this implies there is a way to identify which AI model made the decision. In practice, AI models in IoT could be periodically updated (new firmware, retrained model). We incorporate a simple model management scheme: whenever a model is deployed to devices or edge nodes, a hash of the model binary or its version number is registered on blockchain (this could be another contract or part of the device registry). Then the device includes that model identifier when logging decisions. This creates an immutable record of which model version was responsible for a decision, aiding in accountability (e.g., if later a model is found to be faulty, all decisions made by that version can be traced and reviewed). This idea aligns with emerging practices in AI governance where model provenance is tracked (sometimes called "model lineage"). By utilizing the same blockchain for model registration (or an interoperable one), we ensure model provenance and decision provenance are linked.

In summary, our methodology sets up a secure pipeline: IoT data and AI decisions flow upward through layers, accruing signatures and context, and end up cemented in a blockchain ledger via smart contracts that enforce correctness of origin. The design addresses IoT constraints by using lightweight cryptography and consensus, and ensures that the resulting audit trail is rich enough to meet explainability and compliance needs.

6. Regulatory and Governance Considerations

The technical approach described above is deeply influenced by and aligned with evolving regulatory frameworks and governance principles for AI and IoT systems. In this section, we discuss how our blockchain-based audit trail framework addresses key legal and ethical requirements, focusing on the European context (which is at the forefront of AI and IoT regulation) such as the EU AI Act, GDPR, and the Cyber Resilience Act, as well as broader principles of trust, transparency, and accountability in AI. We also highlight any compliance challenges and how the framework is designed to navigate them.

6.1. Regulatory Alignment

The following table details the regulatory requirements of key EU frameworks and explicitly maps how the proposed blockchain-based audit trail system meets these requirements. It also highlights potential compliance challenges and identifies strategies to mitigate these challenges.

Table 2. Blockchain-Based Framework: Regulatory Requirements, Challenges, and Mitigation Strategies.

Regulation	Key Requirements	Blockchain-Based Framework Implementation	Compliance Challenges	Mitigation Strategies
EU Artificial Intelligence Act (AI Act)	• Automatic event logging (Article 12) [21] • Transparency and explainability [22] • Post-market monitoring capability [23]	• Immutable and automatic logging with timestamps and decision contexts • Model identifiers and decision explainability • Auditable records	Permanent blockchain logs might exceed retention requirements	Flexible design to manage evolving requirements

 doi:10.20944/preprints202504.1789.v1

		suitable for regulatory inspection		
General Data Protection Regulation (GDPR) [24]	• Transparency (Ar- ticles 13-15) • Right to explanation • Data integrity and confidentiality (Articles 5, 17, 32) • Right to erasure (Article 17)	• Detailed decision logging supports transparency • Integrity ensured through hashing and im- mutable records maintained via permissioned access control	Blockchain immutability con- flicts with era- sure rights	Store pseudonymous hashes on-chain, personal data off-chain
EU Cyber Resilience Act (CRA) [25]	• Secure-by-design principles • Lifecycle cybersecurity measures • Incident and vulnerability management	• Strong cryptographic authenti- cation for devices • Immutable security event logging • Verifiable updates for secure lifecycle manage- ment	Ensuring timely and accurate logging of all security events	Blockchain-based immutable records facilitate monitoring and compliance

433

6.2. Ethical and Governance Principles

This table contextualizes the broader ethical and governance principles of trust, transparency, and accountability and demonstrates how the blockchain framework supports these principles through tangible system capabilities and practical impacts [2].

Table 3. Framework Alignment with Trust, Transparency, and Accountability Principles.

Principle	Framework Capability	Practical Impact
Trust	Immutable audit trails of AI de- cisions	Reliable, predictable systems and auditable behaviors

 doi:10.20944/preprints202504.1789.v1

Transparency	Real-time visibility into AI decisions	Enhanced stakeholder confidence and proactive engagement
Accountability	Clear attribution and logging of decisions	Simplified regulatory oversight, improved incident management and dispute resolution

437

In conclusion, our framework is a proactive response to regulatory trends. It is built to future-proof IoT AI systems against upcoming compliance demands by providing infrastructure for transparency and audit. It transforms compliance from a paperwork exercise into an automated technical feature. While some challenges (like GDPR data deletion and defining access policies) require careful implementation, the net effect is that an operator adopting this system is better prepared to meet legal obligations and to build user trust. By baking in accountability at the design phase, we avoid costly retrofits later or the risk of non-compliance. As regulations like the AI Act come into force, solutions like ours could be an enabler for companies to continue using powerful AI in IoT (which they want for efficiency) without falling afoul of the law or public trust. It's a win-win since the technology both satisfies governance needs and improves the system's reliability and trustworthiness for its own sake.

7. Framework and Case Studies

To demonstrate the applicability of our proposed audit trail framework, we outline a general blockchain-based technical framework and then delve into two representative case studies healthcare IoT, and industrial IoT. These scenarios were chosen because they involve critical decisions by AI, have multiple stakeholders who need trust, and are subject to regulatory oversight, making them ideal proving grounds for our approach. In each case study, we describe how the framework would be instantiated, and discuss the benefits and considerations observed.

7.1. Generalized Framework for Blockchain-Based AI Audit Trails in IoT

Across different IoT domains, the core elements of our framework remain consistent. These can be summarized as:

7.1.1. Decentralized Identity and Registry

Every IoT device and AI model is registered on a blockchain ledger with a unique identity (public key or address) [26]. This creates a root of trust-ensuring that any logged decision can be tied back to a known entity. In practice, an initial smart contract serves as the registry (as described in Section 4) where an admin or an enrollment protocol adds new device IDs (with metadata like owner, type, permissions) and model hashes. Table 1 illustrated this concept with a verifiable registry acting as a trust provider to assure data originates from legitimate sources.

7.1.2. On-Device Logging Trigger

IoT devices or edge gateways are configured to trigger a log event whenever an AI decision or other significant event occurs [27]. The trigger could be implemented in software (e.g., a callback in the AI inference code that calls the blockchain logging function) or via an IoT platform rule engine. The point is to make logging an intrinsic part of the decision workflow, not an afterthought.

 doi:10.20944/preprints202504.1789.v1

7.1.3. Blockchain Network and Smart Contracts

The blockchain network (permissioned) is set up connecting relevant parties. Smart contracts (such as the AuditTrailContract) reside on this network to accept decision records [28]. Additional contracts might be present for specific functionality—e.g., a contract to manage incentives or payments, or a contract that automatically flags certain events. Because it's domain-agnostic, the same fundamental contracts could be reused across industries, with possible customization via configuration. For instance, the structure of a DecisionRecord might have optional fields that different deployments fill differently (healthcare might include "patientID" as a pseudonym, etc.).

7.1.4. Access Control and Data Privacy Layer

Each framework deployment defines who can read or query the logs. Typically, writing to the log is restricted to the IoT devices/gateways (through the registry), but reading might be open to all permissioned nodes or limited by role [28]. Fine-grained access control can be achieved by storing sensitive data encrypted on-chain.

7.1.5. Audit and Analytics Tools

On top of the blockchain, we have tools (could be smart contracts or off-chain services) to analyze the logged data [29]. In general, an interface is provided to filter and retrieve logs (e.g., "get all decisions by device X in last 24h" or "find all instances where model Y gave output Z"). These tools might integrate with existing dashboards in that domain (e.g., a patient health record system) so that users may not even know a blockchain is behind the scenes – they just see a log of actions with an extra seal of trust (maybe an icon indicating the log entry is verified and immutable). Using this general framework, we now tailor it to two scenarios to highlight unique considerations and advantages in each.

7.2. Case Study 1: Healthcare IoT (Smart Healthcare and Medical Devices)

7.2.1. Scenario

In healthcare, IoT devices like wearables, implants, or monitoring sensors collect patient data (heart rate, blood glucose, etc.), and AI algorithms assist in diagnosis or triggering alerts [30]. Examples include a smart insulin pump that uses AI to adjust dosage, a remote patient monitoring system that alerts doctors if vitals go out of range, or an AI diagnostic tool analyzing medical IoT sensor data (like an ECG patch streaming heart signals analyzed by an AI for arrhythmias). The decisions here are often life-critical and also sensitive from a privacy standpoint [31]. There is heavy regulation (e.g., FDA in US, MDR in EU) for medical devices, and liability concerns if something goes wrong.

7.2.2. Application of Framework

Each medical IoT device (e.g., a wearable ECG monitor) is registered on a permissioned blockchain shared by, say, the hospital and device manufacturer. The AI model that interprets ECGs is also registered (modelID perhaps corresponds to a certified algorithm version). Whenever the device's AI flags an "arrhythmia detected" event, it logs it on blockchain with details: deviceID, modelID, inputDataHash, decision: arrhythmia alert, timestamp. The treating cardiologist, who has access to the system, can later see this log and also see any subsequent decisions (maybe "alert dismissed by patient" or "alert forwarded to ER"). If a patient suffers an adverse event, investigators can examine whether the AI gave a timely alert and whether it was acted upon. This is crucial for malpractice or product liability cases: the manufacturer could prove "our device did alert at time X as logged immutably, so if action was delayed, it was not due to device failure." Conversely, if the device failed to alert and it should have, the log (or absence thereof) provides evidence.

 doi:10.20944/preprints202504.1789.v1

7.2.3. Benefits in Healthcare IoT

The immutable audit trail can literally save lives by ensuring that no critical alert is lost or tampered with. In traditional systems, a faulty device might not log an event at all (so no one knows it missed an alert). With our framework, even the act of logging (or the lack thereof within expected conditions) is noticeable. It imposes a kind of self-auditing on the device which if an expected periodic heartbeat or report is missing from the ledger, the system can automatically raise a flag (since blockchain events can be monitored in real-time). This could detect device malfunctions more quickly. For patients, the transparency can improve trust in AI-assisted care. Patients are understandably cautious about AI diagnosing them. Knowing that every AI decision is recorded and can be reviewed by humans may reassure them that the AI is not operating in a black box void of oversight. It also helps doctors trust AI outputs more, since they can retrospectively examine what data led to what diagnosis and if needed, contest it with evidence. Over time, these logs can provide valuable data for improving algorithms or for regulatory audits to re-certify devices.

Privacy is a big concern here. We take measures like hashing personal data on-chain and restricting access. Likely the blockchain network is confined to medical professionals bound by confidentiality and perhaps patient representatives. GDPR and health privacy laws (like HIPAA in the US) would necessitate strict access control. However, since the primary goal is audit, not broad sharing, we can tailor it such that only authorized auditors (like a hospital's internal review board or an external regulator in case of incidents) can view patient-specific logs. Patients themselves might be given access to their own data on a patient portal which could present the data from the blockchain (effectively giving them an immutable copy of their device history, which some patients may appreciate for personal record-keeping or for second opinions).

Another advantage is compliance with medical device regulations [32]. Our system can largely automate that – each adverse event detected by the AI is logged immutably, which can feed into the manufacturer's required reporting to authorities. If a recall or safety notice is needed, the manufacturer can query the ledger to find all affected decisions or instances. Additionally, if an AI model is updated (perhaps to fix a bug that caused misdiagnosis), the effect of the update can be tracked by comparing logged decisions before and after across patients.

7.3. Case Study 2: Industrial IoT (Manufacturing and Supply Chain)

7.3.1. Scenario

In an industrial IoT setting (often dubbed Industry 4.0), factories are equipped with sensors and AI-driven control systems for automation. Examples: A predictive maintenance AI monitors vibrations in a machine via IoT sensors [33] and decides when to schedule maintenance; a robotics controller AI decides speed or action adjustments on an assembly line based on sensor inputs; or in supply chains, an AI might reroute shipments or adjust inventory based on IoT tracker data. Here, AI decisions directly impact physical processes, product quality, and safety. If a machine fails or a product defect occurs, companies need to trace back what went wrong. Also, these environments value uptime and efficiency, so quick diagnostics via logs are valuable. Moreover, there can be compliance needs (OSHA safety, ISO standards for quality) that require maintaining logs of operations.

7.3.2. Application of Framework

Each critical machine in a factory has IoT sensors and perhaps an edge controller running AI. All these controllers log their key decisions to a blockchain that is maintained within the factory (and possibly accessible to machine vendors or auditors). For instance, if an AI controlling a chemical process adjusts a valve, it logs: device: ValveController42, model: ChemProcessAIv1.2, input: temp+pressure hash, decision: valve 10 percent open, timestamp. If later an out-of-spec

 doi:10.20944/preprints202504.1789.v1

batch is found, the quality engineers can review logs and see that at a certain time, the valve was opened incorrectly by the AI (or conversely, everything was normal, pointing to a raw material issue). This provides traceability in production – akin to a black box recorder for manufacturing lines. In supply chain, imagine IoT trackers on cargo and an AI that decides to redirect a shipment due to predicted delays (like an AI logistic system). The decision "reroute container 123 via Route B" would be logged and shared among parties (supplier, carrier, buyer). This ensures no disputes like "we didn't authorize that reroute" – the log would show which AI or party did it, under what conditions.

7.3.3. Benefits in Industrial IoT

The blockchain audit trail strongly enhances traceability and quality control. Many industries require a traceability chain for products (from raw materials to final product) especially in sectors like automotive or aerospace. By augmenting the traditional traceability with AI decision logs, one can not only trace which machine processed a part, but also what decision that machine's AI made during processing [34] [35]. This can reveal deeper insights; for example, noticing that an AI vision system rejected 5 percent of parts as defective might indicate an upstream issue if logged properly. Immutable logs also protect against any internal malpractice – e.g., if someone tries to alter maintenance records to hide negligence, the blockchain record would contradict them. This fosters a culture of accountability in operations. From a safety perspective, consider a factory accident: investigators will look at machine logs. With our system, they get reliable logs that can't be doctored (unlike some cases where companies have been found to falsify logs after the fact to avoid liability). Knowing this immutability is in place may also act as a deterrent against cutting corners, because the record will endure. In terms of performance and scalability, factory networks are usually LAN-scale, and blockchains can be run efficiently. Modern factories already have historians (databases logging sensor data). We're complementing those with a tamper-proof layer for decisions. The volume of logged AI decisions might be much smaller than raw sensor data, so it's feasible to record them all on-chain within a factory consortium (perhaps a consortium including the manufacturer and equipment suppliers, so that both can see the logs – helpful if, for example, a machine under warranty fails, both the factory and the manufacturer have the same log evidence).

In the case studies above, the common thread is that the blockchain-based audit trail increases trust among stakeholders, whether they are consumers, patients, operators, or regulators, by providing a shared, indelible history of what the AI did. It also fulfills a documentary need that is otherwise cumbersome – replacing or supplementing manual logbooks or centralized databases that could be manipulated. The case studies show that, regardless of domain, the approach is flexible and valuable.

8. Discussion and Analysis

Implementing immutable blockchain-based audit trails for AI decisions in IoT systems offers clear benefits in terms of transparency and trust, as demonstrated in the case studies. However, it also raises important challenges and considerations. In Table 4, we critically examine the key issues on scalability and performance of the system, privacy concerns and data management, ethical implications of increased transparency, and security aspects, and strategies to mitigate these challenges.

 doi:10.20944/preprints202504.1789.v1

Table 4. Blockchain Implementation Challenges and Mitigation Strategies.

Challenge Area	Issue and Considerations	Mitigation Strategies
Scalability and Performance [19]	Blockchain logging can strain system throughput, storage, and latency.	• Batching multiple events per transaction. • Hierarchical or sharded ledgers. • Edge processing and data compression. • Alternative scalable DLTs (e.g., DAG-based systems).
Privacy and Data Management [2]	Immutable logging raises pri- vacy and GDPR compliance con- cerns.	• Off-chain storage for sensitive data. • Pseudonymous on-chain hashes. • Encryption with user-controlled keys. • Restricted access via permis- sioned blockchains.
Ethical and Social Implications [2]	Transparency might lead to surveillance concerns or misuse.	• Defined governance policies on log usage. • User consent and selective data sharing. • Clear communication on log purpose and ethical guidelines.
Security Considerations [7]	Vulnerabilities in devices, con- sensus mechanisms, and smart contracts.	• Secure hardware and IoT PKI for device authenticity. • Fault-tolerant consensus pro- tocols. • Security audits for smart con- tracts. • Rate limiting and anomaly de- tection to prevent misuse.

9. Future Research Directions

The discussion above reveals several areas where further research and development are warranted:

9.1. Automated Explanation Logging

We log decisions and some context, but an interesting extension is logging explanations generated by explainable AI techniques. Research could explore how to efficiently store why the AI made a decision (e.g., feature importances) on-chain without bloating the log. This would add

 doi:10.20944/preprints202504.1789.v1

a richer layer of provenance, truly capturing decision rationale, but needs summarization techniques to keep data small. There's a trade-off between log detail and size.

9.2. Integration with AI Monitoring Tools

We can integrate our blockchain logs with AI performance monitoring. For example, combine it with systems that detect model drift or bias over time. The blockchain log could feed an AI that scans for bias patterns (meta-AI auditing the AI via the ledger). This synergy of AI and blockchain for continuous auditing is a promising area

9.3. Policy Compliance Smart Contracts

As laws evolve, smart contracts could be made more intelligent to enforce compliance. Imagine a contract that refuses to accept a decision log if it detects that required data (say a justification for a high-risk decision) is missing. Essentially, encode some legal rules into the logging process. This is tricky but aligns with the concept of regulatory technology (RegTech) where systems ensure operations follow rules. It might involve oracles that know the current regulation context.

9.4. Scalability via Layer 2

Investigating advanced layer-2 solutions like state channels or rollups in the context of IoT. Perhaps IoT devices could log to a state channel (fast, off main chain) and then occasionally settle to main chain. This is similar to the hierarchical approach we mentioned and could significantly boost throughput if needed.

9.5. Cross-Domain Audit Trails

Possibly linking audit trails across domains. E.g., a car's IoT AI logs might connect to a smart city's traffic management logs when they interact. How to merge or cross-reference logs from distinct blockchains when an event involves multiple domains is an open question. Interoperability protocols or standards for "audit trail data format" could be beneficial.

9.6. User Interfaces for Audit Data

Researching effective ways to present blockchain audit data to non-technical stake-holders. The raw logs are technical; turning them into human-readable reports, integrating them with existing incident management workflows, is crucial for adoption. This is less about blockchain and more about HCI and design, but important so that the existence of logs actually leads to meaningful human oversight.

10. Conclusion and Future Directions

In this paper, we presented a comprehensive approach to establishing immutable audit trails for AI decisions in IoT systems using blockchain technology. Our work was motivated by the growing need for trust, transparency, and accountability in AI-driven IoT deployments – a need underscored by both practical considerations (e.g., diagnosing system failures, ensuring fair and ethical AI behavior) and emerging regulatory frameworks (such as the EU AI Act's logging requirements and other governance mandates). By combining IoT, AI, and blockchain, we aimed to create a synergistic framework where each technology's strengths compensate for the others' weaknesses: IoT provides data and action, AI provides intelligent decision-making, and blockchain provides a trustworthy record of those decisions.

We proposed a novel framework that is, to our knowledge, one of the first to explicitly intertwine blockchain logging with AI decision-making in IoT. This goes beyond existing literature by not just securing IoT data or enabling AI with blockchain, but specifically focusing on the provenance of AI decisions. We adapted the layered system components (from device level

 doi:10.20944/preprints202504.1789.v1

to application level) and the use of smart contracts to implement an audit trail. The architecture (illustrated in our figures and descriptions) shows how data flows from IoT devices through edge processing to a blockchain ledger, and how that ledger can be accessed by auditors and stakeholders. We also detailed the structure of log records and how consensus and smart contracts are tailored for IoT constraints, providing a blueprint for implementation.

The implications of this work are significant for both industry practitioners and the research community. For practitioners (IoT system architects, AI engineers, compliance officers), our framework provides a pathway to build systems that are "accountability- ready". Instead of retrofitting audit mechanisms after deployment (which often happens today in response to an incident or regulation), our approach advocates designing with auditability in mind from the start. This can drastically reduce compliance costs and improve stakeholder confidence. It encourages interdisciplinary collaboration: between blockchain experts (to optimize the ledger for these use-cases), AI experts (to determine what aspects of AI decisions should be logged to be most useful), and legal scholars (to refine how such logs can serve as evidence or compliance artifacts). We also anticipate that as AI continues to permeate IoT, the need for trust mechanisms will only grow.

In closing, our research demonstrates that blockchain-ledgered audit trails are a powerful enabler for trustworthy AI in IoT, converting the often opaque AI decision- making process into a transparent and verifiable sequence of events. We have shown it is feasible and beneficial, and we provided a structured way to implement it, supported by academic and practical references. As IoT systems become ever more complex and autonomous, we believe that approaches like ours will become not just advantageous but essential – akin to how financial accounting systems are mandatory for businesses, algorithmic accounting systems may become mandatory for AI. By anticipating that future, this work lays down foundational concepts and solutions. We encourage stakeholders in both academia and industry to build upon and deploy these ideas, fostering a future where IoT and AI systems are not only smart and connected but also transparent, accountable, and worthy of the trust we place in them.

References

1. IBM. How Blockchain Adds Trust to AI and IoT, 2020. Accessed: 2025-04-05.
2. Bhumichai, D.; Smiliotopoulos, C.; Benton, R.; Kambourakis, G.; Damopoulos, D. The Con- vergence of Artificial Intelligence and Blockchain: The State of Play and the Road Ahead. *Information* **2024**, *15*. https://doi.org/10.3390/info15050268.
3. Falletti, E. Automated Decisions and Article No. 22 GDPR of the European Union: An Analysis of the Right to an 'Explanation', 2019. Accessed: 2025-04-05.
4. European Union. Article 19: Automatically Generated Logs, 2024. Accessed: 2025-04-05.Schiller, E.; Esati, E.; Stiller, B. IoT-Based Access Management Supported by AI and Blockchains. *Electronics* **2022**, *11*. https://doi.org/10.3390/electronics11182971.
5. Alharbi, S.; Attiah, A.; Alghazzawi, D. Integrating Blockchain with Artificial Intelligence to
6. Secure IoT Networks: Future Trends. *Sustainability* **2022**, *14*. https://doi.org/10.3390/su142316002.
7. Vilchez, P.; Jacques, S.; Freitag, F.; Meseguer, R. LoRaTRUST: Blockchain-Enabled Trust and Accountability Service for IoT Data. *Electronics* **2023**, *12*. https://doi.org/10.3390/electronics1 2091996.
8. Dorri, A.; Kanhere, S.S.; Jurdak, R.; Gauravaram, P. Blockchain for IoT security and privacy: The case study of a smart home. In Proceedings of the 2017 IEEE international conference on pervasive computing and communications workshops (PerCom workshops). IEEE, 2017, pp. 618–623.
9. Banafa, A. IoT and Blockchain Convergence: Benefits and Challenges, 2017. Accessed: 2025-04-05.
10. Fotia, L.; Delicato, F.; Fortino, G. Trust in edge-based internet of things architectures: state of the art and research challenges. *ACM Computing Surveys* **2023**, *55*, 1–34.
11. Marr, B. Artificial Intelligence and Blockchain: 3 Major Benefits of Combining These Two Mega Trends, 2018. Accessed: 2025-04-05.

 doi:10.20944/preprints202504.1789.v1

12. Ananny, M.; Crawford, K. Seeing without Knowing: Limitations of the Transparency Ideal and Its Application to Algorithmic Accountability. *New Media & Society* **2018**, *20*, 973–989. https://doi.org/10.1177/1461444816676645.
13. Sokol, K.; Flach, P. One Explanation Does Not Fit All. *KI - Künstliche Intelligenz* **2020**, *34*, 235–250. https://doi.org/10.1007/s13218-020-00637-y.
14. Akther, A.; Arobee, A.; Adnan, A.A.; Auyon, O.; Islam, A.J.; Akter, F. Blockchain As a Platform for Artificial Intelligence (AI) Transparency, 2025, [arXiv:cs.CR/2503.08699]. Accessed: 2025-04-05.
15. Liang, X.; Shetty, S.; Tosh, D.; Kamhoua, C.; Kwiat, K.; Njilla, L. ProvChain: A Blockchain-Based Data Provenance Architecture in Cloud Environment with Enhanced Privacy and Availability. In Proceedings of the 2017 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID), 2017, pp. 468–477. https://doi.org/10.1109/CCGRID.2017.8.
16. Sharma, M. Hierarchical blockchain-based data provenance in IoT. In Proceedings of the AIP Conference Proceedings. AIP Publishing, 2023, Vol. 2916.
17. Zyskind, G.; Nathan, O.; Pentland, A.S. Decentralizing Privacy: Using Blockchain to Protect Personal Data, 2015. Accessed: 2025-04-05.
18. Regueiro, C.; Seco, I.; Gutiérrez-Agüero, I.; Urquizu, B.; Mansell, J. A Blockchain-Based Audit Trail Mechanism: Design and Implementation. *Algorithms* **2021**, *14*. https://doi.org/10.3390/a14120341.
19. Nie, Z.; Zhang, M.; Lu, Y. HPoC: A Lightweight Blockchain Consensus Design for the IoT. *Applied Sciences* **2022**, *12*. https://doi.org/10.3390/app122412866.
20. Pereira, G.; Chaari, M.Z.; Daroge, F. IoT-Enabled Smart Drip Irrigation System Using ESP32. *IoT* **2023**, *4*. https://doi.org/10.3390/iot4030012.
21. Artificial Intelligence Act, Article 12: Record-Keeping. https://artificialintelligenceact.eu/article/12/, 2024. Accessed April 17, 2025.
22. Artificial Intelligence Act, Article 13: Transparency and Provision of Information to Deployers. https://artificialintelligenceact.eu/article/13/, 2024. Accessed April 17, 2025.
23. Artificial Intelligence Act, Chapter IX: Post-Market Monitoring, Information Sharing and Market Surveillance. https://artificialintelligenceact.eu/chapter/9/, 2024. Accessed April 17, 2025.
24. General Data Protection Regulation (GDPR) – Legal Text. https://gdpr-info.eu/, 2016. Accessed April 17, 2025.
25. Cyber Resilience Act – Full Text of Articles. https://www.european-cyber-resilience-act.com/Cyber_Resilience_Act_Articles.html, 2024. Accessed April 17, 2025.
26. Javed, I.T.; Alharbi, F.; Bellaj, B.; Margaria, T.; Crespi, N.; Qureshi, K.N. Health-ID: A Blockchain-Based Decentralized Identity Management for Remote Healthcare. *Healthcare* **2021**, *9*. https://doi.org/10.3390/healthcare9060712.
27. Gong, L.; Alghazzawi, D.M.; Cheng, L. BCoT Sentry: A Blockchain-Based Identity Authentication Framework for IoT Devices. *Information* **2021**, *12*. https://doi.org/10.3390/info12050203.
28. Alharbi, A. Applying Access Control Enabled Blockchain (ACE-BC) Framework to Manage Data Security in the CIS System. *Sensors* **2023**, *23*. https://doi.org/10.3390/s23063020.
29. Wang, S.; Zhang, Y.; Guo, Y. A Blockchain-Empowered Arbitrable Multimedia Data Auditing Scheme in IoT Cloud Computing. *Mathematics* **2022**, *10*. https://doi.org/10.3390/math10061005.
30. Shukla, M.; Lin, J.; Seneviratne, O. BlockIoT: Blockchain-based Health Data Integration using IoT Devices. *AMIA Annu. Symp. Proc.* **2022**, *2021*, 1119–1128.
31. Vargas, C.; Mira da Silva, M. Case Studies about Smart Contracts in Healthcare. *Digit. Health* **2023**, *9*, 20552076231203571. https://doi.org/10.1177/20552076231203571.
32. U.S. Food and Drug Administration. Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations; Draft Guidance for Industry and Food and Drug Administration Staff; Availability. https://www.federalregister.gov/documents/2025/01/07/2024-31543/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing, 2025. Accessed April 17, 2025.
33. Godbole, R. Blockchain-Enabled AI for Predictive Maintenance in Industrial IoT. *Int. J. Holist. Manag. Perspect.* **2023**, *4*. Accessed April 17, 2025.

 doi:10.20944/preprints202504.1789.v1

34. Ayobami, A. How Blockchain Technology is Revolutionizing Audit and Control in Informa- tion Systems. https://www.isaca.org/resources/news-and-trends/industry-news/2024/how- 781 blockchain-technology-is-revolutionizing-audit-and-control-in-information-systems, 2024. Accessed April 17, 2025.
35. How Walmart Brought Unprecedented Transparency to the Food Supply Chain with Hyper-ledger Fabric. https://8112310.fs1.hubspotusercontent-na1.net/hubfs/8112310/Hyperledger/Printables/Hyperledger_CaseStudy_Walmart_Printable_V4.pdf, 2019. Accessed April 17, 2025.

 doi:10.20944/preprints202504.1789.v1

 doi:10.20944/preprints202504.1789.v1