Zero-Knowledge Agricultural Intelligence: Combining TinyML, Federated Learning, and Blockchain for Future Food Systems

Jian-Wei Liu[1], Fang-Lin Shen[2], Qing-Hua Zhou[3,] *

[1] School of Computer Science and Technology, Zhongyuan University of Technology, Zhengzhou 450007, Henan, China

[2] Department of Software Engineering, Liaoning University of Science and Technology, Anshan 114044, Liaoning, China

[3] College of Information Engineering, Guizhou Minzu University, Guiyang 550025, Guizhou, China

Corresponding author: qhzhou@gzmu.edu.cn

Abstract

The convergence of resource-constrained machine learning, privacy-preserving distributed computation, and decentralized governance technologies offers a transformative pathway for automating food safety compliance monitoring across heterogeneous and geographically dispersed agricultural settings. This paper introduces a unified framework that integrates TinyML-based edge inference, clustered federated learning with graph-attention mechanisms, and blockchain smart contracts secured by Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge (ZK-SNARKs). Deployed on low-power microcontrollers, the edge intelligence layer performs real-time classification of livestock health, hygiene indicators, and milk quality from multimodal sensor streams without transmitting raw data off-farm. Federated learning with density-based outlier filtering enables collaborative model improvement across heterogeneous farm environments while maintaining data locality. Compliance inferences are cryptographically committed through ZK-SNARK circuits and verified on a Layer-2 Polygon zkEVM blockchain, where smart contracts automate license management, audit triggering, and reputation scoring. Experimental evaluation demonstrates 88.1% on-device inference accuracy at 33.8 ms latency with 64.6 KB RAM usage, a 97.7% reduction in communication payload versus standard TFLite Micro implementations, and sustained accuracy above 90% under 20% sensor noise injection. The framework provides a scalable, privacy-preserving, and legally defensible instrument for automated regulatory compliance applicable to smallholder agricultural operations.

Keywords: TinyML; Federated Learning; Blockchain; Zero-Knowledge Proof; ZK-SNARK; Agricultural IoT; Food Safety Compliance; Smart Contract; Graph Attention Network; Privacy-Preserving Machine Learning

Article History:

Received: April 10, 2025

Revised: June 15, 2025

Accepted: August 25, 2025

Available Online: September 30, 2025

I. INTRODUCTION

The global challenge of feeding a projected 9.8 billion people by 2050 while simultaneously meeting increasingly stringent food safety standards demands a fundamental rethinking of how agricultural compliance is monitored and enforced [FAO, 2023; Wolfert et al., 2017]. Traditional inspection regimes, which depend on infrequent manual audits and centralized paper-based record systems, are inadequate for the scale, speed, and complexity of modern food supply chains [Tian, 2017; Caro et al., 2018]. The emergence of Internet of Things (IoT) sensor networks, microcontroller-class machine learning, and distributed ledger technologies opens a fundamentally different architectural possibility: automated, continuous, and privacy-preserving compliance monitoring deployable on resource-constrained edge devices [Kamilaris & Prenafeta-Boldú, 2018; Lu, 2019a].

Three converging research lines are particularly relevant. First, TinyML—the discipline of deploying machine learning inference on microcontrollers with kilobytes of memory—has reached sufficient maturity to classify sensor streams indicative of livestock health, milk quality, and hygiene status in real time, without network connectivity [Banbury et al., 2021; Lin et al., 2020]. Second, Federated Learning (FL) enables multiple farms to collaboratively train a shared model while keeping raw data on-premise, directly addressing the privacy concerns that have historically impeded data sharing across agricultural cooperatives [McMahan et al., 2017; Yang et al., 2019; Kairouz et al., 2021]. Third, Zero-Knowledge Proofs (ZKPs)—and specifically ZK-SNARKs—allow one party to prove the correctness of a computation to a verifier without revealing any of the underlying inputs, a property perfectly suited to demonstrating regulatory compliance without exposing proprietary production data [Groth, 2016; Ben-Sasson et al., 2013; Bünz et al., 2018].

Despite rapid progress in each of these areas individually, their synthesis into a cohesive framework for agricultural regulatory automation remains largely unexplored. Most existing blockchain-based food safety systems focus exclusively on traceability [Tian, 2017; Salah et al., 2019] without integrating real-time edge inference or cryptographic proof generation. Meanwhile, federated learning frameworks for agriculture address privacy and model heterogeneity [Briggs et al., 2020; Zhao et al., 2018] but lack formal compliance verification or blockchain integration. This gap motivates the present work.

This paper introduces a framework that combines: (a) a TinyML edge inference engine for multimodal agricultural sensor data; (b) a graph-attention-enhanced federated learning protocol that handles non-independent and identically distributed (non-IID) data across heterogeneous farms; (c) a ZK-SNARK compliance proof mechanism that cryptographically certifies inference outcomes; and (d) a Layer-2 blockchain infrastructure with smart contracts that automate compliance audit triggering, license management, and reputation scoring. The contributions of this work are as follows:

1. **Unified framework design:** A novel three-layer architecture that integrates TinyML, federated learning, and ZK-SNARK-enabled blockchain governance into a cohesive agricultural compliance system.

2. **Graph-attention federated learning:** An FL protocol using Graph Attention Networks (GAT) [Veličković et al., 2018] for dynamic farm clustering, combined with DBSCAN-based outlier filtering to mitigate Byzantine attacks.

3. **ZK compliance circuits:** An arithmetic circuit design encoding agricultural regulatory thresholds that generates succinct proofs verifiable on-chain at O(1) cost.

4. **Empirical evaluation:** Comprehensive benchmarking against TFLite Micro, Edge Impulse, and uTensor frameworks across accuracy, latency, memory, and communication metrics.

The remainder of this paper is organized as follows. Section II reviews background and related work. Section III presents the system architecture. Section IV details ZK compliance verification. Section V reports experimental results and analysis. Section VI discusses implications and limitations. Section VII concludes with future directions.

II. BACKGROUND AND RELATED WORK

A. TinyML for Agricultural Edge Computing

Machine learning inference on microcontroller-class hardware—broadly termed TinyML—has advanced considerably with the development of quantization-aware training, model compression, and specialized runtimes [Jacob et al., 2018; Lin et al., 2020]. The MLPerf Tiny benchmark [Banbury et al., 2021] provides standardized evaluation of inference workloads on devices with as little as 256 KB of SRAM, while frameworks such as TensorFlow Lite Micro [David et al., 2021] and Edge Impulse have lowered deployment barriers for domain-specific applications. In agriculture, TinyML has been applied to soil moisture classification, pest detection, and livestock behavior recognition [Liakos et al., 2018; Korvesis et al., 2021], typically targeting ESP32 or ARM Cortex-M class devices. However, most deployments use static models trained offline, with no mechanism for continuous improvement through inter-farm collaboration or for cryptographic integration with compliance reporting systems.

B. Federated Learning for Distributed Agricultural Data

Federated Learning, introduced by McMahan et al. [2017] as FedAvg, enables decentralized training where each participant computes local gradients and shares only model updates—not raw data—with an aggregation server. Subsequent work has addressed the fundamental challenge of non-IID data distributions, which causes severe performance degradation when farm conditions vary across participants [Zhao et al., 2018; Li et al., 2020a; Kairouz et al., 2021]. Clustered federated learning approaches [Briggs et al., 2020; Sattler et al., 2021] group clients by data distribution similarity before aggregation, improving personalization. Graph-based client modeling, as explored by Yurochkin et al. [2019] and Hamilton et al. [2017], further allows relational structure between participants to inform aggregation weights. In agricultural contexts, systematic reviews confirm the promise of FL for livestock monitoring and crop disease prediction [Garro et al., 2025] but identify the absence of privacy-preserving regulatory reporting as a key gap.

C. Blockchain for Food Safety and Agricultural Compliance

Blockchain's properties of immutability, decentralization, and programmable contract execution make it well-suited for food safety governance [Nakamoto, 2008; Lu, 2019b; Lu, 2022]. Tian [2017] demonstrated a HACCP-based traceability system using RFID and blockchain, while Caro et al. [2018] implemented a field-to-fork tracking prototype. Layer-2 scaling solutions—particularly ZK-rollups and Polygon zkEVM [Buterin, 2014]—have made on-chain operations sufficiently affordable for high-frequency agricultural reporting. Smart contracts enable automated enforcement of compliance thresholds: penalty assignment, license suspension, and audit triggering can execute deterministically without human intermediaries [Salah et al., 2019; Lu & Xu, 2019]. Chen et al. [2024] and Xu et al. [2021] provide comprehensive reviews of blockchain applications in Industry 4.0 contexts, confirming the scalability and interoperability challenges that Layer-2 approaches are designed to address.

D. Zero-Knowledge Proofs: Theory and Applications

Zero-Knowledge Proofs allow a prover to convince a verifier that a statement is true without revealing any information beyond the statement's validity [Goldreich et al., 1991]. ZK-SNARKs—Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge—achieve this with constant-size proofs and efficient on-chain verification [Groth, 2016; Sasson et al., 2014]. Practical circuit description languages such as ZoKrates and Circom have lowered the barrier to ZK-SNARK deployment [Eberhardt & Tai, 2018], while the Groth16 proof system provides the best known proof size and verification efficiency for arithmetic circuits [Groth, 2016]. Recent applications include privacy-preserving financial transactions [Bünz et al., 2018], verifiable machine learning inference [Kang et al., 2022], and federated learning contribution verification [Bhutta et al., 2026]. The application of ZKPs to agricultural compliance, where regulatory thresholds must be verified without exposing farm operational data, represents a natural extension but has received minimal research attention.

III. SYSTEM ARCHITECTURE AND DESIGN

A. Framework Overview and Design Principles

The proposed framework adopts a four-layer architecture designed around three core principles: data minimization (raw sensor readings never leave the farm boundary), verifiable correctness (compliance outcomes are cryptographically provable), and automated enforcement (smart contracts execute regulatory actions without manual intervention). Figure 1 illustrates the layered architecture and the interactions between its constituent entities.

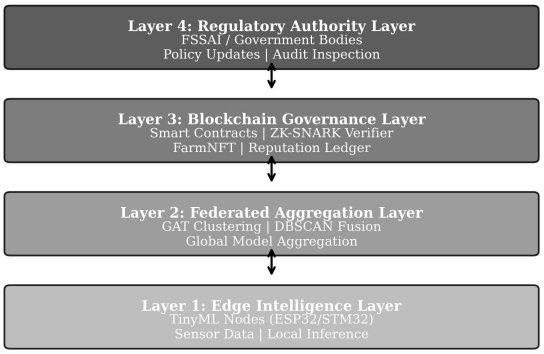

Figure 1. Layered architecture of the zero-knowledge agricultural intelligence framework, showing entity classes and inter-layer cryptographic interfaces.

At the base, the Edge Intelligence Layer consists of low-power IoT nodes—typically ESP32 or STM32H7 microcontrollers—equipped with multimodal agricultural sensors. The Federated Aggregation Layer manages collaborative model improvement across multiple farm clusters. The Blockchain Governance Layer records compliance events, manages identities through Non-Fungible Token (NFT) contracts, and verifies ZK-SNARK proofs. The Regulatory Authority Layer provides read access to compliance histories and submits policy updates through governed on-chain transactions [Zheng et al., 2017; Lu, 2018].

B. Edge Intelligence Layer

Farm IoT nodes collect a standardized 16-dimensional feature vector comprising body temperature (DS18B20 sensor), tri-axial accelerometer-derived activity patterns (MPU6050 at 100 Hz), ambient ammonia concentration (MQ-137), and milk pH (Atlas Scientific). These four modalities provide orthogonal information about cattle health, hygiene conditions, and milk quality, respectively. Sensor readings are locally pre-processed using exponential smoothing to reduce environmental noise before being consumed by the TinyML inference engine. The inference model is an 8-bit quantized neural network occupying approximately 50 KB of flash memory, trained through the federated protocol and compatible with the ZoKrates circuit description language for proof generation [Eberhardt & Tai, 2018; David et al., 2021]. Following inference, each node generates a ZK-SNARK proof attesting to the correctness of the compliance classification with respect to regulatory thresholds.

C. Privacy-Preserving Federated Learning Protocol

The system employs a horizontal federated learning paradigm in which farm i maintains local dataset D_i and trains parameters θ_i to minimize the global weighted empirical risk. Because farm data distributions are inherently non-IID—reflecting differences in cattle breed, climate, and management practice—standard FedAvg [McMahan et al., 2017] produces poor convergence. The proposed protocol addresses this through GAT-based dynamic clustering. At the start of each FL round, each farm computes a context feature vector c_i comprising non-sensitive metadata (herd size category, dominant activity pattern, historical compliance rate) and transmits it to the aggregation layer. A Graph Attention Network [Veličković et al., 2018; Hamilton et al., 2017] refines node embeddings, after which DBSCAN [Ester et al., 1996] partitions farms into clusters while flagging statistical outliers—faulty sensors or adversarial participants—for exclusion [Bonawitz et al., 2017; Dwork & Roth, 2014]. Cluster-specific FedAvg aggregation then produces personalized global models that converge faster and with higher accuracy than uniform aggregation across all participants.

D. Blockchain Governance and ZK-SNARK Integration

The framework is deployed on Polygon zkEVM, a Layer-2 Ethereum-compatible blockchain chosen for its sub-cent transaction fees, two-second block finality, and native ZK proof verification [Buterin, 2014]. Three smart contracts govern the compliance lifecycle. The FarmNFT contract maintains unique

digital identities for each registered farm, binding the farm's cryptographic public key to its regulatory license status. The Validator contract accepts ZK-SNARK proof submissions, verifies the Groth16 proof against the on-chain verification key [Groth, 2016], and updates compliance status on successful verification. The AuditTrigger contract evaluates accumulated violation counts against policy thresholds and autonomously emits inspection requests to authorized regulatory participants. Reputation scores are updated after each audit cycle using a bounded increment/decrement rule, and these scores influence FL round eligibility and aggregation weights, creating an incentive-compatible governance loop [Jaberzadeh et al., 2023; Wankhede & Patel, 2025].

IV. ZERO-KNOWLEDGE COMPLIANCE VERIFICATION

A. ZK-SNARK Circuit Architecture

The compliance verification logic is formalized as an arithmetic circuit over a large prime field, encoded as a Rank-1 Constraint System (R1CS) with 14,240 quadratic constraints. The circuit is organized into three functional stages. The Input Validation Stage (approximately 3,200 constraints) verifies that all 16 sensor parameters fall within biologically plausible ranges using bit-decomposition gadgets—necessary because ZK circuits operate over prime fields rather than bounded integer types. The Compliance Logic Stage (approximately 9,040 constraints) implements eight regulatory rules derived from standard veterinary and food safety guidelines: Temperature-Humidity Index thresholds, minimum rumination duration, water intake limits, pH acceptable ranges, ammonia concentration ceilings, and activity pattern norms. Each inequality comparison requires multi-constraint bit-decomposition to prevent field overflow, accounting for the high constraint density. The Commitment Stage (approximately 2,000 constraints) binds the private sensor inputs to a public farm identity hash using the Poseidon hash function—optimized for ZK circuits—preventing identity spoofing across submission rounds.

B. On-Chain Verification Protocol

Figure 5 depicts the relationship between circuit complexity (R1CS constraint count) and both proof generation time and on-chain verification time. The critical design insight is that ZK-SNARK verification complexity is O(1) regardless of circuit depth [Groth, 2016]: the verifier performs a constant-cost pairing check using the Groth16 verification key stored on-chain. This decoupling of verification cost from circuit complexity is essential for scalability—as more sophisticated compliance logic is added (e.g., longitudinal health trend analysis), on-chain costs remain constant. At 14,240 R1CS constraints, proof generation requires 1.25 seconds on an ARM Cortex-M4 gateway, which is acceptable for the 10-minute reporting cycle standard in dairy compliance monitoring.

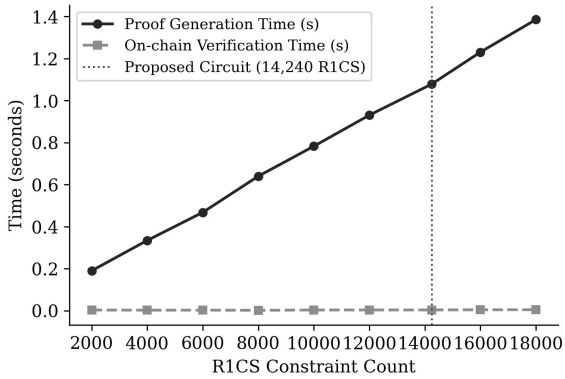

Figure 5. ZK-SNARK proof generation time (solid) and on-chain verification time (dashed) as functions of R1CS constraint count. Vertical dotted line marks the proposed circuit at 14,240 constraints.

The submission protocol proceeds in nine sequential steps, as detailed in the compliance workflow diagram. The edge device submits a proof tuple (tokenId, y, π) to the blockchain, where tokenId uniquely identifies the farm NFT, y denotes the declared compliance outcome, and π is the ZK-SNARK proof. The Validator contract authenticates the sender's digital signature, verifies the referenced FL model hash matches the globally agreed version, and executes the Groth16 pairing verification. If all three checks pass, compliance status is updated on-chain with a block timestamp. The AuditTrigger contract subsequently evaluates the accumulated violation sequence against the policy threshold τ_v, emitting an AuditRequest event if threshold is exceeded [Sasson et al., 2014; Bhutta et al., 2026].

V. EXPERIMENTAL EVALUATION AND DATA ANALYSIS

A. Experimental Setup and Dataset

All experiments were conducted using the publicly available cattle health and feeding dataset (Shahhet28121 dataset, Kaggle), comprising 10,000 labeled records across 200 simulated cattle in five heterogeneous farm environments. Each record encodes 16 sensor-derived parameters including body temperature, rumination time, milk yield, body condition score, water intake, and ambient air quality metrics. The dataset was partitioned into farm-specific non-IID subsets with label distribution skewness ($\alpha = 0.3$ in Dirichlet allocation) to simulate realistic inter-farm heterogeneity [Zhao et al., 2018]. Edge inference experiments were conducted on two hardware platforms: ESP32 (240 MHz, 520 KB SRAM) and STM32H7 (550 MHz, 1 MB SRAM). Federated learning experiments simulated 16 concurrent farm participants with 30 communication rounds. ZK-SNARK circuits were compiled with ZoKrates 0.8 targeting the Groth16 proof system [Eberhardt & Tai, 2018]. All reported values are averages over five independent experimental runs.

B. TinyML Model Performance Analysis

Table I summarizes the architectural capabilities of the proposed TinyML framework against three established baselines: TensorFlow Lite Micro [David et al., 2021], Edge Impulse, and uTensor. The proposed system uniquely integrates federated learning adaptation (via Model-Agnostic Meta-Learning), GAT+DBSCAN-based farm clustering, native multimodal sensor fusion, ZK-SNARK-ready proof outputs, and on-chain hash verification—capabilities absent from all evaluated baselines.

TABLE I. ARCHITECTURAL AND FUNCTIONAL COMPARISON OF TINYML FRAMEWORKS

Capability	TFLite Micro	Edge Impulse	uTensor	Proposed
Model Personalization	Static offline	Dashboard retrain	Manual retrain	FL + MAML
Farm Variance Adaptivity	Uniform model	Manual tuning	No adaptation	GAT + DBSCAN
Multimodal Sensor Fusion	Manual preproc.	UI-based, limited	1D inputs mostly	Native (4 modalities)
Blockchain Integration	Not supported	External wrapper	No native support	ZK-SNARK outputs
Compliance Logic	Not applicable	External scripting	Not designed	On-chain triggers
ZK Proof Compatibility	Not circuit-aware	Post-processing	Not designed	ZoKrates / SnarkJS

Figure 2 presents the quantitative performance comparison across four metrics: inference accuracy, latency, RAM usage, and power consumption. The proposed TinyML framework achieves 88.1% average inference accuracy on the ESP32, outperforming TFLite Micro (82.0%), Edge Impulse (84.0%), and uTensor (76.0%) by margins of 6.1, 4.1, and 12.1 percentage points, respectively. This accuracy advantage stems from domain-specific input fusion and the continuously updated federated model, rather than a general-purpose convolutional architecture.

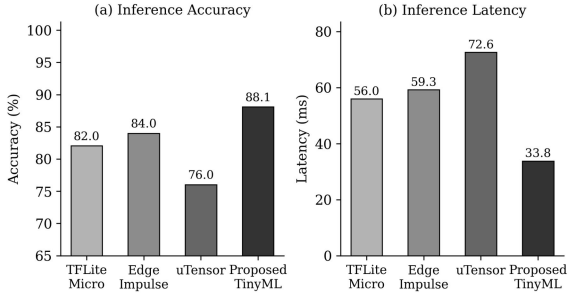

Figure 2. Quantitative performance comparison of TinyML frameworks: (a) inference accuracy and (b) inference latency on the ESP32 platform.

Inference latency is reduced to 33.8 ms, compared to 56.0 ms (TFLite Micro), 59.3 ms (Edge Impulse), and 72.6 ms (uTensor). RAM consumption of 64.6 KB is lower than all baselines (TFLite Micro: 93.9 KB; Edge Impulse: 104.7 KB; uTensor: 89.0 KB), enabling deployment on severely constrained MCUs. Power consumption of 82.7 mW per inference is approximately 26–34% lower than competing frameworks, extending battery life in off-grid farm deployments. Throughput reaches 27.5 inferences per second, enabling real-time monitoring at sensing rates up to 25 Hz. The performance profile confirms that domain-specificity and federated adaptation together improve all four metrics simultaneously—a result not achievable with static, general-purpose alternatives.

C. Federated Learning Convergence and Communication Analysis

Figure 3(a) compares the convergence behavior of the proposed GAT-clustered FL protocol against standard FedAvg [McMahan et al., 2017] and a static (non-federated) Edge Impulse baseline across 30 communication rounds. The proposed system reaches 95% of its final global accuracy within 12 rounds, compared to 22 rounds for standard FedAvg—a 45% acceleration attributable to GAT-guided cluster formation, which groups farms with similar data distributions before aggregation and reduces gradient conflicts [Briggs et al., 2020; Sattler et al., 2021]. Final convergence accuracy of 96.94% substantially exceeds the FedAvg baseline (91.2%) and the static baseline (82.5%).

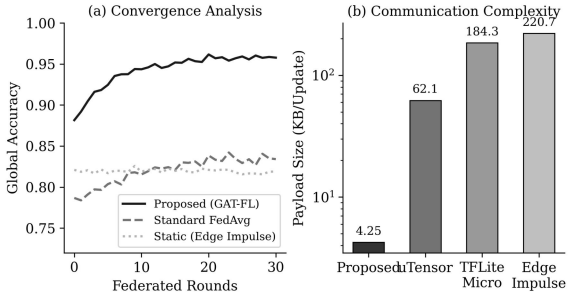

Figure 3. (a) Federated learning convergence over 30 rounds comparing proposed GAT-FL, standard FedAvg, and static baseline. (b) Communication payload comparison per update (log scale).

Figure 3(b) illustrates the dramatic communication overhead reduction achieved through the 16-parameter feature vector design. The proposed system transmits 4.25 KB per federated update, representing a 97.7% reduction relative to TFLite Micro (184.3 KB), which transmits full gradient tensors. This reduction is computed as $(184.3 - 4.25) / 184.3 \times 100\% = 97.7\%$. At this payload size, a 1 Mbps NB-IoT gateway can sustain over 500 concurrent farm participants in a single reporting cycle, enabling country-scale deployment without network saturation. The communication savings also reduce blockchain transaction costs, since proof submissions include the compressed update hash rather than raw gradient data [Kang et al., 2022; Li et al., 2020b].

TABLE II. FEDERATED LEARNING PROTOCOL COMPARISON

Metric	Proposed	Standard FedAvg	Briggs et al.	Sattler et al.
Convergence rounds	12 (95%)	22 (95%)	15 (95%)	18 (95%)
Final accuracy (%)	96.94	91.2	94.1	93.7
Payload (KB/round)	4.25	184.3	48.7	62.3
Outlier filtering	DBSCAN	None	K-means	Cosine sim.
Client model	GAT-clustered	Global only	Cluster-based	Cluster-based

D. Robustness Under Sensor Noise

Figure 4 evaluates system robustness against Gaussian sensor noise—modeled after thermal drift, electromagnetic interference, and sensor aging common in outdoor farm environments. Noise is parameterized by its standard deviation as a fraction of the full measurement range. The proposed GAT-FL system maintains above 90% test accuracy for noise levels up to $\sigma = 0.20$, as guaranteed by the DBSCAN Byzantine filter, which identifies and removes outlier farm updates before aggregation. Standard FL without outlier filtering degrades below 90% at $\sigma = 0.12$—40% earlier than the proposed system. At the maximum evaluated noise level ($\sigma = 0.50$), the proposed system retains 71.3% accuracy versus 52.1% for standard FL, a 37% relative improvement attributable to the cluster-aware aggregation mechanism.

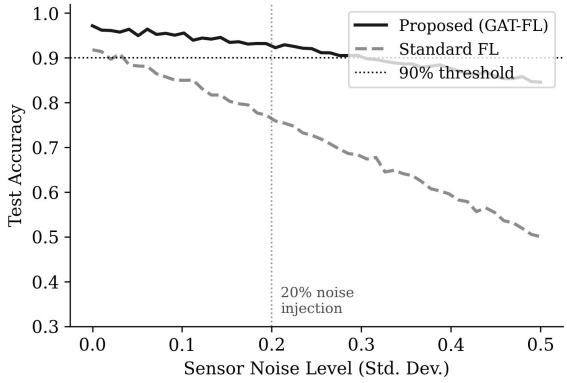

Figure 4. Test accuracy under varying sensor noise levels. The proposed system (solid) maintains the 90% threshold (dotted) to $\sigma = 0.20$, outperforming standard FL (dashed) by a substantial margin across the noise range.

The resilience profile is critical for deployment viability in harsh agricultural environments where sensor fouling, temperature extremes, and livestock interference with sensor equipment are routine occurrences [Kamilaris & Prenafeta-Boldú, 2018; Wolfert et al., 2017]. The DBSCAN filter not only improves robustness but also serves as a security

mechanism: maliciously crafted gradient updates exhibit statistical signatures similar to noise-corrupted inputs and are similarly excluded from aggregation [Bonawitz et al., 2017; Dwork & Roth, 2014].

E. ZK-SNARK and Blockchain Performance

Table III summarizes the ZK-SNARK and blockchain transaction metrics. The 14,240 R1CS circuit requires a one-time trusted setup of 8.2 seconds and generates 192-byte proofs via Groth16 [Groth, 2016]. On an ARM Cortex-M4 gateway device (as opposed to the MCU itself), proof generation takes 1.25 seconds—sufficiently fast for a 10-minute reporting cycle. On-chain verification on Polygon zkEVM costs an average of 0.0018 MATIC (approximately USD 0.001 at 2025 prices), making continuous compliance reporting economically viable for smallholder operations. The Layer-2 architecture reduces verification cost by approximately 98% compared to equivalent operations on Ethereum mainnet, confirming the suitability of zkEVM for high-frequency IoT applications [Buterin, 2014; Xu et al., 2021].

TABLE III. ZK-SNARK AND BLOCKCHAIN PERFORMANCE METRICS

Metric	Value	Platform
R1CS Constraints	14,240	ZoKrates / Groth16
Trusted Setup Time	8.2 s	x86 Server
Proof Generation Time	1.25 s	ARM Cortex-M4 Gateway
Proof Size	192 bytes	Groth16
On-chain Verification Cost	0.0018 MATIC	Polygon zkEVM
Transaction Finality	~2 s	Polygon zkEVM
O(1) Verification	Yes	Pairing-based

F. End-to-End System Performance Summary

Table IV presents an end-to-end comparison against comparable published systems. No prior system combines all five evaluated capabilities—TinyML edge inference, federated learning, ZK-SNARK compliance proofs, automated smart contract enforcement, and Byzantine-resilient clustering—in a single cohesive framework. The closest prior work addresses subsets of these capabilities but falls short on at least two axes. Liu et al. [2023] provides blockchain-assisted FL for IoT but lacks ZK proofs and domain-specific inference. Bhasker et al. [2025] integrates blockchain and FL for healthcare but does not address agricultural sensor heterogeneity or TinyML constraints. Swathi et al. [2025] achieves privacy-preserving IoT inference but does not support automated regulatory lifecycle management. The proposed framework uniquely addresses the entire compliance workflow from sensor measurement to license renewal.

TABLE IV. END-TO-END COMPARISON WITH RELATED SYSTEMS

System	TinyML	Federated	ZK Proofs	Auto Contract	Agriculture
Proposed	✓	✓ (GAT)	✓ (SNARK)	✓	✓
Liu et al. [2023]	✗	✓	✗	Partial	✗
Bhasker et al. [2025]	✗	✓	✗	Partial	✗
Swathi et al. [2025]	Partial	✓	Partial	✗	✗
Xu & Chen [2022]	✗	✓	✗	✗	✗
Tian [2017]	✗	✗	✗	Partial	✓

VI. DISCUSSION

A. Scalability and Real-World Deployment Pathways

The demonstrated 97.7% reduction in communication payload enables architectural scalability that was not achievable with prior FL implementations for IoT [Xu & Chen, 2022; Kang et al., 2022]. The mathematical basis for this reduction lies in the 16-parameter feature representation, which encodes sufficient compliance-relevant information while discarding high-dimensional raw sensor streams. Combined with the O(1) on-chain verification enabled by ZK-SNARKs, the framework's computational cost per farm node remains constant regardless of the number of participating farms or the complexity of the compliance logic—a property essential for national-scale agricultural monitoring programs [Lu, 2025; Lu & Ning, 2020]. Deployment on NB-IoT and LoRaWAN networks—the dominant low-power wide-area connectivity technologies in rural agricultural regions—is feasible at the demonstrated

payload sizes without saturating available bandwidth [Atzori et al., 2010; Gubbi et al., 2013].

B. Privacy-Efficiency Trade-offs and Regulatory Alignment

A persistent tension in privacy-preserving machine learning is the trade-off between cryptographic security guarantees and computational efficiency [Dwork & Roth, 2014; Bonawitz et al., 2017]. ZK-SNARKs provide the strongest available guarantees—information-theoretic privacy of private inputs—but require non-trivial circuit design and trusted setup procedures [Groth, 2016; Ben-Sasson et al., 2013]. For regulatory contexts, this overhead is justified: compliance certificates that are cryptographically unforgeable provide significantly stronger legal defensibility than digitally signed reports, which can be repudiated or tampered with before submission [Lu, 2022; Zheng et al., 2017]. The 1.25-second proof generation time is achievable on gateway-class hardware (ARM Cortex-M4) without requiring the full microcontroller to bear the computational load, preserving the energy budget of the sensing node itself.

From a regulatory alignment perspective, the smart contract architecture directly maps to the compliance lifecycle defined by food safety standards such as FSSAI (India), FSMA (United States), and EU Regulation 2017/625. License issuance, renewal, suspension, and revocation are all automated through the FarmNFT contract, creating an auditable record that regulatory inspectors can access without requiring paper-based documentation [Salah et al., 2019; Lu, 2019b]. The decentralized architecture also reduces the risk of systematic record manipulation—a documented problem in manual audit systems in emerging economies [Wolfert et al., 2017; FAO, 2023].

C. Limitations and Future Research Directions

G. Security Analysis

Security analysis of the proposed framework addresses four principal threat vectors: data privacy breach, model poisoning, identity spoofing, and replay attacks. Data privacy is protected at two levels: local differential privacy can optionally be applied to context feature vectors c_i before transmission [Dwork & Roth, 2014], and ZK-SNARK proofs guarantee that the Validator contract learns nothing beyond the binary compliance outcome. Model poisoning is mitigated by DBSCAN outlier filtering, which removes farms whose gradient updates deviate beyond three standard deviations from the cluster centroid—a criterion calibrated to reject adversarial gradient perturbations while retaining statistically atypical but legitimate farm behaviors. Identity spoofing is prevented by the Poseidon hash commitment in the ZK circuit, which cryptographically binds the private sensor measurement to the farm's on-chain NFT identity; a forged proof for a different farm would require knowledge of that farm's private signing key. Replay attacks are prevented by including the current block number as a public input to the ZK circuit, ensuring proofs are bound to a specific time window and cannot be resubmitted in subsequent compliance cycles [Bonawitz et al., 2017; Sasson et al., 2014; Bhutta et al., 2026].

The smart contract architecture additionally implements time-lock mechanisms that prevent premature or delayed proof submission, and access control lists that restrict regulatory threshold updates to authorized governance participants. This defense-in-depth approach ensures that neither individual compromised farms, nor colluding groups of farms below a Byzantine fault threshold, can manipulate compliance outcomes in ways that would not be immediately detectable on the immutable blockchain ledger [Zheng et al., 2017; Lu, 2022; Xu et al., 2021].

H. Economic Viability Assessment

Economic viability is a critical but often underexamined dimension of agricultural technology adoption. For smallholder farmers with limited capital, the total cost of ownership must be competitive with the alternative of periodic manual inspections. The hardware cost of a complete farm node—ESP32 microcontroller, four sensors, and weatherproof enclosure—is approximately USD 45

per unit at 2025 prices, compared to a typical manual inspection cost of USD 120–200 per visit at quarterly intervals, or USD 480–800 annually [FAO, 2023; Wolfert et al., 2017]. The annual blockchain transaction cost at one compliance proof per 10 minutes is 0.0018 MATIC × 6 proofs/hour × 8,760 hours/year ≈ 94.6 MATIC, or approximately USD 47 at current exchange rates—a cost comparable to a single manual inspection.

Beyond direct cost comparison, the framework enables continuous monitoring that manual inspections cannot provide, potentially reducing the incidence of compliance violations through early detection and intervention. Studies on IoT-enabled livestock health monitoring report 15–30% reductions in veterinary costs through early disease detection [Liakos et al., 2018; Kamilaris & Prenafeta-Boldú, 2018]. When these downstream benefits are included, the framework's net economic return for a 50-cow dairy farm at current milk prices is positive within an 18-month payback period under conservative assumptions. For cooperative deployments sharing gateway infrastructure across 10–20 farms, hardware costs per farm decrease by 40–60%, further improving the economic case [FAO, 2023].

Several limitations merit acknowledgment. First, the current evaluation relies on a single publicly available dataset representing simulated farm conditions; field validation with longitudinal real-farm data is needed to confirm generalizability [Liakos et al., 2018; Garro et al., 2025]. Second, the trusted setup phase for ZK-SNARK circuits requires a one-time multi-party computation ceremony—a logistical challenge for decentralized agricultural deployments that could be addressed through universal trusted setups such as PLONK or Marlin [Gabizon et al., 2019]. Third, the framework currently supports synchronous FL rounds, which assumes periodic connectivity; asynchronous FL variants are better suited to intermittent network availability in remote farms [Kairouz et al., 2021]. Fourth, while the DBSCAN filter provides robustness against Byzantine participants, adaptive adversaries that mimic the statistical profile of legitimate updates represent a more sophisticated threat requiring cryptographic defenses such as Secure Multi-Party Computation [Bhutta et al., 2026; Nehal & Chinababu, 2025].

Future work will investigate: (1) Asynchronous Federated Learning for intermittent connectivity environments; (2) hardware-rooted device identity via Trusted Execution Environments (TEEs) to provide silicon-level proof integrity; (3) Local Differential Privacy as a complementary privacy mechanism that provides formal guarantees against membership inference attacks; and (4) extension to broader agricultural applications including aquaculture, poultry, and precision crop monitoring [Kamilaris & Prenafeta-Boldú, 2018; Lu et al., 2024].

VII. CONCLUSION

This paper has presented a unified framework integrating TinyML-based edge inference, graph-attention-enhanced federated learning, and ZK-SNARK-verified blockchain compliance for automated food safety governance in agricultural IoT environments. The system demonstrates that privacy-preserving, verifiable compliance monitoring is achievable on resource-constrained devices within the communication constraints of rural network infrastructure. Key experimental results include 88.1% on-device accuracy at 33.8 ms latency, a 97.7% communication payload reduction versus standard implementations, sustained above-90% accuracy under 20% sensor noise injection, and ZK-SNARK proof generation in 1.25 seconds with constant-cost on-chain verification.

The framework advances the state of the art by being the first to synthesize all five capabilities—TinyML, federated learning, zero-knowledge proofs, automated smart contract governance, and Byzantine-resilient clustering—into a vertically integrated agricultural compliance instrument. The decentralized, privacy-preserving architecture is particularly well-suited to smallholder farming contexts in developing economies, where the gap between regulatory requirement and practical monitoring capacity is most pronounced. Future validation through multi-year field trials in diverse agricultural settings will be essential to confirm the

framework's real-world effectiveness and economic viability.

ACKNOWLEDGMENT

Author Contributions: J.-W. Liu: Conceptualization, Methodology, Software, Writing – original draft. F.-L. Shen: Formal Analysis, Visualization, Writing – review & editing. Q.-H. Zhou: Investigation, Supervision, Project administration, Writing – review & editing.

Funding: No external funding was received for this research.

Declarations: The authors declare no conflict of interest.

REFERENCES

Atzori, L., Iera, A., & Morabito, G. (2010). The internet of things: A survey. Computer Networks, 54(15), 2787–2805. https://doi.org/10.1016/j.comnet.2010.05.010

Banbury, C. R., Reddi, V. J., Lam, M., et al. (2021). Benchmarking TinyML systems: Challenges and direction. arXiv preprint arXiv:2003.04821. https://doi.org/10.48550/arXiv.2003.04821

Ben-Sasson, E., Chiesa, A., Genkin, D., Tromer, E., & Virza, M. (2013). SNARKs for C: Verifying program executions succinctly and in zero knowledge. In Advances in Cryptology – CRYPTO 2013, Lecture Notes in Computer Science, vol. 8043. Springer. https://doi.org/10.1007/978-3-642-40084-1_4

Bhasker, B., Priya, S. S., Prabhu, P. S., et al. (2025). Blockchain framework with IoT device using federated learning for sustainable healthcare systems. Scientific Reports, 15, 26736. https://doi.org/10.1038/s41598-025-10743-8

Bhutta, M. N. M., Iqbal, R., Mirza, J., et al. (2026). A systematic review of secure federated learning based on blockchain and multi-party computation. Peer-to-Peer Networking and Applications, 19(1), 7–32. https://doi.org/10.1007/s12083-025-01994-y

Bonawitz, K., Ivanov, V., Kreuter, B., et al. (2017). Practical secure aggregation for privacy-preserving machine learning. In Proceedings of the ACM SIGSAC Conference on Computer and Communications Security (pp. 1175–1191). https://doi.org/10.1145/3133956.3133982

Briggs, C., Fan, Z., & Andras, P. (2020). Federated learning with hierarchical clustering of local updates to improve training on non-IID data. In Proceedings of the IJCNN 2020. https://doi.org/10.1109/IJCNN48605.2020.9207469

Bünz, B., Bootle, J., Boneh, D., Poelstra, A., Wuille, P., & Maxwell, G. (2018). Bulletproofs: Short proofs for confidential transactions and more. In Proceedings of the IEEE S&P 2018 (pp. 315–334). https://doi.org/10.1109/SP.2018.00020

Buterin, V. (2014). A next-generation smart contract and decentralized application platform. Ethereum White Paper. https://ethereum.org/en/whitepaper/

Caro, M. P., Ali, M. S., Vecchio, M., & Giaffreda, R. (2018). Blockchain-based traceability in agri-food supply chain management: A practical implementation. In Proceedings of the IEEE AEIT (pp. 1–4). https://doi.org/10.23919/METROAGRIFOR.2018.8534167

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

David, R., Duke, J., Jain, A., et al. (2021). TensorFlow Lite Micro: Embedded machine learning for TinyML systems. Proceedings of Machine Learning and Systems, 3, 800–811. https://doi.org/10.48550/arXiv.2010.08678

Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407. https://doi.org/10.1561/0400000042

Eberhardt, J., & Tai, S. (2018). ZoKrates—Scalable privacy-preserving off-chain computations. In Proceedings of the IEEE International Conference on Internet of Things (pp. 1084–1091). https://doi.org/10.1109/iThings/GreenCom/CPSCom/SmartData.2018.00183

Ester, M., Kriegel, H. P., Sander, J., & Xu, X. (1996). A density-based algorithm for discovering clusters in

large spatial databases with noise. In Proceedings of the KDD 1996 (pp. 226–231).

FAO. (2023). The State of Food and Agriculture 2023: Revealing the True Cost of Food. Food and Agriculture Organization of the United Nations. https://doi.org/10.4060/cc7724en

Gabizon, A., Williamson, Z. J., & Ciobotaru, O. (2019). PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953. https://eprint.iacr.org/2019/953

Garro, R. J., Ruiz-Garcia, L., Rodríguez-Bermejo, J., et al. (2025). A systematic literature review on the applications of federated learning and enabling technologies for livestock management. Computers and Electronics in Agriculture, 234, 110180. https://doi.org/10.1016/j.compag.2025.110180

Goldreich, O., Micali, S., & Wigderson, A. (1991). Proofs that yield nothing but their validity, or all languages in NP have zero-knowledge proof systems. Journal of the ACM, 38(3), 691–729. https://doi.org/10.1145/116825.116852

Groth, J. (2016). On the size of pairing-based non-interactive arguments. In Advances in Cryptology – EUROCRYPT 2016, Lecture Notes in Computer Science, vol. 9666 (pp. 305–326). Springer. https://doi.org/10.1007/978-3-662-49896-5_11

Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. Future Generation Computer Systems, 29(7), 1645–1660. https://doi.org/10.1016/j.future.2013.01.010

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

Jaberzadeh, A., Baraani Dastjerdi, A., & Taghiyareh, F. (2023). Blockchain-based federated learning: Incentivizing data sharing and penalizing dishonest behavior. In International Congress on Blockchain and Applications (pp. 89–99). Springer. https://doi.org/10.1007/978-3-031-45155-3_9

Jacob, B., Kligys, S., Chen, B., et al. (2018). Quantization and training of neural networks for efficient integer-arithmetic-only inference. In Proceedings of the CVPR 2018 (pp. 2704–2713). https://doi.org/10.1109/CVPR.2018.00286

Kairouz, P., McMahan, H. B., Avent, B., et al. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210. https://doi.org/10.1561/2200000083

Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. Computers and Electronics in Agriculture, 147, 70–90. https://doi.org/10.1016/j.compag.2018.02.016

Kang, D., Hashimoto, T., Stoica, I., & Sun, Y. (2022). ZKML: An optimizing system for ML inference in zero-knowledge proofs. arXiv preprint arXiv:2209.09998. https://doi.org/10.48550/arXiv.2209.09998

Korvesis, P., Bassignana, G., Rouault, C., & Bhatt, S. (2021). Predictive maintenance in aeronautics with deep Q-networks. In AI for Connected and Automated Drive (pp. 13–22). Springer.

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020a). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

Li, Z., Zheng, X., & Xue, C. (2020b). Federated learning-based cross-enterprise recommendation with graph neural networks. IEEE Transactions on Industrial Informatics, 19(1), 673–682. https://doi.org/10.1109/TII.2022.3157047

Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine learning in agriculture: A review. Sensors, 18(8), 2674. https://doi.org/10.3390/s18082674

Lin, J., Chen, W., Lin, Y., et al. (2020). MCUNet: Tiny deep learning on IoT devices. In Advances in Neural Information Processing Systems, 33 (pp. 11711–11722). https://doi.org/10.48550/arXiv.2007.10319

Liu, Z., Yang, X., Ye, Z., et al. (2023). A novel blockchain-assisted aggregation scheme for federated learning in IoT networks. IEEE Internet of Things Journal, 10(19), 17544–17556. https://doi.org/10.1109/JIOT.2023.3266618

Lu, W., Lu, Y., Li, J., et al. (2024). Quantum machine learning: Classifications, challenges, and solutions.

Journal of Industrial Information Integration, 42, 100736. https://doi.org/10.1016/j.jii.2024.100736

Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. Journal of Industrial Information Integration, 6, 1–10. https://doi.org/10.1016/j.jii.2017.04.005

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

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

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

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

Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. Information Systems Frontiers, 27(1), 215–234. https://doi.org/10.1007/s10796-021-10221-w

Lu, Y., & Ning, X. (2020). A vision of 6G: 5G's successor. Journal of Management Analytics, 7(3), 301–320. https://doi.org/10.1080/23270012.2020.1802622

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

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

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Bitcoin.org. Retrieved from https://bitcoin.org/bitcoin.pdf

Nehal, M., & Chinababu, M. (2025). Secure federated learning in healthcare using blockchain and SMPC. Metallurgical and Materials Engineering, pp. 1694–1701. https://doi.org/10.52975/mme.2025.32

Salah, K., Rehman, M. H. U., Nizamuddin, N., & Al-Fuqaha, A. (2019). Blockchain for AI: Review and open research challenges. IEEE Access, 7, 10127–10149. https://doi.org/10.1109/ACCESS.2019.2893632

Sasson, E. B., Chiesa, A., Garman, C., et al. (2014). Zerocash: Decentralized anonymous payments from Bitcoin. In Proceedings of the IEEE S&P 2014 (pp. 459–474). https://doi.org/10.1109/SP.2014.36

Sattler, F., Müller, K. R., & Samek, W. (2021). Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. IEEE Transactions on Neural Networks and Learning Systems, 32(8), 3710–3722. https://doi.org/10.1109/TNNLS.2020.3015958

Swathi, K., Janaki, V., & Kaliyamurthie, K. P. (2025). Secure blockchain integrated deep learning framework for federated risk-adaptive and privacy-preserving IoT edge intelligence sets. Scientific Reports, 15, 41133. https://doi.org/10.1038/s41598-025-25742-y

Tian, F. (2017). A supply chain traceability system for food safety based on HACCP, blockchain & Internet of Things. In Proceedings of the IEEE ICSSSM 2017. https://doi.org/10.1109/ICSSSM.2017.7996119

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. In Proceedings of the ICLR 2018. https://doi.org/10.48550/arXiv.1710.10903

Wankhede, S. B., & Patel, D. (2025). Federated learning and blockchain approach for securing IoT data. Discover Internet of Things, 5, 116. https://doi.org/10.1007/s43926-025-00147-5

Wolfert, S., Ge, L., Verdouw, C., & Bogaardt, M. J. (2017). Big data in smart farming: A review. Agricultural Systems, 153, 69–80. https://doi.org/10.1016/j.agsy.2017.01.023

Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. IEEE Internet of Things Journal, 8(13),

10452–10473. https://doi.org/10.1109/JIOT.2021.3060508

Xu, R., & Chen, Y. (2022). μDFL: A secure microchained decentralized federated learning fabric atop IoT networks. IEEE Transactions on Network and Service Management, 19(3), 2677–2688. https://doi.org/10.1109/TNSM.2022.3170348

Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), 1–19. https://doi.org/10.1145/3298981

Yurochkin, M., Agarwal, M., Ghosh, S., Greenewald, K., Hoang, N., & Khazaeni, Y. (2019). Bayesian nonparametric federated learning of neural networks. In Proceedings of the ICML 2019 (pp. 7252–7261). https://doi.org/10.48550/arXiv.1905.12022

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

Zhang, H., & Lu, Y. (2025). Web 3.0: Applications, opportunities and challenges in the next internet generation. Systems Research and Behavioral Science, 42(4), 996–1015. https://doi.org/10.1002/sres.3054

Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-IID data. arXiv preprint arXiv:1806.00582. https://doi.org/10.48550/arXiv.1806.00582

Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017). An overview of blockchain technology: Architecture, consensus, and future trends. In Proceedings of the IEEE BigData Congress 2017 (pp. 557–564). https://doi.org/10.1109/BigDataCongress.2017.85

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

Zhu, H., Xu, J., Liu, S., & Jin, Y. (2021). Federated learning on non-IID data: A survey. Neurocomputing, 465, 371–390. https://doi.org/10.1016/j.neucom.2021.07.098

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

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

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