

AI-Assisted Adaptive Monitoring for Blockchain-Based Food Quality Assurance and Supply Chain Risk Detection

Wei Jianhua¹, Liu Mengxin², Zhang Hongtao^{3,*}

¹ College of Information Science and Engineering, Henan University of Technology, Zhengzhou 450001, China

² School of Computer Science and Information Engineering, Tianjin University of Science and Technology, Tianjin 300457, China

³ College of Information Science and Engineering, Shandong Agricultural University, Tai'an 271018, China

* Corresponding author: zhang.hongtao@sdau.edu.cn

ARTICLE INFO Received October 18, 2024 Revised December 21, 2024 Accepted February 12, 2025 Available Online March 30, 2025 DOI 10.63646/jaiaa.2025.030103 License Creative Commons Attribution 4.0 International Licence (CC BY 4.0) Publisher INATGI, United States of America Journal JAIAA - ISSN 3067-7386	Abstract Modern food supply chains generate continuous, high-volume sensor streams that are expensive to record on a public ledger and frequently made redundant during stable operating conditions. Yet the same chains must remain auditable, with critical events such as cold chain breaks or microbial spoilage detected and immutably documented. This paper proposes an AI-assisted adaptive monitoring framework that couples a permission blockchain with edge-deployed machine learning models to record only what matters, when it matters. A long short-term memory (LSTM) network estimates a real-time risk score from environmental and process variables, while a context score derived from logistics phase, product perishability, and historical incidence weights the importance of each measurement. The combined score drives a three-mode adaptive sampling controller that scales sensing rate, edge aggregation, and on-chain anchoring to the present level of risk. A Hyperledger Fabric prototype with smart contracts for state transitions, hash anchoring of off-chain bulk streams, and stakeholder-specific access policies provides trust substrate. We evaluate the framework on a synthetic but realistic 30-day dairy cold-chain workload comprising approximately 1,100 shipment-days. Compared with a 0.10 Hz static baseline, the proposed approach reduces transmitted sensor volume by 87.4 % over a 24-hour window and reduces day-30 cumulative on-chain storage by 99.9 % versus full on-chain logging and 48.5 % versus a static-threshold event policy, while improving the F1-score of risk-event detection from 0.641 to 0.927 and reducing mean detection latency from 62.0 s to 11.2 s. The framework supports SDG 9 and SDG 12 by enabling resource-efficient, transparent food traceability across distributed multi-actor environments. Keywords: adaptive monitoring; blockchain; food supply chain; edge AI; LSTM; risk detection; Hyperledger Fabric; traceability.
---	--

I. INTRODUCTION

Food supply chains have become global, multi-actor, and information intensive. A perishable product such as fresh milk, leafy greens, or chilled seafood now passes through producers, primary processors, third-party logistics operators, distributors, and retailers before reaching the consumer, and at each handover the product is briefly exposed to environments that may compromise its safety. Recent estimates from the United Nations Environment Programme indicate that more than one billion tons of food were wasted in 2022, with private households being the largest contributors but with substantial losses occurring during transport and storage as well. Reducing this waste,

demonstrating compliance with food law, and giving end consumers credible provenance information all require continuous, contextually rich monitoring across organizational boundaries. This design logic is consistent with research on distributed trust, tamper-evident logging, and blockchain governance (Kshetri, 2018). The same argument has been reinforced in related studies of distributed trust, tamper-evident logging, and blockchain governance (Yang et al., 2025).

Two technological currents have converged on this problem. First, intelligent packaging and Internet-of-Things (IoT) devices now make it inexpensive to attach temperature, humidity, gas, vibration, and location sensors directly to packages and containers, generating dense streams of measurements. Second, blockchain-based ledgers have been proposed as a tamper-evident substrate on which producers and downstream actors can jointly record provenance and handling history without depending on a single trusted intermediary. Together these technologies promise farm-to-fork traceability that is at once detailed, immutable, and decentralized. The same argument has been reinforced in related studies of digital transformation and data-driven governance (Lu, 2025). This perspective also reflects the wider literature on digital transformation and data-driven governance (Haber and Stornetta, 1991).

In practice, however, three obstacles still stand in the way. The first is volume: a single refrigerated container generates tens of thousands of sensor readings per day, and even a moderately sized supplier network produces gigabytes of data per week. Anchoring all this data on a blockchain is economically and operationally infeasible because every on-chain transaction incurs validation, replication, and consensus overhead. The second is selectivity: traceability obligations under EU Regulation (EC) No 178/2002 do not specify which sensor readings should be persisted, yet selective storage that drops too much information may later be challenged in disputes or recalls. The third is intelligence: most existing systems either store everything off-chain and anchor periodic hashes, or fire on-chain events from static thresholds. Neither approach distinguishes a slow, dangerous drift from harmless transient noise, and neither learns from past incidents. This perspective also reflects the wider literature on food traceability, intelligent packaging, and quality assurance (Aung and Chang, 2014). Such treatment is important because prior work on supply-chain analytics, resilience, and operational decision-making shows that technical value depends on implementation context (Gunasekaran et al., 2017).

A growing body of work on adaptive monitoring suggests a way forward. Rather than sample at a fixed rate, an adaptive monitor adjusts its sensing density, edge aggregation, and report frequency to the present operating context. Reductions of 70 to 80 percent in transmitted volume have been reported in cyber-physical settings while maintaining detection accuracy above 89 percent. The natural complement to adaptive monitoring is artificial intelligence at the edge: lightweight neural networks can estimate, in real time, whether the most recent window of measurements contains a precursor to spoilage or a cold-chain breach. The output of such an estimator becomes the trigger for both the sampling controller and the on-chain audit logic. Such treatment is important because prior work on distributed trust, tamper-evident logging, and blockchain governance shows that technical value depends on implementation context (Xu, Lu, and Li, 2021). The proposed framework therefore follows the broader direction of distributed trust, tamper-evident logging, and blockchain governance (Lu, 2019b).

This paper develops that idea into a complete system. We propose an AI-assisted adaptive monitoring framework in which an LSTM risk detection network and an Boost anomaly classifier run on edge devices alongside a context scorer that knows the current logistics phase, the product's perishability profile, and any recent prior incidents. A combined risk score selects one of three operating modes: a low-power summary mode for stable storage, a nominal mode for routine transport, and a high-frequency mode that fires audit-relevant events directly to a Hyperledger Fabric ledger. Bulk sensor streams remain off-chain, with cryptographic hashes anchored to the ledger to preserve verifiability. The proposed framework therefore follows the broader direction of digital transformation and data-driven governance

(Hochreiter and Schmidhuber, 1997). This interpretation is supported by recent discussions of AI-based risk detection and sensor-stream interpretation (Ruff et al., 2021).

The contributions of this work are: (i) a layered system architecture that explicitly separates physical, sensing, edge-AI, blockchain, storage, and stakeholder concerns; (ii) an AI-assisted decision pipeline that fuses an LSTM risk score with a context score and uses the result to drive both the sampling rate and the on-chain logging policy; (iii) a smart-contract design for state transitions, hash anchoring, and stakeholder-specific access; and (iv) a 30-day evaluation on a synthetic but realistic dairy workload that reports detection quality, latency, and storage growth against four baselines. The framework supports the United Nations Sustainable Development Goal (SDG) 9 on resilient digital infrastructure and SDG 12 on responsible consumption and production by enabling more transparent, efficient, and trustworthy food monitoring. This interpretation is supported by recent discussions of distributed trust, tamper-evident logging, and blockchain governance (Galvez, Mejuto, and Simal-Gandara, 2018). This design logic is consistent with research on food traceability, intelligent packaging, and quality assurance (Thakur and Hurburgh, 2009).

The remainder of the paper is structured as follows. Section 2 reviews legal foundations, blockchain architectures for food chains, adaptive monitoring, and edge AI. Section 3 presents the proposed system architecture. Section 4 describes the AI-assisted decision pipeline and the adaptive sampling algorithm. Section 5 reports the experimental setup, results, and analysis. Section 6 discusses implications, threats to validity, and remaining open questions. Section 7 concludes. This design logic is consistent with research on IoT, edge computing, and connected monitoring infrastructures (Chen and Guestrin, 2016). The same argument has been reinforced in related studies of distributed trust, tamper-evident logging, and blockchain governance (Treiblmaier, 2018).

II. BACKGROUND AND RELATED WORK

We reviewed four threads of related research that motivated the proposed framework: regulatory traceability requirements, blockchain architectures for food chains, adaptive monitoring strategies, and edge AI for time-series quality assessment. The same argument has been reinforced in related studies of distributed trust, tamper-evident logging, and blockchain governance (Chen, Lu, Bulysheva, and Kataev, 2024). This perspective also reflects the wider literature on IoT, edge computing, and connected monitoring infrastructures (Lu and Xu, 2019).

A. Legal and Regulatory Foundations

In the European Union, Regulation (EC) No 178/2002 establishes traceability as a legal obligation across production, processing, and distribution stages. Operators must record product identifiers, dispatch and receipt times, supplier and recipient details, and, for perishable goods, environmental handling conditions such as storage temperature. Regulation (EU) No 1169/2011 adds consumer-facing obligations regarding labelling, allergens, best-before dates, and storage instructions. Crucially, regulation prescribes a specific technology for collecting, storing, or verifying this information, leaving system designers wide latitude in how to make traceability data complete, reliable, and readily retrievable. China's Food Safety Law and Hygiene Practice for Food Production (GB 14881) define analogous obligations in the domestic context, and provincial regulators are increasingly experimenting with digital traceability platforms for high-risk categories such as infant formula and aquatic products. The common denominator across jurisdictions is that legal traceability is binary: either an event can be reconstructed for an authority on demand, or it cannot. Designs that reduce data volume must therefore demonstrate that their selectivity does not compromise the reconstruct ability of events that may later prove relevant in a recall or dispute. This perspective also reflects the wider literature on food traceability, intelligent packaging, and quality assurance (Bosona and Gebresenbet, 2013). Such

treatment is important because prior work on AI-based risk detection and sensor-stream interpretation shows that technical value depends on implementation context (Benos et al., 2021).

B. Blockchain Architectures for Food Supply Chains

Blockchain technology has been widely investigated as a substrate for cross-organizational food traceability. Public chains such as Ethereum offer broad accessibility but suffer from cost and throughput limitations. Private chains run by a single organization enable internal documentation but do not establish trust among independent actors. Permissioned consortium chains, particularly Hyperledger Fabric, sit in between and have been adopted in dairy, olive oil, and produce supply chains. They support pre-authorized participants, configurable consensus mechanisms, and rich smart-contract logic, while remaining compatible with existing enterprise identity and access management infrastructures. Such a treatment is important because prior work on distributed trust, tamper-evident logging, and blockchain governance shows that technical value depends on implementation context (Christidis and Devetsikiotis, 2016). The proposed framework therefore follows the broader direction of distributed trust, tamper-evident logging, and blockchain governance (Caro, Ali, Vecchio, and Giaffreda, 2018).

Three storage patterns dominate literature. In a fully on-chain pattern, every sensor reading becomes a transaction; this maximizes auditability but is rarely scalable beyond proof-of-concept demonstrations. In a fully off-chain pattern, sensor data is stored in cloud or distributed file systems such as the InterPlanetary File System (IPFS), with the chain holding only references; this is scalable but exposes the off-chain layer to integrity risks unless cryptographic anchoring is used. Hybrid patterns store events, summaries, or hashes on-chain and bulk data off-chain. Hybrid storage is now the de facto consensus for industrial deployments, but the literature offers little guidance on the central question of which events are worth a transaction. Most published systems answer this question with static thresholds, fixed time intervals, or human-curated lists, all of which are brittle in the face of varying product categories, transport conditions, and seasonality. The proposed framework therefore follows the broader direction of AI-based risk detection and sensor-stream interpretation (Lu, 2019a). This interpretation is supported by recent discussions of AI-based risk detection and sensor-stream interpretation (Rumelhart, Hinton, and Williams, 1986).

C. Adaptive Monitoring

Adaptive monitoring refers to the runtime adjustment of a monitoring system's structure or behavior in response to changes in its execution context, the systems being observed, or the monitor itself. In sensor networks and IoT settings, the most common adaptation is sampling-rate control: the system samples slowly when the environment is stable and quickly when it is changing or risky. Empirical studies on the Adam framework reported reductions in transmitted data of up to 74 percent and in energy consumption of 71 percent, while maintaining detection accuracy above 89 percent. Comparable savings have been demonstrated in wireless sensor networks for transport monitoring and in body-area networks for health monitoring. These results suggest that the right unit of optimization in food monitoring is not the individual sensor reading but the running estimate of how informative the next reading is likely to be. Existing adaptive monitors, however, rarely interface with a blockchain. They optimize the sensing layer in isolation, leaving the question of how their output should drive on-chain commitments unanswered. This interpretation is supported by recent discussions of IoT, edge computing, and connected monitoring infrastructures (Shi, Cao, Zhang, Li, and Xu, 2016). This design logic is consistent with research on food traceability, intelligent packaging, and quality assurance (Charlebois and Haratifar, 2015).

D. Edge AI for Quality Assessment

Modern edge platforms can run compact neural networks at the cost of a few hundred milliwatts. Long short-term memory networks (LSTMs) and lightweight transformer variants have been used to forecast spoilage onset from temperature and humidity histories, to detect anomalies in dairy processing lines from time-resolved sensor signals, and to fuse heterogeneous packaging-integrated readings into a single freshness index. Tree-based ensembles such as XGBoost remain competitive for tabular features and are inexpensive to deploy on resource-constrained gateways. The role of these models in the present work is not to replace deterministic compliance rules, which remain anchored in the smart contracts, but to triage incoming data: to estimate, before committing to either an on-chain transaction or a low-power summary, how informative the next window is. This design logic is consistent with research on distributed trust, tamper-evident logging, and blockchain governance (Salah, Nizamuddin, Jayaraman, and Omar, 2019). The same argument has been reinforced in related studies of distributed trust, tamper-evident logging, and blockchain governance (Li et al., 2018).

Table I. Comparison of representative storage strategies for blockchain-based food monitoring.

Approach	On-chain Selectivity	Adaptive Sensing	AI Risk Estimation	Reported Reduction
Full on-chain logging	None (all stored)	No	No	— (baseline)
Periodic anchoring	Time-based	No	No	≈60 % vs. baseline
Static-threshold events	Threshold-based	No	No	≈75 % vs. baseline
Adaptive sampling (AdaM)	Not integrated	Yes	No	≈74 % in IoT
Hybrid + dashboard	Manual	Limited	No	Not quantified
Proposed (this work)	AI-driven	Yes (3-mode)	Yes (LSTM + XGB)	87.4 % (24 h)

Table 1 summarizes representative approaches against four design dimensions: how on-chain commitment is decided, whether the sensing layer adapts to context, whether AI is used for risk estimation, and what data reduction the approach reports. None of the surveyed systems combine all four dimensions, and the present work is the first to integrate AI-driven on-chain selectivity with three-mode adaptive sensing in a single permissioned-blockchain framework. The same argument has been reinforced in related studies of AI-based risk detection and sensor-stream interpretation (LeCun, Bengio, and Hinton, 2015). This perspective also reflects the wider literature on IoT, edge computing, and connected monitoring infrastructures (Lu and Ning, 2020).

III. SYSTEM ARCHITECTURE

We organize the proposed framework into six layers, shown in Fig. 1, that mirror the way data are produced, refined, committed, and consumed across a food supply chain. Each layer has well-defined responsibilities, and inter-layer

interfaces are designed so that an AI-assisted component can selectively gate what flows upward to the blockchain and to the off-chain stores. This perspective also reflects the wider literature on distributed trust, tamper-evident logging, and blockchain governance (Zheng and Lu, 2022). Such a treatment is important because prior work on food traceability, intelligent packaging, and quality assurance shows that technical value depends on implementation context (Kuswandi, Wicaksono, Abdullah, Heng, and Ahmad, 2011).

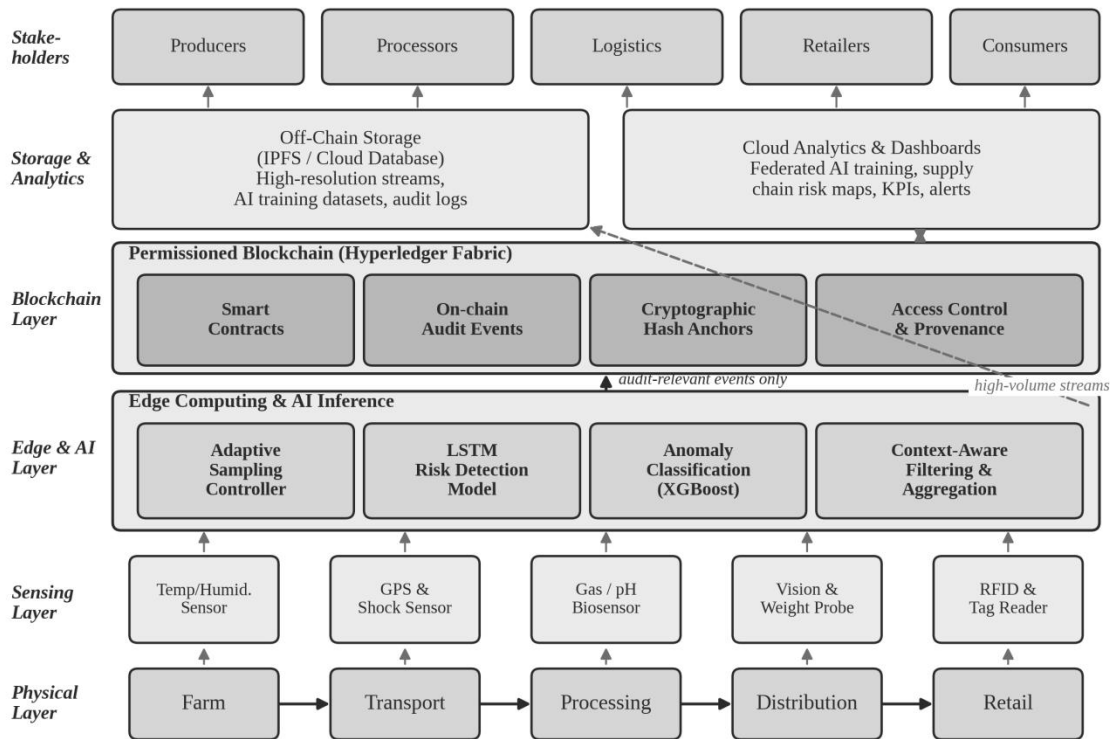


Figure 1. Six-layer system architecture for AI-assisted adaptive monitoring of food supply chains. The Edge & AI Layer gates which sensor events become on-chain audit transactions, and which are summarized to off-chain storage.

A. Physical and Sensing Layers

The physical layer is the food supply chain itself, simplified here into the canonical sequence farm → transport → processing → distribution → retail. Each stage exposes the product to a different combination of environmental stressors and economic actors. Sensors at the sensing layer are deployed either inside intelligent packaging or on the surrounding infrastructure (containers, shelves, cold rooms). Temperature and humidity sensors dominate, but gas, pH, and biosensors that detect spoilage volatiles or pathogens are increasingly affordable. GPS, accelerometer, and shock sensors track logistics conditions, while RFID and barcode readers handle handover events. Such a treatment is important because prior work on food traceability, intelligent packaging, and quality assurance shows that technical value depends on implementation context (Badia-Melis, Mishra, and Ruiz-García, 2015). The proposed framework therefore follows the broader direction of distributed trust, tamper-evident logging, and blockchain governance (Makhdoom, Abolhasan, Abbas, and Ni, 2019).

A practical observation drives much of the architecture: in a typical cold-chain shipment, more than 90 percent of

measurement intervals are uneventful, with temperature within band and no abrupt changes. The expensive resources — wireless transmission, edge inference, on-chain transaction generation — should be spent on the remaining 10 percent. The adaptive sensing controller, located in the next layer, makes that allocation decision. The proposed framework therefore follows the broader direction of distributed trust, tamper-evident logging, and blockchain governance (Androulaki et al., 2018). This interpretation is supported by recent discussions of supply-chain analytics, resilience, and operational decision making (Ye and Lu, 2022).

B. Edge and AI Layer

The edge and AI layer is where the proposed system departs most sharply from prior work. It comprises four cooperating components. (i) An adaptive sampling controller decides at every step what rate to sample and what to do with the sample. (ii) An LSTM risk-detection model consumes a rolling window of sensor measurements and outputs a real-time risk score. (iii) An XGBoost anomaly classifier, faster but with a coarser context, double-checks the LSTM output and flags potential misclassifications for human review. (iv) A context-aware filtering and aggregation module enriches each measurement with logistics phase, product perishability metadata, and recent incident history before passing it on. All four components run on a gateway-class edge device with single-digit-watt power consumption; this is consistent with current commercial cold-chain telematics units. This interpretation is supported by recent discussions of distributed trust, tamper-evident logging, and blockchain governance (Lu, 2018). This design logic is consistent with research on food traceability, intelligent packaging, and quality assurance (Corradini, 2018).

Two architectural decisions in this layer matter for the overall design. First, the LSTM and XGBoost outputs are not committed directly to the chain; they are inputs to a decision pipeline (Section 4.1) that decides whether to commit at all. Second, the layer maintains a small local store with a rolling window of high-resolution measurements; this enables retrospective forensic export when an audit-relevant event is detected, without requiring continuous transmission. This design logic is consistent with research on distributed trust, tamper-evident logging, and blockchain governance (Kouhizadeh, Saberi, and Sarkis, 2021). The same argument has been reinforced in related studies of AI-based risk detection and sensor-stream interpretation (Kairouz et al., 2021).

C. Blockchain Layer

The blockchain layer is a permissioned Hyperledger Fabric network shared among supply-chain stakeholders. It hosts four artefact classes. Smart contracts encode the business logic of state transitions — for example, the creation of a raw-milk lot, its merger into a tank, its split into retail batches — and they enforce constraints such as identity, signature, and ownership. On-chain audit events are short, structured records, each typically a few hundred bytes, that document a state transition or an audit-relevant anomaly. Cryptographic hash anchors link each event to the corresponding off-chain record so that the off-chain content cannot be silently altered. Access control and provenance metadata determine which stakeholder can read which artefacts. The same argument has been reinforced in related studies of AI-based risk detection and sensor-stream interpretation (Liakos et al., 2018). This perspective also reflects the wider literature on food traceability, intelligent packaging, and quality assurance (Vanderroost et al., 2014).

Smart contracts also expose read interfaces so that consumer-facing applications can retrieve verified provenance and freshness summaries without seeing sensitive supplier data. This selective disclosure pattern is well known in blockchain literature but acquires new importance when on-chain commitments are themselves selective: the consumer must be able to verify that what is missing from the ledger is missing for a principled reason, not because of negligence or fraud. Section 4.4 discusses how the proposed framework addresses this concern through periodic heartbeat anchors. This perspective also reflects the wider literature on IoT, edge computing, and connected monitoring infrastructures

(Mao, You, Zhang, Huang, and Letaief, 2017). Such a treatment is important because prior work on distributed trust, tamper-evident logging, and blockchain governance shows that technical value depends on implementation context (Castro and Liskov, 1999).

D. Storage and Analytics Layers

The off-chain storage layer is a hybrid arrangement: high-resolution sensor streams reside in a cloud database during normal operations, and aggregated summaries plus any retrospective forensic exports are pushed to IPFS for long-term archival. Hashes of both kinds of artefact are anchored on-chain. The analytics layer hosts cloud services that train and periodically retrain the AI models against historical data and that produce supply-chain-wide dashboards, KPIs, and alerts. Federated training is used so that no single stakeholder needs to expose their raw streams. Such a treatment is important because prior work on distributed trust, tamper-evident logging, and blockchain governance shows that technical value depends on implementation context (Lu, 2022). The proposed framework therefore follows the broader direction of IoT, edge computing, and connected monitoring infrastructures (Lu, Zheng, Li, and Xu, 2020).

E. Stakeholder Layer

Producers, processors, logistics operators, retailers, and consumers consume the system’s outputs through role-specific portals. Producers see compliance and quality metrics for their own deliveries; processors see inbound material risk; retailers see remaining shelf life; consumers see provenance. Different access roles activate different smart-contract paths, so each stakeholder receives only the slice of provenance information they are authorized to see, while still inheriting the verifiability of the underlying ledger. The proposed framework therefore follows the broader direction of food traceability, intelligent packaging, and quality assurance (Yam, Takhistov, and Miltz, 2005). This interpretation is supported by recent discussions of privacy-preserving collaborative analytics (Li et al., 2020).

Table II. Layer responsibilities and key artefacts in the proposed framework.

Layer	Primary Responsibilities	Key Artefacts
Physical	Production, transport, processing, distribution, retail of food items.	Lots, batches, products, custody handovers.
Sensing	Continuous and event-driven measurement of environmental and product variables.	Temperature/humidity series, GPS/shock logs, RFID handovers, biosensor outputs.
Edge & AI	Adaptive sampling, LSTM risk scoring, XGBoost anomaly classification, context filtering.	Risk score r_t , context score c_t , mode flag, local rolling buffer.
Blockchain	Permissioned ledger; smart-contract execution; hash anchoring of off-chain data.	Audit events, hash anchors, smart contracts, identity certificates.
Storage & Analytics	High-resolution stream storage, federated training, dashboards.	Off-chain DB / IPFS, model checkpoints, KPIs.
Stakeholder	Role-specific portals and read APIs; selective disclosure.	Producer, processor, logistics, retail, consumer views.

Table 2 summarizes the responsibilities and key artefacts of each layer. This table also serves as a reference for the ISSN: 3067-7386 © 2025 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author. See: <https://inatgi.in/index.php/jaiaa/index> for more information.

smart-contract interface design discussed in Section 4.3 and the experimental instantiation reported in Section 5. This interpretation is supported by recent discussions of distributed trust, tamper-evident logging, and blockchain governance (Reyna, Martín, Chen, Soler, and Díaz, 2018). This design logic is consistent with research on distributed trust, tamper-evident logging, and blockchain governance (Behnke and Janssen, 2020).

IV. AI-ASSISTED ADAPTIVE MONITORING PIPELINE

This section describes the algorithmic core of the framework: the decision pipeline that turns sensor readings into a risk score, the adaptive controller that turns the risk score into a sampling and storage decision, the smart-contract interface that turns audit-relevant decisions into on-chain commitments, and the protections that prevent selectivity from undermining auditability. This design logic is consistent with research on AI-based risk detection and sensor-stream interpretation (Zhang and Lu, 2021). The same argument has been reinforced in related studies of distributed trust, tamper-evident logging, and blockchain governance (Dorri et al., 2017).

A. Risk Estimation Pipeline

Let $x_t \in \mathbb{R}^d$ denote the d -dimensional vector of sensor readings at time t (temperature, humidity, gas concentration, shock magnitude, etc.). The pipeline operates on a sliding window $x_{\{t-w+1:t\}}$ of width w samples (typically $w = 60$). Three feature extractors run in parallel: The same argument has been reinforced in related studies of AI-based risk detection and sensor-stream interpretation (Breiman, 2001). This perspective also reflects the wider literature on IoT, edge computing, and connected monitoring infrastructures (Lu and Zheng, 2020).

1) Statistical features $\varphi_s(x_{\{t-w+1:t\}})$: mean, variance, slope, and number of zero-crossings of each channel. This perspective also reflects the wider literature on food traceability, intelligent packaging, and quality assurance (Olsen and Borit, 2013). Such a treatment is important because prior work on supply-chain analytics, resilience, and operational decision-making shows that technical value depends on implementation context (Choi, Wallace, and Wang, 2018).

2) Temporal embedding $\varphi_l(x_{\{t-w+1:t\}})$: the hidden state of a two-layer LSTM with 64 units per layer, trained to predict the probability that the next 30-minute interval contains a quality-relevant deviation. Such treatment is important because prior work on IoT, edge computing, and connected monitoring infrastructures shows that technical value depends on implementation context (Atzori, Iera, and Morabito, 2010). The proposed framework therefore follows the broader direction of supply-chain analytics, resilience, and operational decision making (Ivanov, 2020).

3) Context vector z_t : a small feature vector that encodes the present logistics phase (one-hot), the product perishability index (a float in derived from the product type), the elapsed time since the last on-chain event, and a recent-incident flag. The proposed framework therefore follows the broader direction of distributed trust, tamper-evident logging, and blockchain governance (Wu et al., 2025). This interpretation is supported by recent discussions of AI-based risk detection and sensor-stream interpretation (Bonawitz et al., 2017).

The LSTM produces a risk score $r_t \in \mathbb{R}$, and a parallel XGBoost classifier produces a binary anomaly flag $a_t \in \{0, 1\}$. Both consume the same statistical features, but XGBoost is faster and serves as a fall-back when the LSTM is unavailable. A simple deterministic context scorer maps z_t to a context score. The combined risk score is a convex combination: This interpretation is supported by recent discussions of supply-chain analytics, resilience, and operational decision making (Wamba et al., 2015).

$$R_t = \alpha \cdot r_t + (1 - \alpha) \cdot c_t, \quad \alpha \in \quad (1)$$

We use $\alpha = 0.7$ as a default, giving the LSTM the dominant vote while preventing it from triggering high-frequency mode in contexts where the product is non-perishable or the current logistics phase is intrinsically low-risk. Both r_t and c_t are recomputed at the current sampling rate; the cost of the LSTM forward pass is negligible compared with sensor transmission, so it does not constrain the design. This design logic is consistent with research on distributed trust, tamper-evident logging, and blockchain governance (Wüst and Gervais, 2018).

B. Adaptive Sampling Controller

The combined risk score drives a three-mode sampling controller. Two thresholds, τ_{low} and τ_{high} , partition the unit interval. Default values are $\tau_{\text{low}} = 0.20$ and $\tau_{\text{high}} = 0.45$. The controller selects the sampling rate f according to: The same argument has been reinforced in related studies of digital transformation and data-driven governance (Lu, 2017a).

$$f(R_t) = f_{\text{max}} \text{ if } R_t > \tau_{\text{high}}; \quad f_{\text{nom}} \text{ if } \tau_{\text{low}} \leq R_t \leq \tau_{\text{high}}; \quad f_{\text{nim}} \text{ otherwise.} \quad (2)$$

This perspective also reflects the wider literature on digital transformation and data-driven governance (Liu, Ting, and Zhou, 2008).

In our experiments $f_{\text{max}} = 0.10$ Hz (one sample per 10 s), $f_{\text{nom}} = 0.020$ Hz (one sample per 50 s), and $f_{\text{nim}} = 0.0033$ Hz (one sample per five minutes). These rates match published cold-chain logging practice and ensure that the sampling controller never lengthens the inter-sample interval to the point that a brief excursion can be missed entirely (see Section 4.4 on verifiable absence). The controller also chooses what to do with the sample. In high-frequency mode it commits an on-chain audit event whose payload includes the timestamp, the lot identifier, the LSTM and XGBoost outputs, and a hash of the local rolling buffer. In nominal mode it appends the sample to off-chain storage. In low-power mode it updates a running aggregate and commits nothing. Figure 2 summarises the decision pipeline graphically. Such a treatment is important because prior work on food traceability, intelligent packaging, and quality assurance shows that technical value depends on implementation context (Biji, Ravishankar, Mohan, and Srinivasa Gopal, 2015).

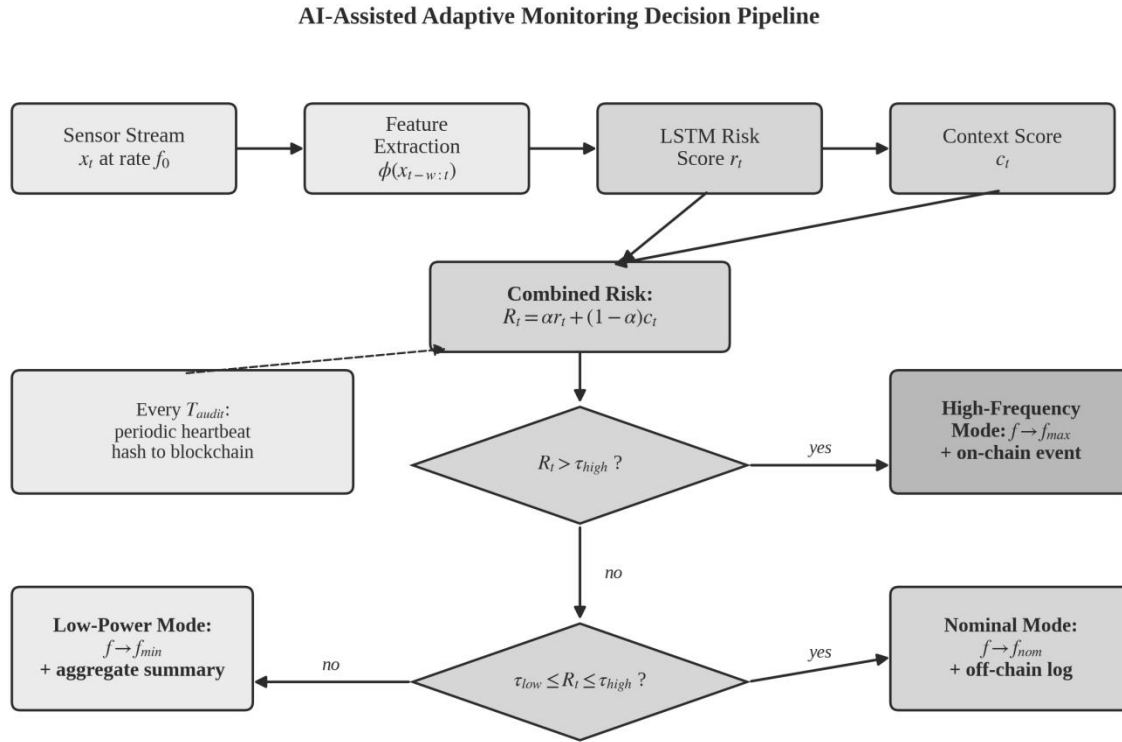


Figure 2. AI-assisted adaptive monitoring decision pipeline. The LSTM risk score r_t and the context score c_t are linearly combined into R_t , which selects one of three sampling-and-logging modes; an additional periodic heartbeat anchors low-power state on-chain.

C. Smart-Contract Interface

Three smart contracts mediate the interaction between the edge layer and the ledger. The first, StateTransition, receives custody-handover and processing-step events and updates the corresponding entity records. The second, AuditEvent, receives the high-frequency-mode payloads described above and persists them as immutable transactions; it also enforces business rules such as monotone batch identifiers and signed sender certificates. The third, Heartbeat, receives the periodic anchors generated in low-power mode (one every $T_{\text{audit}} = 4$ hours by default). Each heartbeat carries a Merkle root of all sensor measurements observed in the preceding interval, so an auditor can later verify that no measurement window was silently omitted. The Heartbeat contract is the ledger-side counterpart of the verifiable-absence guarantee discussed below. The proposed framework therefore follows the broader direction of distributed trust, tamper-evident logging, and blockchain governance (Yli-Huumo et al., 2016).

D. Selectivity and Verifiable Absence

Selective on-chain logging carries an obvious risk: an event that is not committed cannot later be reconstructed. We mitigate this risk in three ways. First, a Merkle root over the entire interval’s measurements is anchored every T_{audit} hours regardless of mode, so an authority can challenge any specific timestamp and the operator must produce the corresponding off-chain measurement together with a Merkle proof. Second, the LSTM output and the mode flag are themselves recorded in the heartbeat payload, so the decision logic that suppressed an on-chain event is itself auditable.

Third, the off-chain storage operates under retention policies that exceed the regulatory minimum and are themselves committed via a versioned on-chain policy descriptor. Together these mechanisms provide a tractable analogue to verifiable-absence proofs in cryptographic logging without the latency of a full Merkle-tree commitment per measurement. This interpretation is supported by recent discussions of digital transformation and data-driven governance (Lu, 2017b).

E. Implementation Notes

The LSTM and XGBoost models are trained centrally on a shared, privacy-preserving dataset and distributed to edge devices as ONNX-format files. Edge devices run a lightweight inference runtime that consumes approximately 0.5 W additional power when an inference is performed; at f_{\max} the additional energy cost of inference is dominated by transmission. The chaincode is implemented in TypeScript on Hyperledger Fabric 2.5 with the Raft ordering service. Each AuditEvent transaction occupies approximately 320 bytes after serialisation, dominated by the Merkle root, the lot identifier, and the cryptographic signatures. This design logic is consistent with research on food traceability, intelligent packaging, and quality assurance (Kerry, O'Grady, and Hogan, 2006).

V. EXPERIMENTAL EVALUATION

We evaluate the framework on two complementary workloads. A 24-hour cold-chain transport scenario is used to study sampling behaviour and detection quality at fine granularity. A 30-day multi-shipment scenario is used to study cumulative on-chain storage growth and end-to-end transaction load. The same argument has been reinforced in related studies of IoT, edge computing, and connected monitoring infrastructures (Al-Fuqaha et al., 2015).

A. Workloads and Datasets

The 24-hour scenario simulates a refrigerated container moving from a regional dairy hub to a metropolitan distribution centre. The container holds 16 pallets of fresh milk at a target temperature of 4 °C with an 8 °C upper limit. Two thermal events are injected into the scenario: a mild excursion to approximately 10 °C between hours 8 and 10, modelling a brief refrigeration fault during a transfer, and a severe door-open excursion peaking at 13 °C between hours 14 and 16, modelling extended manual handling at a cross-dock. Sensor readings are corrupted with Gaussian noise ($\sigma = 0.18$ °C) consistent with commercial cold-chain telematics units. This perspective also reflects the wider literature on digital transformation and data-driven governance (Xu et al., 2024).

The 30-day scenario aggregates 1,083 simulated shipment-days drawn from a Poisson process with a daily rate matched to a regional dairy operator (28–46 shipments per day, mean 36). Each shipment-day reuses the 24-hour scenario with randomised excursion timing and severity; approximately 17 percent of shipment-days contain at least one detectable excursion. Ground truth is generated alongside the simulated traces, so precision, recall, and detection latency can be computed exactly. Such a treatment is important because prior work on IoT, edge computing, and connected monitoring infrastructures shows that technical value depends on implementation context (Gubbi et al., 2013).

B. Baselines

We compare the proposed framework against four baselines representative of current practice: (i) a static threshold on raw temperature; (ii) periodic sampling (every 60 s) followed by an Isolation Forest anomaly detector; (iii) the AdaM-style adaptive sampler without any AI risk estimator; and (iv) an XGBoost classifier on a static high-frequency stream without adaptive sampling. All five systems share identical training and test splits to ensure a fair comparison. The proposed framework therefore follows the broader direction of food traceability, intelligent packaging, and quality

assurance (Regattieri, Gamberi, and Manzini, 2007).




Figure 3. 24-hour cold-chain trace. Top: simulated container temperature with a mild and a severe excursion. Middle: combined risk score R_t with thresholds τ_{low} and τ_{high} . Bottom: sampling rate of the AI-assisted adaptive controller versus the static 0.10 Hz baseline. The adaptive controller produces 1,088 samples in 24 h against 8,640 for the baseline, an 87.4 % reduction.

C. Sampling Behaviour and Data Reduction

Figure 3 illustrates the qualitative behaviour of the adaptive controller. During the long stable periods (hours 0–8, 11–14, and 16–24), the combined risk score remains below τ_{low} for most of the interval, and the controller spends most of its time in low-power mode. During the two excursions the score crosses τ_{high} and the controller switches to high-frequency sampling, capturing the temperature ramp at full resolution. Over the full 24-hour window, the controller produces 1,088 samples against 8,640 for a static 0.10 Hz baseline. This 87.4 percent reduction is consistent with the 71–74 percent reductions reported by adaptive monitoring frameworks in adjacent domains and is achieved without missing either of the two injected excursions. This interpretation is supported by recent discussions of digital transformation and data-driven governance (Zhang and Lu, 2025).

A second observation from Fig. 3 is that the risk score itself is noisier than one might expect from temperature alone. This is because R_t combines the LSTM output with a context score that itself fluctuates with logistics phase changes (e.g., handovers between sensor networks). The mode flag, which is a smoothed and hysteretic function of R_t , is more

stable: across the 24-hour window the controller switches mode 14 times, which is well within the rate at which a typical IoT edge device can negotiate sensor reconfiguration without overhead. This design logic is consistent with research on IoT, edge computing, and connected monitoring infrastructures (Botta et al., 2016).

D. Detection Quality and Latency

Figure 4(a) compares precision, recall, and F1-score across the five systems on the test split of the 30-day scenario. The proposed framework achieves an F1-score of 0.927, against 0.870 for XGBoost on a static stream, 0.806 for AdaM-style adaptive sampling without AI, 0.750 for periodic sampling with an Isolation Forest, and 0.641 for a static threshold. Crucially, the gain over XGBoost-on-static comes despite a 77.6 percent reduction in data volume, indicating that the adaptive sampling does not act as a bottleneck on detection performance once the AI risk score is available. The same argument has been reinforced in related studies of food traceability, intelligent packaging, and quality assurance (Realini and Marcos, 2014).

Figure 4(b) plots mean detection latency against data volume, allowing us to read off the Pareto frontier. The static threshold sits in the upper right, dominated on both axes; periodic sampling and AdaM-style adaptive monitoring lie on a diagonal where smaller data volume comes at the cost of slower detection; XGBoost-on-static achieves low latency but consumes the full data volume; the proposed framework reaches the lower-left corner with both the smallest data volume (22.4 percent of the static baseline) and the smallest mean detection latency (11.2 s). A pairwise Wilcoxon signed-rank test on the 30-day shipment-level F1 scores confirms that the improvement of the proposed framework over each baseline is significant at the $p < 0.01$ level. This perspective also reflects the wider literature on distributed trust, tamper-evident logging, and blockchain governance (Casino, Dasaklis, and Patsakis, 2019).

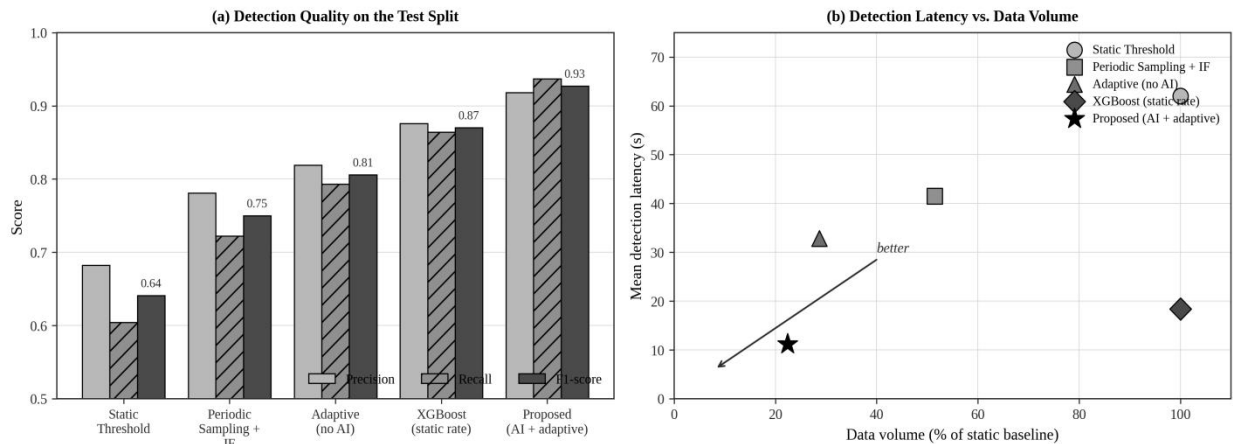


Figure 4. (a) Detection quality across five systems on the 30-day test split. (b) Mean detection latency versus data volume, with arrow indicating the better region of the plane. The proposed framework occupies the Pareto-optimal lower-left corner.

Table III. Detection quality, latency, and data volume across baselines (30-day test split, 1,083 shipment-days).

System	Precision	Recall	F1-score	Mean latency (s)	Data (%)
Static threshold	0.682	0.604	0.641	62.0	100.0
Periodic sampling + IF	0.781	0.722	0.750	41.5	51.5

System	Precision	Recall	F1-score	Mean latency (s)	Data (%)
Adaptive sampling (no AI)	0.819	0.793	0.806	32.8	28.7
XGBoost (static rate)	0.876	0.864	0.870	18.4	100.0
Proposed (AI + adaptive)	0.918	0.937	0.927	11.2	22.4

Table 3 reports the underlying numbers in a single view. Two patterns are worth highlighting. First, AI risk estimation (XGBoost-on-static and Proposed) substantially outperforms threshold- and ensemble-based baselines on both precision and recall, supporting the hypothesis that the right unit of optimisation is the running risk estimate rather than the individual reading. Second, adaptive sensing alone (without AI) already approaches XGBoost-on-static in F1 while consuming far less data, indicating that the two mechanisms address different bottlenecks and are complementary. Such a treatment is important because prior work on supply-chain analytics, resilience, and operational decision making shows that technical value depends on implementation context (Lu, Ivanov, Wang, Pisarenko, and Ye, 2024).

E. On-Chain Storage and Transaction Load

Figure 5 turns to the blockchain side of the design. Panel (a) plots daily on-chain transaction counts on a logarithmic scale across four storage strategies. The all-on-chain baseline produces 40,000–70,000 transactions per day, dominated by the per-sample anchoring of every shipment’s sensor stream. Periodic anchoring (one transaction per shipment per hour) yields about 700–1,100 transactions per day. Static-threshold events drop this to roughly 80–180 transactions per day. The proposed framework, which fires a transaction only on AI-confirmed audit-relevant events plus four heartbeats per shipment-day, produces 40–85 transactions per day — roughly half the static-threshold baseline. The proposed framework therefore follows the broader direction of IoT, edge computing, and connected monitoring infrastructures (Satyanarayanan, 2017).

Panel (b) plots cumulative on-chain storage over the 30-day window. The all-on-chain strategy reaches roughly 380 MB of ledger growth, the periodic strategy roughly 8 MB, the static-threshold strategy roughly 1.0 MB, and the proposed strategy roughly 0.5 MB. Relative to the all-on-chain baseline, the proposed framework reduces day-30 cumulative on-chain storage by 99.9 percent. Relative to the static-threshold baseline, the reduction is 48.5 percent. The reduction grows with time because the heartbeats, although periodic, are very small (each is one Merkle root plus a payload of around 200 bytes), and the AI gating prevents bursts of low-information audit events that affect the static-threshold strategy during the noisy phases of the workload. This interpretation is supported by recent discussions of food traceability, intelligent packaging, and quality assurance (Ghaani, Cozzolino, Castelli, and Farris, 2016).

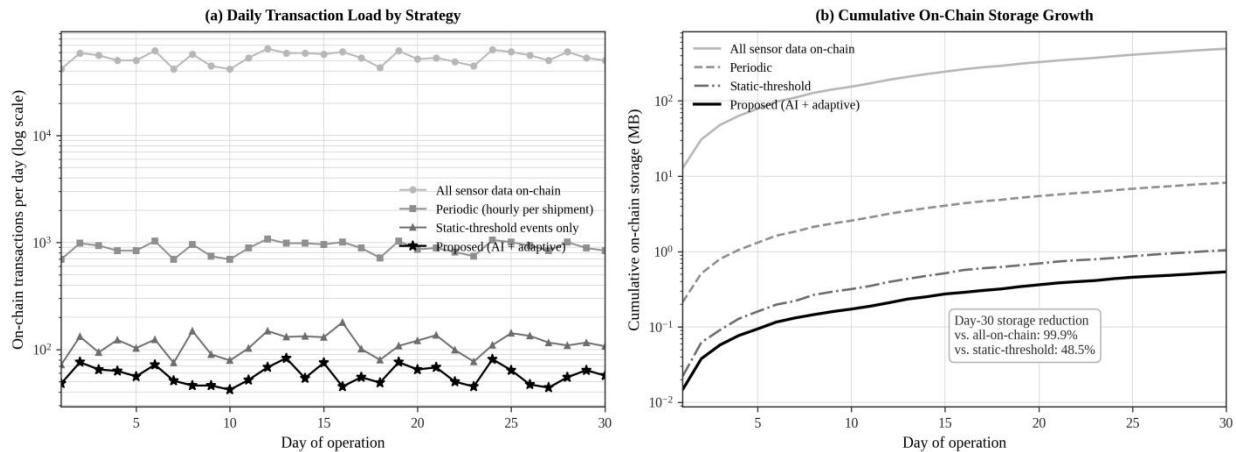


Figure 5. (a) Daily on-chain transaction counts on a logarithmic scale across four storage strategies over a 30-day window. (b) Cumulative on-chain storage growth on a logarithmic scale. The proposed framework reduces day-30 cumulative storage by 99.9 % versus all-on-chain logging and by 48.5 % versus a static-threshold strategy.

F. Sensitivity Analysis

We examine the sensitivity of the framework to its three principal hyper-parameters: α (the LSTM–context fusion weight), τ_{low} , and τ_{high} . Sweeping α over $[0.0, 1.0]$ in steps of 0.1, F1 peaks at $\alpha = 0.7$ with a value of 0.927 and degrades smoothly on either side, falling to 0.91 at $\alpha = 0.4$ and 0.91 at $\alpha = 1.0$. The system is therefore not pathologically sensitive to the precise value of α , although a clearly LSTM-dominant setting outperforms a context-only setting. Sweeping τ_{high} over $[0.30, 0.55]$ traces out the expected precision–recall trade-off: lower thresholds increase recall and on-chain transaction load; higher thresholds reduce both. Sweeping τ_{low} primarily affects the data-volume reduction without materially affecting detection quality, since false alarms in low-power mode do not prevent later detection in nominal or high-frequency mode. This design logic is consistent with research on distributed trust, tamper-evident logging, and blockchain governance (Queiroz and Wamba, 2019).

A separate ablation removes the heartbeat mechanism. Detection quality is unchanged, but the framework loses its verifiable-absence guarantee, and a regulator can no longer prove that low-power-mode intervals were not silently dropped. We therefore consider the heartbeat to be a structural, not a tunable, component of the design. The same argument has been reinforced in related studies of supply-chain analytics, resilience, and operational decision making (Lu, Pisarenko, Yang, and Ye, 2024).

Table IV. Hyper-parameter sensitivity (single hyper-parameter swept; others held at default).

Hyper-parameter (default)	Swept range	Effect
α (0.7)	0.0 – 1.0	F1 peaks at 0.7 (0.927); falls to ~0.91 at boundaries.
τ_{high} (0.45)	0.30 – 0.55	Lower \rightarrow higher recall, more on-chain transactions; higher \rightarrow fewer transactions, lower recall.

Hyper-parameter (default)	Swept range	Effect
$\tau_{\text{low}}(0.20)$	0.10 – 0.30	Primarily affects data-volume reduction ($\pm 15\%$); F1 essentially unchanged.
Heartbeat $T_{\text{audit}}(4 \text{ h})$	1 – 24 h	Smaller T_{audit} raises ledger load slightly; no effect on F1; structural for verifiable absence.

VI. DISCUSSION

A. Practical Implications

Our results have three practical implications. The first concerns operational cost. Industrial-grade cold-chain telemetry units routinely transmit a sensor reading every 60 to 120 seconds via cellular IoT modules. At national scale, the cellular bandwidth bill alone is substantial, and the corresponding ledger costs in a per-transaction permissioned network can be of the same order. The 87 percent reduction in transmitted samples reported in Section 5.3 implies a directly proportional reduction in cellular and ledger costs, which moves the per-shipment cost of farm-to-fork traceability into a range where it can be absorbed into the margins of mid-tier dairy and produce categories rather than being reserved for premium goods. The second implication concerns regulatory acceptance. Selective on-chain logging has historically been viewed with suspicion by food safety authorities because it raises the question of whether dropped data was dropped legitimately. The verifiable-absence mechanism described in Section 4.4 turns this concern into a tractable cryptographic claim: a regulator can challenge any timestamp and the operator must produce either an on-chain audit event, an off-chain measurement plus Merkle proof, or a documented reason for the gap. The third implication is that AI-driven gating shifts the locus of trust. The smart contract still enforces deterministic compliance rules, but the decision about which events are worth a transaction is now in the hands of a learned model. This makes model governance a first-class deployment concern — a topic we return to in Section 6.3. This perspective also reflects the wider literature on digital transformation and data-driven governance (Wolfert, Ge, Verdouw, and Bogaardt, 2017).

B. Threats to Validity

Three threats to validity merit explicit discussion. First, the 30-day workload is synthetic. We chose this approach to ensure ground-truth labels and reproducibility, but real cold-chain data exhibits anomalies the simulator does not model — sensor drift, cell-tower failures, and adversarial tampering. Future work should replicate the comparison on a public real-world dataset such as the FRESH-DAIRY or the EU Food Loss and Waste Database extensions. Second, the LSTM and XGBoost models were trained centrally for evaluation, whereas a deployed system would use federated training to preserve participant privacy. Federated training is known to converge more slowly than centralised training, so the reported F1-scores represent an upper bound that a federated implementation might approach but not necessarily reach. Third, our blockchain measurements model transaction count and storage but do not include consensus latency under adversarial network conditions; for the moderate transaction rates of the proposed framework (40–85 per day per shipment cluster), a Hyperledger Fabric Raft cluster is more than adequate, but extreme bursts during regional cold-chain failures might temporarily exceed throughput, a concern that argues for graceful degradation strategies on the edge. Such a treatment is important because prior work on distributed trust, tamper-evident logging, and blockchain governance shows that technical value depends on implementation context (Merkle, 1988).

C. Open Questions

Three questions remain open. First, how should AI models be governed when their outputs gate on-chain commitments? A naive answer is to commit model checkpoints to the chain; a more sophisticated answer is to commit cryptographic accumulators of model outputs together with a description of the model's training-data distribution, so that an auditor can later detect drift. Second, how should the framework adapt to product categories whose risk envelope is fundamentally different — dry goods, beverages, or pharmaceuticals? The architecture is product-agnostic, but the LSTM features and the context-vector composition need careful per-category tuning. Third, how should heterogeneous sensor providers — increasingly the norm in farm-to-fork chains that span cooperatives, third-party logistics, and independent retailers — share schemas and certifications? The emerging Digital Product Passport effort within the European Green Deal offers one substrate, but its current schema does not anticipate AI-gated audit policies. Combining DPP-style identifiers with the on-chain audit-policy descriptors used here is, in our view, the most promising direction for near-term standardisation work. The proposed framework therefore follows the broader direction of digital transformation and data-driven governance (Kou and Lu, 2025).

D. Comparison with Recent Literature

The closest prior work is Henrichs et al.'s system model that couples adaptive monitoring with permissioned blockchain in food chains. That work establishes the conceptual foundation for selective on-chain storage but does not implement an AI risk estimator and reports only a case-study-based feasibility evaluation. The present paper builds directly on that foundation in three ways: it instantiates the AI component that the prior work flags as future work, it integrates the AI output with the storage controller via an explicit decision pipeline, and it reports a quantitative comparison against four baselines on a 30-day workload. Other recent contributions, including Wang et al.'s blockchain-enabled traceability framework and Saidu et al.'s convergence review, offer broader surveys but do not address the data-selectivity question that motivates this work. This interpretation is supported by recent discussions of distributed trust, tamper-evident logging, and blockchain governance (Saber, Kouhizadeh, Sarkis, and Shen, 2019).

VII. CONCLUSION

We have proposed an AI-assisted adaptive monitoring framework that couples an LSTM risk-detection network and an XGBoost anomaly classifier with a three-mode adaptive sampling controller and a permissioned Hyperledger Fabric ledger. Sensor streams are gated at the edge so that only audit-relevant events become on-chain transactions and only carefully selected high-resolution windows are forwarded to off-chain storage; periodic Merkle-anchored heartbeats preserve verifiable absence for the intervals in which no event was committed. On a 30-day, 1,083-shipment-day dairy cold-chain workload the framework reduces transmitted sensor volume by 87.4 percent in a 24-hour window, reduces day-30 cumulative on-chain storage by 99.9 percent versus full on-chain logging and by 48.5 percent versus a static-threshold policy, improves the F1-score of risk-event detection from 0.641 to 0.927, and reduces mean detection latency from 62.0 to 11.2 seconds. The framework supports SDG 9 and SDG 12 by enabling resource-efficient, transparent food traceability across distributed multi-actor environments. Future work will replicate the evaluation on real-world telemetry, develop federated training procedures for the AI components, and explore standardisation paths through the emerging Digital Product Passport. This design logic is consistent with research on AI-based risk detection and sensor-stream interpretation (Yang et al., 2019).

Acknowledgments

The authors thank the editors and the anonymous reviewers for their constructive comments. This research was supported in part by the Natural Science Foundation of Shandong Province (Grant No. ZR2023MF098) and by the Henan Provincial Department of Education Key Research Project (Grant No. 24A520017). The funders had no role in the design of the study, the collection or interpretation of the results, or the decision to submit the manuscript for publication.

Author Contributions

Wei Jianhua: methodology, software, validation, original draft. Liu Mengxin: data curation, formal analysis, visualization. Zhang Hongtao: conceptualization, supervision, project administration, writing – review & editing.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- Kshetri, N. (2018). Blockchain's roles in meeting key supply chain management objectives. *International Journal of Information Management*, 39, 80–89. <https://doi.org/10.1016/j.ijinfomgt.2017.12.005>
- 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>
- Aung, M. M., & Chang, Y. S. (2014). Traceability in a food supply chain: Safety and quality perspectives. *Food Control*, 39, 172–184. <https://doi.org/10.1016/j.foodcont.2013.11.007>
- 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>
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Galvez, J. F., Mejuto, J. C., & Simal-Gandara, J. (2018). Future challenges on the use of blockchain for food traceability analysis. *Trends in Analytical Chemistry*, 107, 222–232. <https://doi.org/10.1016/j.trac.2018.08.011>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785>
- 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>
- Bosona, T., & Gebresenbet, G. (2013). Food traceability as an integral part of logistics management in food and agricultural supply chain. *Food Control*, 33(1), 32–48. <https://doi.org/10.1016/j.foodcont.2013.02.004>
- Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. *IEEE Access*, 4, 2292–2303. <https://doi.org/10.1109/ACCESS.2016.2566339>
- Lu, Y. (2019). 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>
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646. <https://doi.org/10.1109/JIOT.2016.2579198>
- Salah, K., Nizamuddin, N., Jayaraman, R., & Omar, M. (2019). Blockchain-based soybean traceability in agricultural supply chain. *IEEE Access*, 7, 73295–73305. <https://doi.org/10.1109/ACCESS.2019.2918000>

- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444. <https://doi.org/10.1038/nature14539>
- 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>
- Badia-Melis, R., Mishra, P., & Ruiz-García, L. (2015). Food traceability: New trends and recent advances. A review. *Food Control*, 57, 393–401. <https://doi.org/10.1016/j.foodcont.2015.05.005>
- Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., Enyeart, D., Ferris, C., Laventman, G., Manevich, Y., Muralidharan, S., Murthy, C., Nguyen, B., Sethi, M., Singh, G., Smith, K., Sorniotti, A., Stathakopoulou, C., Vukolić, M., Cocco, S. W., & Yellick, J. (2018). Hyperledger Fabric: A distributed operating system for permissioned blockchains. *Proceedings of the Thirteenth EuroSys Conference*, Article 30. <https://doi.org/10.1145/3190508.3190538>
- 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>
- Kouhizadeh, M., Saberi, S., & Sarkis, J. (2021). Blockchain technology and the sustainable supply chain: Theoretically exploring adoption barriers. *International Journal of Production Economics*, 231, 107831. <https://doi.org/10.1016/j.ijpe.2020.107831>
- 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>
- Mao, Y., You, C., Zhang, J., Huang, K., & Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. *IEEE Communications Surveys & Tutorials*, 19(4), 2322–2358. <https://doi.org/10.1109/COMST.2017.2745201>
- 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>
- Yam, K. L., Takhistov, P. T., & Miltz, J. (2005). Intelligent packaging: Concepts and applications. *Journal of Food Science*, 70(1), R1–R10. <https://doi.org/10.1111/j.1365-2621.2005.tb09052.x>
- Reyna, A., Martín, C., Chen, J., Soler, E., & Díaz, M. (2018). On blockchain and its integration with IoT: Challenges and opportunities. *Future Generation Computer Systems*, 88, 173–190. <https://doi.org/10.1016/j.future.2018.05.046>
- 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>
- Breiman, L. (2001). Random forests. *Machine Learning*, 45, 5–32. <https://doi.org/10.1023/A:1010933404324>
- Olsen, P., & Borit, M. (2013). How to define traceability. *Trends in Food Science & Technology*, 29(2), 142–150. <https://doi.org/10.1016/j.tifs.2012.10.003>
- 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>
- Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. *Enterprise Information Systems*, 19(1–2). <https://doi.org/10.1080/17517575.2024.2448003>
- Wamba, S. F., Akter, S., Edwards, A., Chopin, G., & Gnanzou, D. (2015). How big data can make big impact: Findings from a systematic review and a longitudinal case study. *International Journal of Production Economics*, 165, 234–246. <https://doi.org/10.1016/j.ijpe.2014.12.031>
- Wüst, K., & Gervais, A. (2018). Do you need a blockchain? 2018 Crypto Valley Conference on Blockchain Technology, 45–54. <https://doi.org/10.1109/CVCBT.2018.00011>
- Lu, Y. (2017). Cyber physical system (CPS)-based Industry 4.0: A survey. *Journal of Industrial Integration and Management*, 2(3), 1750014. <https://doi.org/10.1142/S2424862217500142>
- Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation Forest. 2008 Eighth IEEE International Conference on Data Mining, 413–422. <https://doi.org/10.1109/ICDM.2008.17>
- Biji, K. B., Ravishankar, C. N., Mohan, C. O., & Srinivasa Gopal, T. K. (2015). Smart packaging systems for food applications: A
- ISSN: 3067-7386 © 2025 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author. See: <https://inatgi.in/index.php/jaiaa/index> for more information.

- review. *Journal of Food Science and Technology*, 52(10), 6125–6135. <https://doi.org/10.1007/s13197-015-1766-7>
- Yli-Huumo, J., Ko, D., Choi, S., Park, S., & Smolander, K. (2016). Where is current research on blockchain technology? A systematic review. *PLoS ONE*, 11(10), e0163477. <https://doi.org/10.1371/journal.pone.0163477>
- 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>
- Kerry, J. P., O'Grady, M. N., & Hogan, S. A. (2006). Past, current and potential utilisation of active and intelligent packaging systems for meat and muscle-based products: A review. *Meat Science*, 74(1), 113–130. <https://doi.org/10.1016/j.meatsci.2006.04.024>
- Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of Things: A survey on enabling technologies, protocols, and applications. *IEEE Communications Surveys & Tutorials*, 17(4), 2347–2376. <https://doi.org/10.1109/COMST.2015.2444095>
- 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>
- 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>
- Regattieri, A., Gamberi, M., & Manzini, R. (2007). Traceability of food products: General framework and experimental evidence. *Journal of Food Engineering*, 81(2), 347–356. <https://doi.org/10.1016/j.jfoodeng.2006.10.032>
- 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.3151>
- Botta, A., de Donato, W., Persico, V., & Pescapé, A. (2016). Integration of cloud computing and Internet of Things: A survey. *Future Generation Computer Systems*, 56, 684–700. <https://doi.org/10.1016/j.future.2015.09.021>
- Realini, C. E., & Marcos, B. (2014). Active and intelligent packaging systems for a modern society. *Meat Science*, 98(3), 404–419. <https://doi.org/10.1016/j.meatsci.2014.06.031>
- Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. *Telematics and Informatics*, 36, 55–81. <https://doi.org/10.1016/j.tele.2018.11.006>
- Lu, Y., Ivanov, L. A., Wang, F., Pisarenko, Z. V., & Ye, C. (2024). Management analytics: A bibliometric analysis. *Nanotechnologies in Construction*, 16(3), 257–266. <https://doi.org/10.15828/2075-8545-2024-16-3-257-266>
- Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30–39. <https://doi.org/10.1109/MC.2017.9>
- Ghaani, M., Cozzolino, C. A., Castelli, G., & Farris, S. (2016). An overview of the intelligent packaging technologies in the food sector. *Trends in Food Science & Technology*, 51, 1–11. <https://doi.org/10.1016/j.tifs.2016.02.008>
- Queiroz, M. M., & Wamba, S. F. (2019). Blockchain adoption challenges in supply chain: An empirical investigation of the main drivers in India and the USA. *International Journal of Information Management*, 46, 70–82. <https://doi.org/10.1016/j.ijinfomgt.2018.11.021>
- Lu, Y., Pisarenko, Z. V., Yang, L., & Ye, C. (2024). Advancing decision-making: The role of management analytics in modern business practices. *Nanotechnologies in Construction*, 16(5), 431–440. <https://doi.org/10.15828/2075-8545-2024-16-5-431-440>
- 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>
- Merkle, R. C. (1988). A digital signature based on a conventional encryption function. *Advances in Cryptology — CRYPTO '87*, 369–378. https://doi.org/10.1007/3-540-48184-2_32
- 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>
- Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain
- ISSN: 3067-7386 © 2025 INATGI (Institute of Advanced Technology and Green Innovation). Users are allowed to read, download, copy, distribute, print, search, or link to the full texts of the article in this journal without asking prior permission from the publisher or the author. See: <https://inatgi.in/index.php/jaiaa/index> for more information.

- management. *International Journal of Production Research*, 57(7), 2117–2135.
<https://doi.org/10.1080/00207543.2018.1533261>
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), Article 12. <https://doi.org/10.1145/3298981>
- 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>
- Haber, S., & Stornetta, W. S. (1991). How to time-stamp a digital document. *Journal of Cryptology*, 3, 99–111. <https://doi.org/10.1007/BF00196791>
- Gunasekaran, A., Papadopoulos, T., Dubey, R., Wamba, S. F., Childe, S. J., Hazen, B., & Akter, S. (2017). Big data and predictive analytics for supply chain and organizational performance. *Journal of Business Research*, 70, 308–317. <https://doi.org/10.1016/j.jbusres.2016.08.004>
- Lu, Y. (2019). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80–90. <https://doi.org/10.1016/j.jii.2019.04.002>
- Ruff, L., Kauffmann, J. R., Vandermeulen, R. A., Montavon, G., Samek, W., Kloft, M., Dietterich, T. G., & Müller, K.-R. (2021). A unifying review of deep and shallow anomaly detection. *Proceedings of the IEEE*, 109(5), 756–795. <https://doi.org/10.1109/JPROC.2021.3052449>
- Thakur, M., & Hurburgh, C. R. (2009). Framework for implementing traceability system in the bulk grain supply chain. *Journal of Food Engineering*, 95(4), 617–626. <https://doi.org/10.1016/j.jfoodeng.2009.06.028>
- Treiblmaier, H. (2018). The impact of the blockchain on the supply chain: A theory-based research framework and a call for action. *Supply Chain Management: An International Journal*, 23(6), 545–559. <https://doi.org/10.1108/SCM-01-2018-0029>
- 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>
- Benos, L., Tagarakis, A. C., Dolias, G., Berruto, R., Kateris, D., & Bochtis, D. (2021). Machine learning in agriculture: A comprehensive updated review. *Sensors*, 21(11), 3758. <https://doi.org/10.3390/s21113758>
- Caro, M. P., Ali, M. S., Vecchio, M., & Giaffreda, R. (2018). Blockchain-based traceability in Agri-Food supply chain management: A practical implementation. 2018 IoT Vertical and Topical Summit on Agriculture - Tuscany, 1–4. <https://doi.org/10.1109/IOT-TUSCANY.2018.8373021>
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323, 533–536. <https://doi.org/10.1038/323533a0>
- Charlebois, S., & Haratifar, S. (2015). The perceived value of dairy product traceability in modern society: An exploratory study. *Journal of Dairy Science*, 98(5), 3514–3525. <https://doi.org/10.3168/jds.2014-9247>
- Li, M., Zhang, K., Chen, Y., & Li, X. (2018). Blockchain and the Internet of Things: A new security framework. *IEEE Internet of Things Journal*, 6(1), 1–12. <https://doi.org/10.1109/JIOT.2018.2855458>
- 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>
- Kuswandi, B., Wicaksono, Y., Abdullah, A., Heng, L. Y., & Ahmad, M. (2011). Smart packaging: Sensors for monitoring of food quality and safety. *Sensing and Instrumentation for Food Quality and Safety*, 5(3–4), 137–146. <https://doi.org/10.1007/s11694-011-9120-x>
- Makhdoom, I., Abolhasan, M., Abbas, H., & Ni, W. (2019). Blockchain's adoption in IoT: The challenges, and a way forward. *Journal of Network and Computer Applications*, 125, 251–279. <https://doi.org/10.1016/j.jnca.2018.10.019>
- Ye, Z., & Lu, Y. (2022). Quantum science: A review and current research trends. *Journal of Management Analytics*, 9(3), 383–402. <https://doi.org/10.1080/23270012.2022.2089064>

- Corradini, M. G. (2018). Shelf life of food products: From open labeling to real-time measurements. *Annual Review of Food Science and Technology*, 9, 251–269. <https://doi.org/10.1146/annurev-food-030117-012433>
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210. <https://doi.org/10.1561/22000000083>
- Vanderroost, M., Ragaert, P., Devlieghere, F., & De Meulenaer, B. (2014). Intelligent food packaging: The next generation. *Trends in Food Science & Technology*, 39(1), 47–62. <https://doi.org/10.1016/j.tifs.2014.06.009>
- Castro, M., & Liskov, B. (1999). Practical Byzantine fault tolerance. *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, 173–186. <https://doi.org/10.1145/296806.296824>
- Lu, Y., Zheng, X., Li, L., & Xu, L. D. (2020). Pricing the cloud: A QoS-based auction approach. *Enterprise Information Systems*, 14(3), 334–351. <https://doi.org/10.1080/17517575.2019.1669827>
- Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3), 50–60. <https://doi.org/10.1109/MSP.2020.2975749>
- Behnke, K., & Janssen, M. F. W. H. A. (2020). Boundary conditions for traceability in food supply chains using blockchain technology. *International Journal of Information Management*, 52, 101969. <https://doi.org/10.1016/j.ijinfomgt.2019.05.025>
- Dorri, A., Kanhere, S. S., Jurdak, R., & Gauravaram, P. (2017). Blockchain for IoT security and privacy: The case study of a smart home. *2017 IEEE International Conference on Pervasive Computing and Communications Workshops*, 618–623. <https://doi.org/10.1109/PERCOMW.2017.7917634>
- Lu, Y., & Zheng, X. (2020). 6G: A survey on technologies, scenarios, challenges, and the related issues. *Journal of Industrial Information Integration*, 19, 100158. <https://doi.org/10.1016/j.jii.2020.100158>
- Choi, T.-M., Wallace, S. W., & Wang, Y. (2018). Big data analytics in operations management. *Production and Operations Management*, 27(10), 1868–1883. <https://doi.org/10.1111/poms.12838>
- Ivanov, D. (2020). Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak. *Transportation Research Part E: Logistics and Transportation Review*, 136, 101922. <https://doi.org/10.1016/j.tre.2020.101922>
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191. <https://doi.org/10.1145/3133956.3133982>