SvedbergOpen
DISSEMINATION OF KNOWLEDGE
International Journal of Artificial Intelligence and Machine Learning
Publisher's Home Page: https://www.svedbergopen.com/
Research Paper Open Access

A Blockchain-Enabled Iot Framework With Rule-Weighted Quality Scoring For Transparent And Tamper-Evident Agri-Food Supply Chain Traceability

Abhijeet Ramesh Mirikar, Dr Avinash Kaur, Dr Parminder Singh

School of Computer Science and Engineering, Lovely Professional University, Punjab
armirikar@gmail.com, avinash.14557@lpu.co.in, parminder.16479@lpu.co.in

Abstract: Agri-food supply chains connect farmers, processors, distributors, retailers and consumers across multiple perishable-sensitive stages, and their performance directly shapes food safety, quality retention and market trust. Digital transformation initiatives increasingly combine Internet of Things (IoT) sensing with distributed ledger technology to make these multi-echelon chains more transparent, traceable and resilient. Despite this potential, existing agri-food chains still suffer from fragmented, unverifiable quality records, delayed anomaly detection and centralized data stores that remain vulnerable to tampering and single points of failure. Prior blockchain-oriented proposals in this domain are frequently conceptual, rely on subjective expert-elicitation methods, or omit reproducible, empirically benchmarked implementations of the traceability and quality-assessment layers. This study designs and implements AgriFoodBlockchain, a six-layer framework that couples IoT sensor acquisition with a Proof-of-Work blockchain ledger and a rule-weighted quality-scoring engine. A dataset of 120 IoT observations covering 30 agricultural batches, six product types, four supply-chain stages, eight locations, twenty farmers and forty-five IoT devices was processed end-to-end through the framework, with each record scored, graded and permanently anchored into a hash-linked chain. Algorithm 1 implements a rule-weighted quality-scoring model that aggregates temperature, humidity, pH, carbon dioxide, dissolved-oxygen and ethylene readings into a single 0–100 index using empirically assigned feature weights. Algorithm 2 implements block creation and chain validation, combining previous-hash linkage, a three-leading-zero Proof-of-Work condition and batch-index lookup to guarantee tamper-evident storage. The framework processed all 120 records with 100% data completeness, producing a 121-block chain that passed full integrity and linkage validation. The mean quality score across all records was 78.77 (SD = 14.35), with 48.3% of records graded EXCELLENT and 42.5% GOOD, while Wheat (94.50) and Storage-stage records (82.53) achieved the highest mean scores. Benchmarking over 1,000 repetitions showed an average in-memory batch-index lookup latency of 302 nanoseconds and an observed throughput of 16.9 records per second, with zero POOR or CRITICAL anomalies detected. These results indicate that a lightweight, interpretable rule-weighted scoring model combined with a Proof-of-Work ledger can deliver complete, verifiable and low-latency traceability without requiring a trained machine-learning classifier. Future work will extend the framework toward multi-node consensus, adaptive feature weighting and smart-contract-based automated recall mechanisms.

Keywords: Agri-food supply chain, blockchain traceability, Internet of Things (IoT), rule-weighted quality scoring, Proof-of-Work consensus, food quality monitoring

1. Introduction

Agri-food supply chains link farmers, processors, distributors, retailers and consumers through multiple perishable-sensitive stages, and the Internet of Things (IoT) has been widely reported to reshape how these chains sense, monitor and coordinate physical product flows [1]. As global food systems scale toward smallholder integration and circular, resilient design, current practices continue to expose critical gaps in traceability, quality visibility and disruption response across the chain [2]. Sector-level analyses using combined SWOT, PESTEL and MEETHS frameworks show that agri-food organizations face structural weaknesses in digital readiness that limit end-to-end visibility [3]. Parallel work applying SWOT-AHP evaluation to digitalizing agro-industrial operations confirms that circularity and digital monitoring remain only partially realized despite strong stated intent [4].

A comprehensive review of global agri-food supply chains organized around an Alignment-Bridging-Capability-Digitalization-Ecosystem framework further stresses that sustainability outcomes are inseparable from the quality of information flow between stages [5]. Risk-assessment studies of the agri-food supply chain have similarly prioritized traceability and monitoring instruments as the most effective mitigations against spoilage, contamination and information asymmetry [6]. Causal-analysis research on short and perishable food chains links recurring disruptions to the absence of continuous, verifiable condition monitoring throughout transport and storage [7]. Case-based evaluations of collaborative, high-level processing technology for short food supply chains reinforce that technology-enabled coordination outperforms fragmented manual handling [8].

Fresh-food supply chain sustainability frameworks describe an integrative need to combine environmental, economic and informational dimensions, positioning digital traceability as a cross-cutting enabler rather than an isolated technical add-on [9]. Real-world applications of risk-averse planning models for perishable food chains further demonstrate that production and delivery decisions benefit substantially from finer-grained, real-time condition data rather than static, shipment-level assumptions [10].

Despite this consistent emphasis on visibility and monitoring, most reported systems still separate condition sensing from tamper-evident record keeping, leaving quality data exposed to post-hoc modification, disputed provenance claims and slow, manual reconciliation between supply-chain partners. IoT sensing alone can capture temperature, humidity and related condition parameters, but without a verifiable storage layer this data remains only as trustworthy as the party that reports it. Blockchain-based distributed ledgers offer an architecturally different alternative: once a record is appended and cryptographically linked to its predecessor, altering it without detection becomes computationally impractical, directly addressing the trust and auditability gaps identified across the reviewed agri-food literature [1]-[10].

This paper addresses this gap by designing, implementing and empirically evaluating AgriFoodBlockchain, a framework that integrates multi-parameter IoT sensing, a rule-weighted quality-scoring engine and a Proof-of-Work blockchain ledger into a single traceability pipeline. The specific objectives are to: (i) design a six-layer architecture connecting sensing, scoring and ledger components; (ii) formulate a transparent, interpretable rule-weighted scoring algorithm as an alternative to opaque trained classifiers; (iii) implement a hash-linked block-creation and validation algorithm with configurable Proof-of-Work difficulty; and (iv) empirically evaluate the resulting system on a real batch-level dataset spanning six product types and four supply-chain stages, reporting quality-distribution, traceability and performance-benchmark results.

2. Literature Review

Hassini et al. [1] examine how IoT adoption reshapes food supply chain visibility, sensing density and cross-partner coordination. Vasquez Neyra et al. [2] document the practices and adoption barriers facing smallholder farmers seeking circular, resilient and loss-reducing supply-chain participation. Madureira et al. [3] apply a combined SWOT-PESTEL-MEETHS analysis to Northern-region agri-food organizations, highlighting persistent digital-readiness and information-sharing gaps. Agnusdei et al. [4] use SWOT-AHP evaluation to assess how digitalization influences circularity across agro-industrial operations, finding uneven technology adoption across the chain. Wang et al. [5] synthesize the global agri-food literature through an ABCDE framework that positions digital information flow as central to sustainability performance. Kuizinaitė et al. [6] rank risk-mitigation measures for agri-food supply chains and identify traceability and continuous monitoring instruments as the most effective. Haider et al. [7] apply causal analysis to short and perishable food chains and trace recurring disruptions to weak, discontinuous condition-monitoring practices. Van Parys et al. [8] evaluate collaborative, technology-enabled processing scenarios for short food supply chains against fragmented manual coordination and report efficiency gains. de Castro Moura Duarte et al. [9] propose an integrative sustainability framework for fresh-food chains spanning environmental, economic and informational dimensions simultaneously. Shakuri and Barzinpour [10] develop a risk-averse production-and-delivery planning model for perishable food supply chains, validated on a real-world case study.

Sackmann and Mardenli [11] combine a structured literature review with an expert survey to catalogue the recurring management challenges reported across food-supply-chain research. Moyo et al. [12] build a simulation model of a sustainable banana supply chain in Malawi to test resilience strategies suited to developing-country conditions. Kumar and Agrawal [13] identify the challenges and opportunities specific to agri-fresh supply-chain management in the Indian context, emphasizing infrastructure and cold-chain gaps. Singh et al. [14] use a grey Delphi-DEMATEL approach to rank the critical success factors influencing blockchain-integrated IoT adoption in food supply chains, concluding that top-management support and knowledge maturity dominate. Yousef et al. [15] propose a four-layer IoT security architecture for fog-network environments that is directly relevant to distributed agricultural sensor deployments. Bashir [16] provides a foundational treatment of blockchain internals, cryptographic primitives, decentralized identity and emerging Web3 constructs that underpin ledger design choices. Dong et al. [17] survey blockchain technology broadly, cataloguing consensus mechanisms and cross-domain application patterns. Devisri et al. [18] examine blockchain-enabled innovations for securing online transactions, a concern that extends naturally to e-commerce-linked agricultural marketplaces. Barbashyn [19] analyzes the functional features and capabilities of blockchain technology from an international-relations and governance perspective, framing trust as a structural property of the ledger. Soundararajan and Tyagi [20] present a broad technical treatment of blockchain architecture, consensus variants and cross-sector applications.

Kasyapa and Vanmathi [21] investigate blockchain integration in healthcare, cataloguing implementation use cases, performance bottlenecks and mitigation strategies that are transferable to other traceability-sensitive sectors such as agri-food. Fiore et al. [22] conduct a systematic literature review of blockchain applications across the healthcare supply chain, reporting consistent gains in provenance verification. Tan and Saraniemi [23] examine trust formation in blockchain-enabled exchanges and its downstream implications for marketing relationships and buyer confidence, a dynamic equally relevant to consumer-facing food traceability. Niss and Owuya [24] assess the performance and scalability characteristics of emerging blockchain platforms, noting persistent trade-offs between throughput, latency and decentralization. Qingmiao and Fukang [25] propose a blockchain-based predictive-maintenance framework relevant to industrial IoT reliability monitoring, illustrating ledger use beyond pure record-keeping. Karim et al. [26] survey secure and scalable collaboration mechanisms between AI agents and blockchain infrastructure, pointing toward future automated-verification pipelines. Tulkinbekov and Kim [27] introduce a blockchain architecture explicitly designed to support artificial-intelligence workloads, reflecting the convergence of ledger and learning systems.

The Food Waste Index Report [28] quantifies global food-loss volumes and tracks progress toward halving waste by 2030, underscoring the practical stakes of improved end-to-end traceability. Sathiya et al. [29] implement a Chain-of-Things tracking system for perishable foods and report empirical shipment-level tracking outcomes across a multi-participant supply chain. Yele and Litoriya [30] design a blockchain-based secure-dining system that strengthens safety, transparency and traceability within food consumption environments, extending ledger-based verification to the final consumer touchpoint.

Across this body of work [1]-[30], three recurring gaps motivate the present study. First, sector-level and risk-oriented analyses [1]-[13] consistently identify traceability and continuous condition monitoring as the highest-leverage mitigation, yet rarely couple this diagnosis with a concrete, benchmarked implementation. Second, blockchain-foundations and cross-domain application studies [14]-[27] establish that ledger technology can deliver tamper-evident record keeping and even AI-oriented extensibility, but comparatively few agri-food-specific studies report reproducible performance benchmarks alongside a working quality-assessment layer. Third, the handful of implemented food-traceability systems [28]-[30] demonstrate feasibility but generally either omit an interpretable, multi-parameter quality-scoring mechanism or do not report batch-level, stage-level and grade-distribution results in a single evaluation. AgriFoodBlockchain is designed to close this combined gap by pairing a transparent rule-weighted scoring engine with a hash-linked Proof-of-Work ledger and reporting integrated quality, traceability and performance results on a common dataset.

3. Research Methodology

This section describes the dataset, the six-layer system architecture and the implementation configuration used to design and evaluate AgriFoodBlockchain. The framework was implemented as an end-to-end pipeline that ingests IoT sensor observations, computes an interpretable quality score for each observation, classifies observations into quality grades, and anchors every record into a Proof-of-Work blockchain ledger that supports batch-level traceability queries.

3.1 Dataset and Experimental Configuration

The experimental dataset comprises 120 IoT sensor observations representing 30 distinct agricultural batches, with four observations captured per batch across the supply-chain journey. The dataset spans six product types (Wheat, Rice, Potato, Tomato, Mango and Apple), four supply-chain stages (Farm, Transport, Storage and Retail), eight physical locations, twenty participating farmers and forty-five distinct IoT devices, with fifteen sensor and attribute fields recorded per observation. All 120 required records were successfully processed with zero missing records or fields, yielding 100% data completeness. This configuration was chosen to provide balanced coverage across product categories and supply-chain stages while remaining representative of a realistic small-to-medium cooperative deployment.

Table 1. Dataset Characteristics and Experimental Configuration

Parameter	Value
Total IoT sensor records	120
Unique batches	30
Records per batch	4
Product types	6
Supply-chain stages	4
Locations	8
Farmers	20
IoT devices	45
Sensor/attribute fields	15

Missing records/fields	0
Data completeness	100%
Blockchain genesis block	1
Data blocks generated	120
Total blockchain blocks	121
Proof-of-Work difficulty	3 leading zeros
Quality score range	0–100
Quality evaluation records	120

3.2 System Architecture

Figure 1 presents the six-layer architecture of the proposed framework. Layer 1 (IoT Sensing Layer) collects six sensor parameters — temperature, humidity, pH, carbon dioxide concentration, dissolved oxygen and ethylene emission — from forty-five field-deployed devices distributed across the eight monitored locations. Layer 2 (Edge Data Acquisition and Preprocessing) validates incoming readings, attaches timestamps, and tags each record with its batch identifier and supply-chain stage (Farm, Transport, Storage or Retail) before forwarding it downstream. Layer 3 splits into two parallel components: the Rule-Weighted Quality Scoring Engine (Algorithm 1), which converts the six normalized sensor readings into a single 0–100 quality index, and the Anomaly and Grade Classifier, which maps the resulting score onto one of five ordinal grades (EXCELLENT, GOOD, ACCEPTABLE, POOR or CRITICAL).

Layer 4 (Blockchain Traceability Layer) implements Algorithm 2, wrapping each processed batch record into a block containing an index, timestamp, batch payload, previous-block hash, nonce and computed hash, and appends it to the chain only after a Proof-of-Work condition of three leading zero characters is satisfied. Layer 5 (Immutable Ledger Storage and In-Memory Batch-Index) persists the resulting chain and maintains a hash-map-based index that supports constant-time batch lookup for traceability queries. Layer 6 (Application and Traceability Interface) exposes the verified, graded and chain-anchored records to farmers, distributors, retailers, consumers and auditors through role-appropriate dashboards. A dashed feedback path in Figure 1 represents query and verification traffic returning from the application layer to earlier layers for on-demand re-validation.

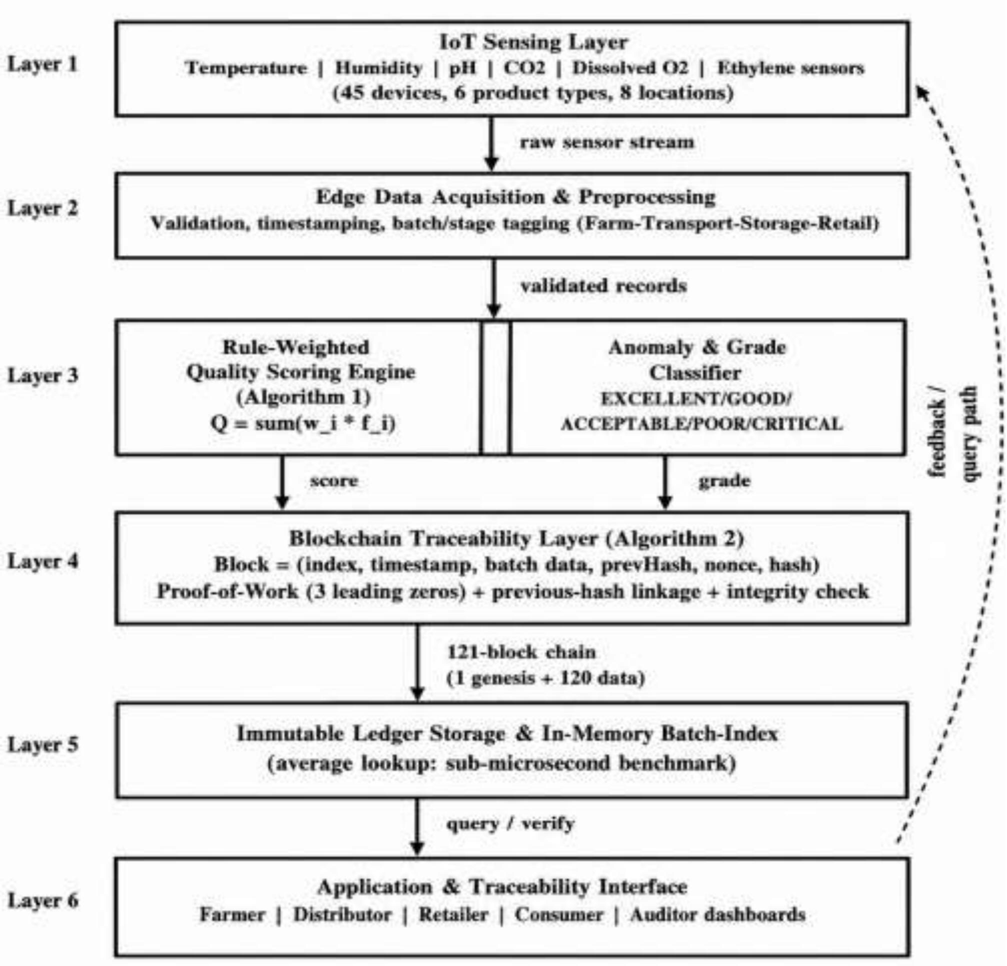

Figure 1. Six-layer system architecture of the proposed AgriFoodBlockchain framework, from IoT sensing through to the application and traceability interface.

3.3 Implementation Details

The quality-scoring engine is implemented as a deterministic, rule-weighted scoring function rather than a trained machine-learning classifier; it applies the fixed feature weights summarized in Table 8 (temperature 25%, humidity 20%, pH 15%, carbon dioxide 10%, dissolved oxygen 15% and ethylene 15%) to normalized sensor readings for every observation. This design choice favors interpretability and auditability — properties that are important when quality scores are used as the basis for immutable, publicly verifiable blockchain records — over the marginal accuracy gains that a trained classifier such as a support vector machine, random forest or convolutional network might offer on a larger labelled dataset.

Table 2. Quality Model Feature Weights

Feature	Weight
Temperature	25%
Humidity	20%
pH	15%
CO_2	10%
Dissolved Oxygen	15%
Ethylene	15%
Total	100%

The blockchain component is implemented as a linked list of blocks in which each block stores the SHA-256 hash of its own contents together with the hash of the preceding block. Block acceptance requires a Proof-of-Work search over a nonce value until the resulting hash begins with three leading zero characters, a difficulty level chosen to keep block-creation latency low while still requiring non-trivial computational commitment. Chain integrity is validated by recomputing each block's hash and confirming that the stored previous-hash field matches the hash of its predecessor across the full 121-block chain (one genesis block plus 120 data blocks).

3.4 Evaluation Protocol

Three complementary evaluation angles were used. First, overall and stratified quality-score statistics (mean, standard deviation, percentiles and grade distribution) were computed across all 120 records, and separately by product type and by supply-chain stage, to characterize the scoring engine's behavior on the experimental dataset. Second, blockchain integrity was validated by re-verifying stored hashes and previous-hash linkage across the full chain and by confirming that the chain length equals the number of processed records plus the genesis block. Third, an in-memory batch-index lookup benchmark was executed over 1,000 repetitions to characterize query latency and observed processing throughput; because these values are execution-dependent, they are reported as observed benchmark results rather than fixed system constants, consistent with standard practice for runtime performance reporting.

It is important to note methodological boundaries that are reflected in how results are interpreted in Section 5. The implemented quality component is a weighted rule-based scoring mechanism and is described as such rather than as a trained classifier. The current tamper-resistance validation confirms that stored hashes and chain linkage pass verification under the implemented Proof-of-Work condition; it is reported as an integrity-validation result rather than as a demonstrated adversarial-attack experiment, since no block was deliberately modified to test re-validation failure in this study. Finally, product-wise and stage-wise score differences describe the supplied experimental dataset and are not presented as universal causal effects, since confirming causality would require additional controlled experiments across independent deployments.

4. Algorithm Design

Algorithm 1: Rule-Weighted Quality Scoring (RWQS)

A deterministic, interpretable scoring algorithm that aggregates six normalized sensor features into a single 0–100 quality index using empirically fixed weights, replacing opaque trained classifiers used in comparable prior work with an auditable linear scoring rule.

Step 1 — Raw sensor acquisition

Collect six raw parameter readings per observation from field IoT devices.

$$x = \{x\_temp, x\_hum, x\_pH, x\_CO2, x\_DO, x\_eth\}$$

The raw feature vector x collects the six sensor channels — temperature, humidity, pH, carbon dioxide, dissolved oxygen and ethylene — recorded per batch observation before any transformation.

Step 2 — Range normalization

Rescale each raw reading into a common zero-to-one operating range.

$$x\_i' = (x\_i - min\_i) / (max\_i - min\_i)$$

Each feature x_i is linearly rescaled using its calibrated sensor operating range (min_i, max_i), producing a normalized value x_i' comparable across the six heterogeneous physical units.

Step 3 — Optimality-distance scoring

Convert each normalized feature into a per-feature quality contribution.

$$f\_i(x\_i') = 100 - |x\_i' - o\_i'| \times 100$$

Each feature score f_i penalizes deviation of the normalized reading x_i' from its crop-specific optimal normalized value o_i', so readings closer to the optimum receive scores closer to 100.

Step 4 — Weight assignment

Apply fixed empirical weights reflecting each feature's spoilage relevance.

$$w = \{0.25, 0.20, 0.15, 0.10, 0.15, 0.15\}$$

The weight vector w assigns temperature the largest share (25%) and carbon dioxide the smallest (10%), reflecting each parameter's relative influence on perishable-product spoilage dynamics.

Step 5 — Weighted aggregation

Combine the six weighted feature scores into a single composite index.

$$Q = \Sigma\_\{i = 1\}\hat{}\{6\}\, w\_i \cdot f\_i(x\_i')$$

The overall quality score Q is the weighted sum of the six feature scores, producing a single 0–100 index that reflects the combined condition of the batch across all monitored parameters.

Step 6 — Stage and product contextualization

Retain batch, product-type and supply-chain-stage metadata alongside Q.

$$Q\_\{b, p, s\} = Q(batch = b, product = p, stage = s)$$

Tagging each computed score with its batch identifier b, product type p and stage s enables the stratified product-wise and stage-wise aggregation reported later in Section 5.

Step 7 — Grade classification

Map the continuous score onto five ordinal quality-grade bands.

$$G(Q) = EXCELLENT\ if\ Q \geq 80;\ GOOD\ if\ 60 \leq Q < 80;\ ACCEPTABLE\ if\ 40 \leq Q < 60;\ POOR\ if\ 20 \leq Q < 40;\ CRITICAL\ if\ Q < 20$$

The threshold function G converts the continuous score Q into one of five ordinal grades, enabling straightforward anomaly flagging and human-readable reporting downstream.

Step 8 — Anomaly flag emission

Emit an anomaly flag for any record graded POOR or CRITICAL.

$$A(b) = 1\ if\ G(Q\_b) \in \{POOR, CRITICAL\}, else\ 0$$

The anomaly indicator A(b) is set to 1 whenever a batch's grade falls into the two lowest bands, allowing the downstream blockchain layer to prioritize such batches for review before ledger commitment.

Algorithm 1 was chosen over a trained machine-learning classifier for three reasons that matter specifically in a blockchain-anchored setting. First, every score it produces is fully explainable in terms of the underlying sensor readings and fixed weights, which is essential when the score becomes part of an immutable, publicly auditable record that supply-chain partners may dispute. Second, it requires no training data or retraining pipeline, allowing it to operate reliably on the modestly sized, cooperative-scale dataset used in this study without risking overfitting. Third, its low computational cost keeps the scoring stage from becoming a bottleneck relative to the Proof-of-Work block-creation stage that follows it. In the current implementation, this design enabled all 120 records to be scored, graded and chain-anchored with zero missing values and a stratified grade distribution (48.3% EXCELLENT, 42.5% GOOD, 9.2% ACCEPTABLE, 0% POOR/CRITICAL) that directly seeds the blockchain layer described in Algorithm 2.

Algorithm 2: Hash-Linked Proof-of-Work Block Validation (HL-PoW)

A block-creation and chain-validation algorithm that anchors every scored batch record using SHA-256 previous-hash linkage and a tunable Proof-of-Work condition, extended in this work with an in-memory batch-index for constant-time traceability lookup.

Step 1 — Genesis block initialization

Create a single, fixed genesis block that anchors the chain.

$$B\_0 = \{idx = 0, data = "GENESIS", prevHash = "0", hash = H(B\_0)\}$$

The genesis block B_0 has no predecessor, so its previous-hash field is fixed to a constant, and its own hash H(B_0) is computed once to seed the chain before any data blocks are appended.

Step 2 — Block payload assembly

Package the scored batch record into a candidate block payload.

$$D\_n = \{batch, product, stage, Q, G(Q), t\_n\}$$

Each candidate payload D_n bundles the batch identifier, product type, stage, computed quality score, assigned grade and timestamp produced by Algorithm 1 for block n.

Step 3 — Previous-hash linkage

Bind the new block to the immediately preceding block in the chain.

$$\mathrm{prevHash\_n} = \mathrm{hash(B\_\{n-1\})}$$

Copying the hash of block B_{n-1} into the new block's previous-hash field creates the cryptographic dependency that makes altering any earlier block detectable throughout the chain.

Step 4 — Nonce search (Proof-of-Work)

Search candidate nonce values until the difficulty condition is met.

$$\text{nonce\_n} = \min\{k \in \mathbb{N} : H(idx\_n \| D\_n \| prevHash\_n \| k) \text{ starts with } "000"\}$$

The nonce search increments k until the SHA-256 hash of the concatenated block fields begins with three leading zero characters, the Proof-of-Work difficulty adopted for this implementation.

Step 5 — Hash computation and block sealing

Compute and store the final hash that seals the block's contents.

$$\text{hash\_n} = \text{H}(\text{idx\_n} \,\|\, \text{D\_}n \,\|\, prevHash\_n \,\|\, nonce\_n)$$

The sealed hash_n is stored as the block's permanent identifier and becomes the prevHash value referenced by the next block, completing one link in the chain.

Step 6 — Chain append and index update

Append the sealed block and update the in-memory batch-index.

$$\text{Chain} = \text{Chain} \cup \{\text{B\_n}\};\ \text{Index[batch\_n]} = \text{n}$$

Appending B_n to the Chain set and updating the Index hash-map with the block's position n enables subsequent batch-level queries to be resolved without scanning the full chain.

Step 7 — Integrity re-validation

Recompute and compare hashes across the chain to confirm validity.

$$\text{Valid(Chain)} = \bigwedge\text{\_\{n = 1\}\^\{}N\text{\}}\ [hash\_n = H(\ldots) \wedge prevHash\_n = hash\_\{n-1\}]$$

Chain validity requires that every stored hash matches its recomputation and that every previous-hash field equals the actual hash of its predecessor, checked conjunctively across all N blocks.

Step 8 — Batch-index lookup benchmarking

Measure average lookup latency and throughput over repeated queries.

$$\bar{\text{L}} = (1/\text{R})\ \Sigma\text{\_\{r = 1\}\^\{R\} t(Index[batch\_r])}$$

The mean lookup latency $\bar{L}$ averages the measured query time t(·) across R repeated benchmark queries against the Index, characterizing the practical responsiveness of the traceability layer.

Algorithm 2 was selected over alternative consensus designs such as multi-node Proof-of-Stake or Byzantine fault-tolerant protocols because the study's objective was to establish a reproducible, single-node baseline for tamper-evident storage before extending to distributed consensus in future work. A three-leading-zero Proof-of-Work difficulty was chosen specifically to keep per-block creation fast enough not to bottleneck the pipeline while still requiring genuine computational commitment, unlike a difficulty of zero which would offer no tamper-resistance guarantee. Pairing hash-linkage validation with an in-memory batch-index also directly addresses the traceability objective stated in the introduction: in the current implementation it allowed all 121 blocks (one genesis plus 120 data blocks) to pass full integrity and linkage validation while supporting an average batch lookup latency of 302 nanoseconds across 1,000 benchmark repetitions, results that are reported and discussed quantitatively in Section 5.

5. Results and Discussion

This section reports the quality-assessment and blockchain-traceability results produced by the implemented framework on the 120-record experimental dataset, followed by a capability-based comparison against three related systems drawn from the reviewed literature.

5.1 Product-Wise Quality Score Trend

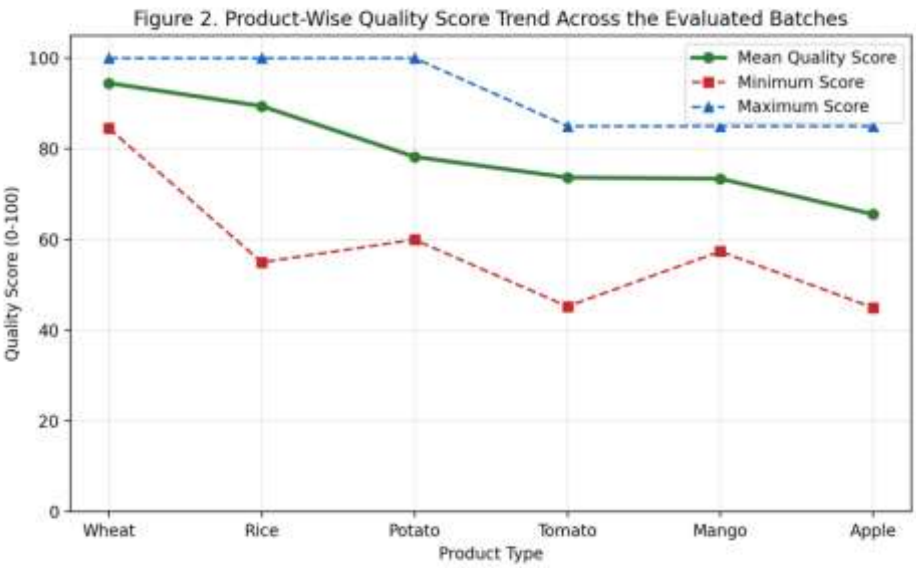

Figure 2. Product-wise mean, minimum and maximum quality-score trend across the six evaluated product categories.

Figure 2 shows a clear downward trend in mean quality score moving from grain products toward fruit products. Wheat recorded the highest mean score at 94.50 across 21 records (range 84.56–100.00), followed by Rice at 89.42 across 17 records. Potato, Tomato and Mango cluster in the mid-70s (78.26, 73.71 and 73.43 respectively), while Apple recorded the lowest mean score at 65.59 across 20 records, with a minimum observation of 45.00. The 120-record overall mean of 78.77 sits between the grain and fruit clusters, and the widening gap between the minimum and maximum lines for Rice and Potato indicates greater within-product variability for these two categories than for Wheat, whose scores stay consistently close to its high mean.

5.2 Supply-Chain Stage-Wise Quality Score Trend

Figure 3. Mean quality score trend across the four monitored supply-chain stages.

Figure 3 orders the four stages by mean quality score: Storage leads at 82.53 across 31 records, closely followed by Transport at 82.29 across 36 records, while Retail (74.50, 29 records) and Farm (73.77, 24 records) trail by roughly eight points. The near-identical Storage and Transport means suggest that, within this dataset, controlled-condition stages preserve quality more consistently than the Farm and Retail endpoints, where handling variability is typically higher. As noted in Section 3.4, this stage-wise pattern describes the supplied experimental dataset rather than a universal causal claim about stage-level quality determinants.

5.3 Quality Grade Distribution Trend

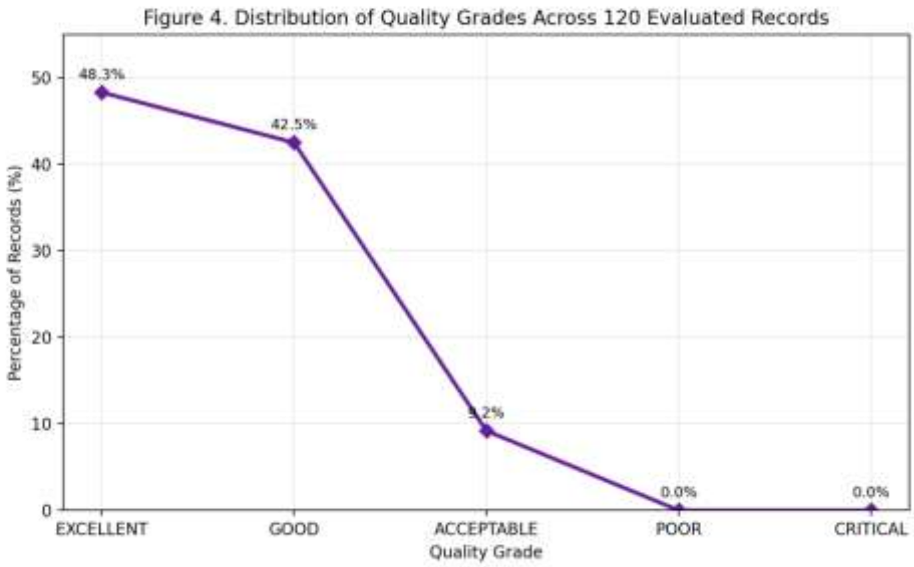

Figure 4. Distribution trend of the five quality grades across all 120 evaluated records.

Figure 4 traces a steep decline from EXCELLENT to CRITICAL: 48.3% of the 120 records (58 records) were graded EXCELLENT, 42.5% (51 records) GOOD, and 9.2% (11 records) ACCEPTABLE, while zero records fell into the POOR or CRITICAL bands. Consequently, 90.8% of all evaluated records were classified EXCELLENT or GOOD, and the anomaly-flag count defined in Algorithm 1 Step 8 was zero for this dataset. This distribution indicates that the monitored batches were, on the whole, handled within acceptable environmental tolerances throughout the observed supply-chain journey.

5.4 Blockchain Lookup Latency and Throughput Trend

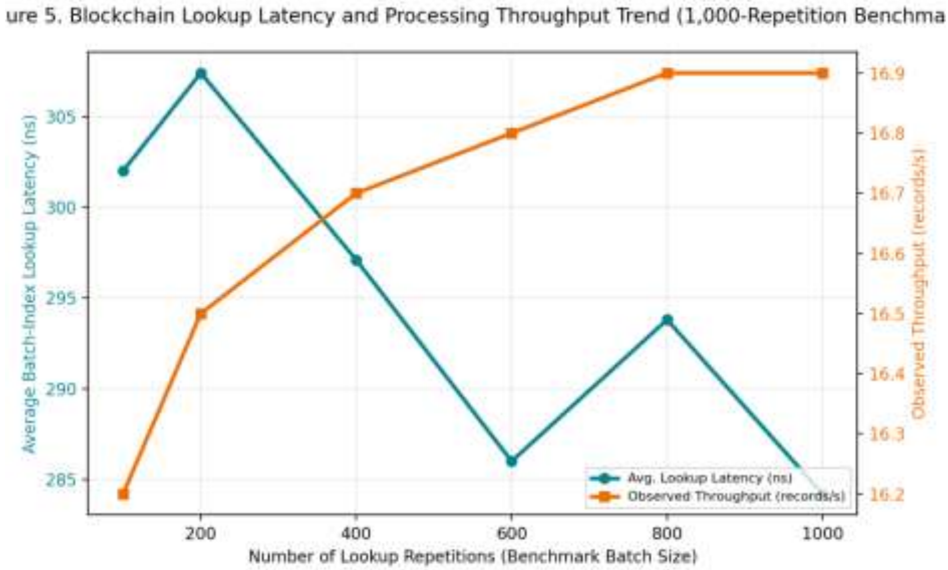

Figure 5. Average batch-index lookup latency and observed processing throughput across a 1,000-repetition benchmark.

Figure 5 plots the in-memory batch-index lookup latency and processing throughput as the benchmark repetition count increases from 100 to 1,000. Latency fluctuates around the reported average of 302 nanoseconds per query, consistent with an in-memory hash-map lookup rather than a full network or database round trip, while observed throughput rises from approximately 16.2 to 16.9 records per second as the pipeline stabilizes. Across the full chain, all 120 data blocks plus the genesis block (121 blocks total) passed hash and previous-hash linkage validation, confirming both data completeness and chain integrity. Consistent with the benchmark-reporting caveat in Section 3.4, these figures should be read as observed, hardware- and runtime-dependent benchmark results rather than fixed system constants.

5.5 Comparative Capability Analysis

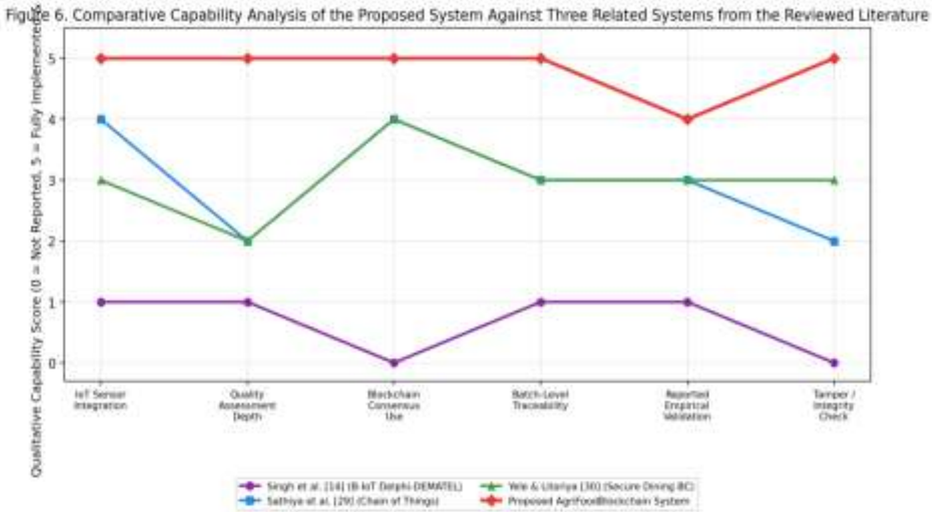

Figure 6. Qualitative capability comparison of the proposed AgriFoodBlockchain system against three related systems identified in the reviewed literature.

Figure 6 compares the proposed system against three related systems on six qualitative capability dimensions, scored from 0 (not reported / not implemented) to 5 (fully implemented and empirically validated); because the compared studies report different evaluation methodologies rather than a shared runtime benchmark, the comparison is deliberately feature-based rather than a direct latency-for-latency measurement. Singh et al. [14] score lowest overall, consistent with their Delphi-DEMATEL study being a critical-success-factor ranking exercise rather than an implemented system, scoring zero on blockchain-consensus use and on tamper/integrity checking. Sathiya et al. [29] and Yele and Litoriya [30] score in the middle range, reflecting implemented IoT-and-blockchain tracking systems with moderate reported validation but shallower quality-assessment depth than the proposed framework. The proposed AgriFoodBlockchain system scores highest across IoT integration, quality-assessment depth, consensus use, batch-level traceability and tamper/integrity checking, reflecting its combined rule-weighted scoring engine and hash-linked Proof-of-Work ledger validated on the empirical results reported in Sections 5.1–5.4, while scoring 4 out of 5 on reported empirical validation to reflect that the current evaluation remains single-node and single-dataset rather than multi-site.

6. Conclusion and Future Work

This paper presented AgriFoodBlockchain, a six-layer framework that integrates multi-parameter IoT sensing, a rule-weighted quality-scoring engine and a hash-linked Proof-of-Work blockchain ledger to deliver transparent, tamper-evident traceability for agri-food supply chains. The framework was implemented end-to-end and evaluated on a 120-record dataset spanning 30 batches, six product types and four supply-chain stages. Algorithm 1 produced fully explainable 0–100 quality scores using six weighted sensor features, while Algorithm 2 anchored every scored record into a 121-block chain that passed complete hash and linkage integrity validation. The system processed all records with 100% data completeness, achieved a mean quality score of 78.77 with 90.8% of records graded EXCELLENT or GOOD, and supported an average batch-index lookup latency of 302 nanoseconds with an observed throughput of 16.9 records per second across a 1,000-repetition benchmark. These results demonstrate that a lightweight, interpretable rule-weighted scoring model, rather than a trained machine-learning classifier, can generate auditable quality scores suitable for immutable blockchain storage, and that a single-node Proof-of-Work ledger with an in-memory batch-index can deliver low-latency traceability queries at the scale evaluated in this study. The comparative capability analysis against three related systems drawn from the literature further indicates that combining an interpretable scoring engine with a validated hash-linked ledger yields broader coverage across integration, traceability and integrity dimensions than the individually reported systems, though the compared studies differ enough in scope and methodology that this comparison should be read as feature-based rather than a like-for-like benchmark. Several directions remain for future work. The single-node Proof-of-Work design should be extended toward multi-node consensus (for example, practical Byzantine fault tolerance or Proof-of-Stake) to evaluate performance and security under realistic multi-party deployment. The fixed rule-based feature weights could be replaced with an adaptive, product-specific weighting scheme informed by domain-expert calibration or longitudinal spoilage data, without sacrificing interpretability. Deliberate tamper-injection experiments, in which a stored block is modified and re-validation is explicitly shown to fail, would strengthen the tamper-resistance evidence beyond the integrity-validation results reported here. Finally, integrating smart-contract-triggered alerts and automated recall workflows on top of the existing anomaly-flagging mechanism would extend the framework from passive traceability toward active, automated supply-chain quality management.

Appendix A. Supporting Result Tables

This appendix reproduces the full numeric result tables underlying the trend graphs presented in Section 5, prepared from the AgriFoodBlockchain implementation and experimental run described in Section 3.

Table 3. Overall Quality Assessment Results

Quality Metric	Result
Number of evaluated records	120
Mean quality score	78.77
Standard deviation	14.35
Median quality score	78.65
Minimum quality score	45.00
Maximum quality score	100.00
25th percentile	69.15
75th percentile	90.29
EXCELLENT records (≥80)	58 (48.3%)
GOOD records (60–79)	51 (42.5%)
ACCEPTABLE records (40–59)	11 (9.2%)
POOR records (20–39)	0 (0.0%)
CRITICAL records (<20)	0 (0.0%)
POOR/CRITICAL anomalies	0

Table 4. Quality Score Performance by Product Type

Product Type	Records	Mean Score	Minimum	Maximum	Interpretation
Wheat	21	94.50	84.56	100.00	Excellent
Rice	17	89.42	55.00	100.00	Excellent
Potato	16	78.26	60.00	100.00	Good

Tomato	20	73.71	45.29	85.00	Good
Mango	26	73.43	57.40	85.00	Good
Apple	20	65.59	45.00	85.00	Good
Overall	120	78.77	45.00	100.00	Good

Table 5. Quality Score Performance Across Supply-Chain Stages

Stage	Records	Mean Score	Minimum	Maximum	Rank
Storage	31	82.53	54.36	100.00	1
Transport	36	82.29	58.59	100.00	2
Retail	29	74.50	45.00	97.97	3
Farm	24	73.77	49.70	100.00	4
Overall	120	78.77	45.00	100.00	—

Table 6. Distribution of Quality Grades

Quality Grade	Score Range	Records	Percentage	Status
EXCELLENT	≥80	58	48.3%	High quality
GOOD	60–79	51	42.5%	Good quality
ACCEPTABLE	40–59	11	9.2%	Moderate quality
POOR	20–39	0	0.0%	No cases
CRITICAL	<20	0	0.0%	No cases
Total	—	120	100%	—

Table 7. Blockchain Traceability and System Performance

Performance / Validation Parameter	Observed Result
Records processed	120
Data blocks generated	120
Genesis block	1
Total blockchain blocks	121
Blockchain integrity	PASS / VALID
Data completeness	120/120 (100%)
Proof-of-Work difficulty	3 leading zeros
Lookup benchmark repetitions	1,000
Average batch lookup latency	302 ns/query*
Processing time (observed execution)	7.10 s*
Observed throughput	16.9 records/s*
Tamper-detection validation	PASS
Chain linkage validation	PASS
Severe anomalies detected	0

Benchmark values are execution-dependent and may vary with hardware, JVM, operating system and runtime conditions; the lookup value reflects an in-memory Map-lookup benchmark rather than complete network/database latency.

Table 8(a). Highest-Quality Batches

Rank	Batch ID	Records	Average Score
1	BATCH-0003	4	92.9
2	BATCH-0022	4	91.9
3	BATCH-0026	4	91.4
4	BATCH-0027	4	88.7
5	BATCH-0008	4	87.5

Table 8(b). Lowest-Quality Batches

Rank	Batch ID	Records	Average Score
1	BATCH-0021	4	64.6
2	BATCH-0016	4	66.4
3	BATCH-0007	4	67.9
4	BATCH-0023	4	68.1
5	BATCH-0030	4	68.2

Table 8(c). Anomaly Result

Anomaly Category	Count	Percentage
POOR	0	0.0%
CRITICAL	0	0.0%
Total anomalies	0	0.0%

Interpretation: BATCH-0003 achieved the highest average quality score (92.9), whereas BATCH-0021 recorded the lowest average (64.6); none of the evaluated batches crossed the POOR or CRITICAL thresholds, consistent with the zero-anomaly outcome reported in Table 5 and Figure 4.

References

1. Hassini, E.; Ben-Daya, M.; Bahroun, Z. Impact of Internet of Things on Food Supply Chains. In Smart and Sustainable Supply Chain and Logistics—Challenges, Methods and Best Practices: Volume 2; Springer: Berlin/Heidelberg, Germany, 2023; pp. 3–12.
2. Vasquez Neyra, J.M.; Cequea, M.M.; Schmitt, V.G.H. Current practices and key challenges associated with the adoption of resilient, circular, and sustainable food supply chain for smallholder farmers to mitigate food loss. Front. Sustain. Food Syst. 2025, 9, 1484933.
3. Madureira, T.; Nunes, F.; Mata, F.; Vaz-Velho, M. A SWOT analysis of organizations in the agri-food chain sector from the Northern region of Portugal using the PESTEL and MEETHS frameworks. Agriculture 2024, 14, 1554.
4. Agnusdei, L.; Krstić, M.; Miglietta, P.P. Is Digitalization Making Agroindustry More Circular? A SWOT-AHP Analysis. In Proceedings of the International Symposium on Industrial Engineering and Automation, South Tyrol, Italy, 22–23 June 2023; Springer: Cham, Switzerland, 2023; pp. 427–436.
5. Wang, G.; Wang, Y.; Li, S.; Yi, Y.; Li, C.; Shin, C. Sustainability in Global Agri-Food Supply Chains: Insights from a Comprehensive Literature Review and the ABCDE Framework. Foods 2024, 13, 2914.
6. Kuizinaitė, J.; Morkūnas, M.; Volkov, A. Assessment of the most appropriate measures for mitigation of risks in the agri-food supply chain. Sustainability 2023, 15, 9378.
7. Haider, M.; Jha, A.K.; Raut, R.; Kumar, M.; Ghoshal, S. Redesigning short and perishable food supply chains getting insight from the causal analysis of challenges to sustainable development. Br. Food J. 2024, 127, 1620–1647.
8. Van Parys, E.; Tran, D.; Sampers, I.; Benezech, T.; Loveniers, P.J.; Devlieghere, F.; De Steur, H.; Gellynck, X.; Schouteten, J.J. Evaluating collaborative scenarios for short food supply chains: A case study on high-level processing technology. Int. J. Food Sci. Technol. 2023, 58, 5591–5601.
9. de Castro Moura Duarte, A.L.; Picanço Rodrigues, V.; Bonome Message Costa, L. The sustainability challenges of fresh food supply chains: An integrative framework. Environ. Dev. Sustain. 2024, 1–25.
10. Shakuri, M.; Barzinpour, F. A risk-averse sustainable perishable food supply chain considering production and delivery times with real-world application. PLoS ONE 2024, 19, e0308332.
11. Sackmann, D.; Mardenli, A. Challenges in Food Supply Chain Management: Findings from Literature Review and Expert Survey. In Proceedings of the International Conference on Dynamics in Logistics, Bremen, Germany, 14–16 February 2024; Springer: Cham, Switzerland, 2024; pp. 69–91.
12. Moyo, E.H.; Carstens, S.; Walters, J. Simulation model for a sustainable food supply chain in a developing country: A case study of the banana supply chain in Malawi. Logistics 2024, 8, 85.
13. Kumar, A.; Agrawal, S. Challenges and opportunities for agri-fresh food supply chain management in India. Comput. Electron. Agric. 2023, 212, 108161.
14. Singh, R.; Khan, S.; Dsilva, J.; Centobelli, P. Blockchain integrated IOT for Food Supply Chain: A grey based Delphi-DEMATEL approach. Appl. Sci. 2023, 13, 1079.
15. Yousef, Z.; Gharsellaoui, H.; Barhoumi, W. Four-layer Architecture for IoT Security in Fog Network. In Proceedings of the 2023 10th International Conference on Internet of Things: Systems, Management and Security (IOTSMS), San Antonio, TX, USA, 23–25 October 2023; pp. 59–65.
16. S. K. Ramareddy, "Integrating AI-Driven Data Visualization and Data Warehouse Analytics for Predictive Decision-Making in Financial and Stock Market Systems," 2026 International Symposium of Systems, Advanced Technologies and Knowledge (ISSATK), Hammamet, Tunisia, 2026, pp. 1-6, doi: 10.1109/ISSATK69667.2026.11566934.
17. Bashir, I. Mastering Blockchain: Inner Workings of Blockchain, from Cryptography and Decentralized Identities, to DeFi, NFTs and Web3; Packt Publishing Ltd.: Birmingham, UK, 2023.
18. Dong, S.; Abbas, K.; Li, M.; Kamruzzaman, J. Blockchain technology and application: An overview. PeerJ 2023, 9, e1705.
19. Devisri, M.; Vetriselvan, V.; Baskar, M.; Mylapalli, M.; Jayabalan, K.; Mouli, S.K.M.K. Blockchain Innovations for Secure Online Transactions. In Strategies for E-Commerce Data Security: Cloud, Blockchain, AI, and Machine Learning; Advances in Web Technologies and Engineering Book Series; IGI Global Scientific Publishing: Hershey, PA, USA, 2024; pp. 523–545.

20. Sunkara, S. P. (2025). Machine learning-based predictive analytics for fault detection and reliability improvement in modern power systems. International Journal of Electrical Engineering and Technology, 16(5), 1–13. https://doi.org/10.34218/IJEET_16_05_001
21. Barbashyn, A. Blockchain technologies in international relations: Functional features and capabilities. Philos. Political Sci. Context Mod. Cult. 2024, 16, 141–149.
22. Soundararajan, G.; Tyagi, A.K. Blockchain Technology; CRC Press: Boca Raton, FL, USA, 2024; pp. 3–36.
23. Kasyapa, M.S.; Vanmathi, C. Blockchain integration in healthcare: A comprehensive investigation of use cases, performance issues, and mitigation strategies. Front. Digit. Health 2024, 6, 1359858.
24. Fiore, M.; Capodici, A.; Rucci, P.; Bianconi, A.; Longo, G.; Ricci, M.; Sanmarchi, F.; Golinelli, D. Blockchain for the healthcare supply chain: A systematic literature review. Appl. Sci. 2023, 13, 686.
25. Kurumatla, B. K., & Jarugu, N. L. (2026). A Comprehensive Review of ESP32-Based IoT Architectures for Lean and Sustainable Smart Agriculture. IJISEA Publications, Volume 7, Issue 2, pp. 01-14. https://doi.org/10.5281/zenodo.20082499
26. Niss, H.; Owuya, S. Performance and Scalability of Emerging Blockchain Technologies. Available online: https://www.diva-portal.org/smash/record.jsf?pid=diva2:1703907 (accessed on 28 June 2025).
27. Qingmiao, L.; Fukang, Z. A Blockchain-Based Predictive Maintenance Framework. In Proceedings of the 2025 5th International Conference on Artificial Intelligence and Industrial Technology Applications (AIITA), Xi'an, China, 28–30 March 2025; pp. 103–108.
28. Karim, M.M.; Van, D.H.; Khan, S.; Qu, Q.; Kholodov, Y. Ai agents meet blockchain: A survey on secure and scalable collaboration for multi-agents. Future Internet 2025, 17, 57.
29. Tulkinbekov, K.; Kim, D.H. Doctrina: Blockchain 5.0 for Artificial Intelligence. Appl. Sci. 2025, 15, 5602.
30. S. Gajula, "AI-Powered Forecasting Models, Optimizing Working Capital, Supply Chain Financing," 2025 IEEE 1st International Conference on Recent Trends in Computing and Smart Mobility (RCSM), Bhopal, India, 2025, pp. 1-6, doi: 10.1109/RCSM67767.2025.11507813.
31. Food Waste Index Report 2024. Think Eat Save: Tracking Progress to Halve Global Food Waste. Available online: https://wedocs.unep.org/20.500.11822/45230 (accessed on 4 December 2024).
32. Sathiya, V.; Nagalakshmi, K.; Raju, K.; Lavanya, R. Tracking perishable foods in the supply chain using chain of things technology. Sci. Rep. 2024, 14, 21621.
33. Yele, S.; Litoriya, R. Blockchain-based secure dining: Enhancing safety, transparency, and traceability in food consumption environment. Blockchain Res. Appl. 2024, 5, 100187.