

Date of publication xxxx 00, 0000, date of current version xxxx 00, 0000.

Digital Object Identifier xx.xxxx/ACCESS.xxxx.DOI

The Quantum-Enhanced Agri-Ledger (QAL): A Framework for Preemptive Sensing and Verifiable Sustainability

A HARSHA KUMAR¹, N SARAN¹, KUMARAN K¹, SARANYA G¹, DISHA DANIEL², AND HIMESWAR POTNURU²

¹School of Computer Science and Engineering, Vellore Institute of Technology, Chennai, India (e-mail: ahkharsha@gmail.com; n.sarancs@gmail.com; kumaran.k@vit.ac.in; saranya.g@vit.ac.in)

²School of Electronics Engineering, Vellore Institute of Technology, Chennai, India (e-mail: dishadaniel24@gmail.com; himesh2103@gmail.com)

Corresponding author: Kumaran K (e-mail: kumaran.k@vit.ac.in).

ABSTRACT The confluence of climate change and population growth poses unprecedented challenges to global food security. Existing smart agriculture solutions often rely on lagging indicators of crop stress and lack robust mechanisms for verifying sustainable practices. This paper introduces the Quantum-Enhanced Agri-Ledger (QAL), a framework designed to shift agricultural management from a reactive to a preemptive paradigm. Its potential is evaluated through a comprehensive simulation study using a highly complex, multi-variate synthetic dataset. The framework integrates three innovations: (1) a proposed high-sensitivity Quantum Dot Spectrometry Sensor (QDSS) model for in-situ detection of plant stress volatile organic compounds (VOCs); (2) a privacy-preserving Federated Learning (FL) model that leverages high-dimensional sensor data to predict crop health; and (3) a novel blockchain consensus mechanism, Dynamic Proof-of-Stake with Sustainability Slashing (dPoS-SS), to create an immutable ledger for sustainable farming actions. Simulation results, aggregated over multiple independent runs, show that the full QAL model, featuring a feed-forward neural network, achieves a mean stress classification accuracy of **96.71% \pm 0.47%**. While the high-complexity data model narrowed the gap in yield prediction, the QAL model still demonstrated a lower Root Mean Square Error (RMSE) of **1.31 \pm 0.20 tons/ha** compared to baseline conventional models. Similarly, the dPoS-SS model suggests substantially higher energy efficiency than traditional Proof-of-Work systems. The QAL framework is presented as a blueprint for a future generation of transparent, incentive-driven ecosystems for a sustainable and secure food future.

INDEX TERMS Precision Agriculture, Quantum Dot Sensor, Federated Learning, Blockchain, Preemptive Intervention, Sustainable Farming, Crop Stress Phenotyping

I. INTRODUCTION

THE agricultural sector is at a critical juncture, tasked with ensuring global food security for a population projected to reach 9.7 billion by 2050, a challenge requiring an estimated 70% increase in food production [1]. This task is compounded by the escalating pressures of climate volatility, soil degradation, and water scarcity. In developing nations, particularly India where agriculture supports over 58% of the population, these challenges are acutely felt. While visions for integrating the Internet of Things (IoT), Artificial Intelligence (AI), and blockchain technology are compelling, a significant gap persists between concept and field-ready implementation. This is particularly evident in the need for

sensors capable of real-time, molecular-level monitoring and a secure framework for verifying and incentivizing sustainable on-farm practices.

Current smart farming solutions predominantly rely on sensors that measure bulk environmental properties like temperature (T), humidity (H), and soil macronutrients. These metrics are often **lagging indicators of plant stress**; by the time a measurable change is detected, irreversible physiological damage may have occurred. Plants, however, emit dynamic profiles of Volatile Organic Compounds (VOCs) as an early warning of biotic and abiotic stress [2]. The ability to detect and decode these VOC signatures offers a crucial window for **preemptive intervention**, shifting the

paradigm from reactive damage control to proactive health management.

To bridge this data gap, we introduce the **Quantum-Enhanced Agri-Ledger (QAL)**, a multi-layered framework designed to shift agricultural management from a reactive to a preemptive paradigm. This paper presents the conceptual framework and a comprehensive simulation-based validation of the QAL system. It integrates a theoretical model of a Quantum Dot Spectrometry Sensor (QDSS) with a simulated privacy-preserving Federated Learning architecture to explore the potential of this new economic model.

This proactive approach supports progress toward United Nations Sustainable Development Goal (SDG) 2 (Zero Hunger) by enabling precise yield forecasting to secure food supplies. Furthermore, the data-driven optimization of chemical inputs aids in SDG 12 (Responsible Consumption and Production) by reducing the environmental footprint of industrial agriculture. The primary contributions are threefold:

- 1) **A Quantum Dot Spectrometry Sensor (QDSS Model):** A model for a low-cost, field-deployable sensor array is proposed. This model uses specifically functionalized quantum dots, tuned to the spectral lines of key stress-related VOCs and soil micronutrients, to generate high-dimensional data with high sensitivity and specificity.
- 2) **A Privacy-Preserving Federated Learning Architecture:** A machine learning pipeline is designed to use the high-dimensional data from the QDSS model for predicting crop stress and yield. This **Federated Stress-Phenotyping Model (FSPM)** is trained using a Federated Learning (FL) architecture to preserve **data sovereignty**, allowing farmers to benefit from a collectively improved global model without sharing their sensitive operational data.
- 3) **A Novel Consensus Mechanism for Verifiable Sustainability:** A concept called Proof-of-Sustainable-Practice (PoSP) is introduced, implemented via a mechanism named **Dynamic Proof-of-Stake with Sustainability Slashing (dPoS-SS)**. This mechanism is designed to enable farmers to create immutable records of sustainable actions, which are then validated using sensor data and rewarded, creating a direct financial incentive for sustainability.

This paper is organized as follows: Section II reviews related work. Section III presents the architecture of the QAL framework and the basis of the QDSS model. Section IV details the core algorithms. Section V describes the simulation environment used to evaluate the framework. Section VI presents and discusses the simulation results, and Section VII concludes with a summary, a discussion of limitations, and directions for future research.

II. RELATED WORK

The QAL framework represents a synthesis of advances in sensing, distributed intelligence, and decentralized verification. This review critically evaluates the current state-of-the-

art across these domains to pinpoint unresolved limitations that motivate the design of the integrated QAL architecture. The following subsections analyze existing work, with comparative summaries in Tables 1, 2, and 3.

A. THE SENSING FRONTIER: FROM PROXIES TO PREEMPTIVE DIAGNOSTICS

The foundation of modern precision agriculture rests on data acquisition. While the proliferation of IoT devices [3], [4] and advanced remote sensing techniques like UAV-based hyperspectral imaging [5] have enabled detailed field-scale mapping, their utility is fundamentally constrained. As Mahlein describes [6], many imaging methods detect disease only after physiological changes are well underway. **The critical gap remains their reliance on lagging indicators.** Whether it is a shift in canopy reflectance or soil moisture, the signal indicates that stress has already occurred. The frontier for preemptive action lies in detecting biochemical precursors like VOCs. While electronic-nose technologies have been explored [7], [8], they often lack the field-deployable specificity needed to deconvolve complex VOC mixtures into actionable stress signatures, a limitation the QDSS model is designed to address.

Recent integrations of UAV-mounted diffraction spectroscopy and fiber Bragg grating sensors [9] have advanced real-time forest health monitoring. While effective for physiological strain and canopy reflectance, our work extends this domain by targeting specific volatile organic compounds (VOCs) via quantum dot sensing for pre-symptomatic stress detection. Furthermore, recent work on reconfigurable holographic surfaces [10] highlights that hardware impairments and phase shift errors are inevitable in practical wireless links. This reinforces the QAL strategy of performing heavy data processing at the Edge Layer, minimizing the reliance on pristine uplink conditions for raw data transmission.

B. THE INTELLIGENCE LAYER: FROM CENTRALIZED MODELS TO DATA-RICH FEDERATED LEARNING

The application of AI in agriculture is well-documented [11], but most deep learning models rely on centralized data, creating significant privacy and data ownership risks for farmers. Federated Learning (FL) [12], [13] provides a powerful architectural solution to this problem. However, the performance of any AI, federated or not, is capped by its input data. **The unresolved limitation of current FL applications in agriculture is not the architecture but the data quality.** Even advanced multi-modal networks [14] are constrained by data from conventional sensors that capture proxies of plant health, not the underlying biochemical state. The FSPM aims to address this data-quality gap by pairing the FL architecture with a high-fidelity data source.

C. THE VERIFICATION BACKBONE: FROM PASSIVE LEDGERS TO ACTIVE INCENTIVE ENGINES

Blockchain technology has been proposed to enhance agri-food traceability [15]–[17]. The evolution from energy-

intensive Proof-of-Work (PoW) [18] to efficient Proof-of-Stake (PoS) [19] has made it more sustainable. However, the primary unresolved challenge is the "oracle problem" [20]: reliably and securely feeding real-world data onto the blockchain. Existing systems often function as passive ledgers reliant on manual or centralized data entry, which is vulnerable to error or manipulation. They record claims about practices but cannot autonomously verify them. The dPoS-SS mechanism is conceptualized to address this by integrating cryptographically signed sensor data directly into the consensus process, transforming the blockchain from a passive recorder into an active verification engine.

Beyond basic PoS consensus, recent research has focused on incentive-compatible designs for resource-constrained and UAV-assisted networks. Tang et al. [21] propose a PoS blockchain for secured data collection in UAV-assisted IoT, where a multi-agent reinforcement learning framework optimizes stake investment and profit sharing. Complementary surveys on blockchain-enabled UAV systems [22] and comparative analyses of consensus mechanisms [23] further highlight the importance of carefully designing staking, rewards, and slashing rules to ensure both security and economic sustainability.

To contextualize the QAL framework within the broader ecosystem of existing agricultural technologies, **Table 4** benchmarks the proposed system against current commercial UAV and IoT platforms. This comparison highlights the structural shift from reactive monitoring to the preemptive, molecular-level diagnostics proposed in this work.

As illustrated, while commercial platforms excel in spatial mapping or bulk soil monitoring, they predominantly rely on lagging indicators. The QAL framework distinguishes itself by targeting the biochemical precursors of stress, enabling intervention before visible symptoms manifest.

D. SUMMARY OF FRAMEWORK CONTRIBUTIONS

As detailed in Tables 1, 2, and 3, the QAL framework is designed to address fundamental limitations across sensing, intelligence, and verification domains. The integrated approach aims to shift agriculture from reactive to proactive monitoring, pursue superior AI performance while maintaining data privacy, and create a blockchain system that helps mitigate the oracle problem for on-farm verification. This synergy establishes a paradigm for sustainable precision agriculture whose potential is explored in this work.

III. SYSTEM ARCHITECTURE

The QAL framework is a multi-layered, decentralized system designed to integrate field data with a robust, transparent, and incentive-driven digital ecosystem. The modularity of the framework is visualized in the four-layer architecture depicted in Figure 1. While Figure 1 illustrates the high-level architecture, the specific operational logic is defined mathematically: the sensing physics follows the modified Stern-Volmer model (Eq. 1), the decision logic is governed by the Federated Learning protocol detailed in Algorithm

2, and the verification rules are enforced by the dPoS-SS Consensus in Algorithm 3.

A. PERCEPTION LAYER

This foundational layer is responsible for in-situ data acquisition and forms the core of the framework's IoT network. It is comprised of a distributed network of autonomous, low-power IoT sensing nodes that provide a comprehensive, real-time view of plant health and environmental conditions. A key component of this layer is the QDSS Array, which is designed to detect specific VOCs and soil micronutrients. This provides a high-dimensional spectral fingerprint of plant stress, enabling early detection of stressors before visual symptoms appear. Complementing this, the IoT nodes are also equipped with traditional environmental sensors to measure bulk parameters such as soil moisture, temperature, humidity, and soil macronutrients (NPK), providing crucial environmental context.

B. EDGE LAYER

The Edge Layer serves as the on-farm computational gateway, a crucial component of the IoT architecture that enables edge computing. Acting as an intermediary between the sensor nodes and the cloud, its primary functions are real-time data pre-processing and privacy preservation. The Edge Gateway Device aggregates high-volume data streams from the Perception Layer. By performing initial cleaning, normalization, compression and encryption locally, it reduces the amount of data transmitted to the cloud, lowering bandwidth costs and latency. Most importantly, this device hosts the local client for the FL model, ensuring that raw, sensitive farm data never leaves the premises.

C. PLATFORM LAYER

The Platform Layer is the cloud-based backend that orchestrates the system's analytical and decision-making processes. The FL Server manages the core machine learning pipeline. It sends a global model to Edge Gateways for local training and aggregates only the model updates from participating farms to refine the global model. To ensure the integrity of the blockchain, the system relies on External Oracle Services, which provide independent data from sources like weather APIs and satellite imagery for validating on-chain claims.

D. BLOCKCHAIN LAYER

This decentralized, immutable ledger forms the backbone of the QAL framework. The QAL Blockchain Network utilizes the PoSP consensus mechanism, implemented as dPoS-SS, to immutably record verified sustainable actions. These actions trigger Smart Contracts which automatically issue tokenized rewards to farmers upon successful validation. A decentralized network of Validators verifies claims by comparing submitted data against trusted external oracle data, achieving consensus through a trust-weighted voting process that is both energy-efficient and secure. As an illustrative funding model, the token economy could potentially be supported

Table 1: Comparative Analysis of Sensing Technologies in Agriculture

Feature	Hyperspectral Imaging [5]	IoT Environmental Sensors [3]	E-Nose Systems [7]	QAL QDSS (Proposed)
Detection Timing	Post-symptom	Post-symptom	Limited Pre-symptom	Design Goal: Pre-symptom detection
Diagnostic Precision	Spatial mapping	Indirect indicators	Low specificity	Design Goal: Molecular-level fingerprinting
Field Deployment	UAV-based	Stationary nodes	Limited field use	Design Goal: Low-cost, edge-ready
Validation Status / Limitations	High cost, requires clear atmospheric conditions	Lagging indicator, subject to placement error	Prone to drift and environmental interference	Simulated model, not physically validated

Table 2: Comparative Analysis of AI Systems in Precision Agriculture

Feature	Centralized Deep Learning [11]	Federated Learning [12]	Multi-modal Networks [14]	QAL FSPM (Proposed)
Data Privacy	High-risk	Privacy-preserving	Centralized risk	Privacy-by-design architecture
Model Performance	High accuracy but data-limited	Good performance, constrained by setting	Advanced fusion but proxy-data limited	Simulated Potential: Superior accuracy via molecular data
Scalability	Requires central infrastructure	Distributed computation	High computational demands	Edge-optimized design
Validation Status / Limitations	Requires central data ownership, high trust	Complex aggregation, potential for non-IID data issues	Computationally expensive, centralized	Performance is contingent on hypothetical sensor data

Table 3: Comparative Analysis of Blockchain Systems for Agricultural Verification

Feature	Supply Chain Traceability [15]	Standard PoS with Oracles [19]	PoW with Oracles [18]	QAL dPoS-SS (Proposed)
Functionality	Passive ledger	Energy-efficient consensus	Secure but high energy cost	Active incentive engine
Trust Model	Trust in data entry	Oracle-dependent	Trust in hash rate	Integrated trust via sensors
Incentive Structure	Transaction fees	Staking rewards	Mining rewards	Sustainability rewards
Validation Status / Limitations	Vulnerable to erroneous data entry, dependent on input fidelity	Centralization risk in oracle providers	Extremely high energy consumption	Conceptual, dependent on reliable sensor data as oracles

Table 4: Comparison of QAL with Commercial and IoT Agricultural Platforms

Feature	Commercial Multispectral UAV (e.g., DJI P4 / Plant-O-Meter [24])	Commercial IoT Soil Sensors (e.g., Monnit Wireless [25])	Proposed QAL Framework
Detection Basis	Canopy Reflectance (NDVI/NDRE)	Bulk Dielectric Permittivity (Moisture)	Specific Biochemicals (VOCs)
Timing	Reactive (Visual/Spectral change)	Reactive (Moisture deficit)	Preemptive (Metabolic shift)
Data Ownership	Often Vendor-Cloud Locked	Centralized Cloud Dashboard	Federated (On-premise)
Verification	Brand Trust	None	Blockchain Consensus

by supply chain premiums for certified sustainable produce and the tokenization of verified carbon credits, though a full economic validation is outside the scope of this technical study.

E. THE QUANTUM DOT SPECTROMETRY SENSOR (QDSS) MODEL

The core sensing innovation of the framework is the QDSS. A quantum dot (QD) is a semiconductor nanocrystal whose excitons are confined in all three spatial dimensions. This property, known as **quantum confinement**, allows the dot's emission energy and corresponding color to be precisely tuned by controlling its physical size. This tunability is the foundation for creating an array of sensors, each highly specific to the spectral absorption of different target molecules.

The sensor's operating principle, shown in Figure 2, relies on **fluorescence quenching**. Each QD has a baseline fluores-

cence efficiency, or quantum yield. When target molecules, such as VOCs, bind to the functionalized surface of a QD, they introduce new pathways for non-radiative energy decay. This process effectively "quenches" or dims the dot's fluorescence. The degree of quenching is proportional to the concentration of the target analyte, providing a measurable signal.

While this study relies on simulation, the sensor model is informed by candidate materials reported in the literature. Specifically, sol-gel encapsulation of polymer-coated ZnS/CdSe quantum dots has been demonstrated to effectively detect organic vapors through fluorescence quenching [26]. Similarly, core-shell CdSe/ZnS nanocrystals have shown high sensitivity in electrochemical sensor applications for phenolic compounds [27], providing a material basis for the proposed array design.

For the simulation, the total quenching effect in a multi-

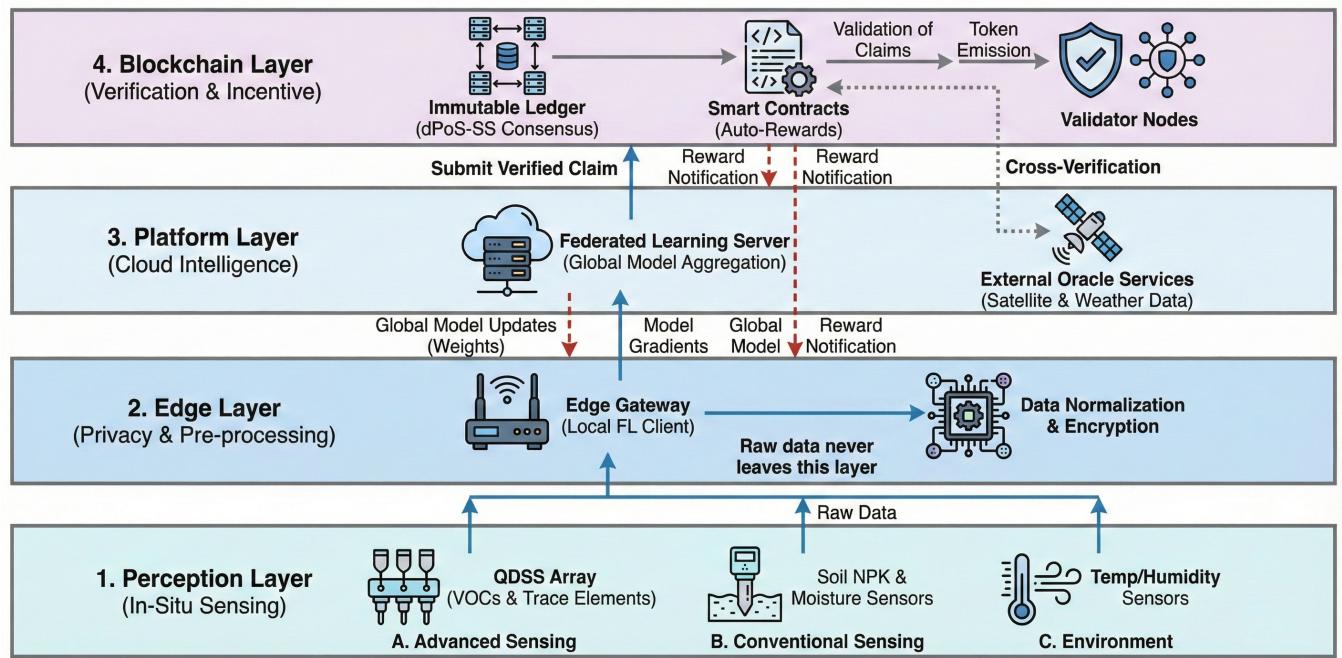


Figure 1: An overview of the QAL four-layer architecture. The diagram illustrates the data pipeline, starting with in-situ data capture at the Perception Layer. Data is then pre-processed at the Edge Layer to preserve privacy. The cloud-based Platform Layer aggregates insights using Federated Learning, and the Blockchain Layer provides an immutable ledger for verifying and incentivizing sustainable practices.

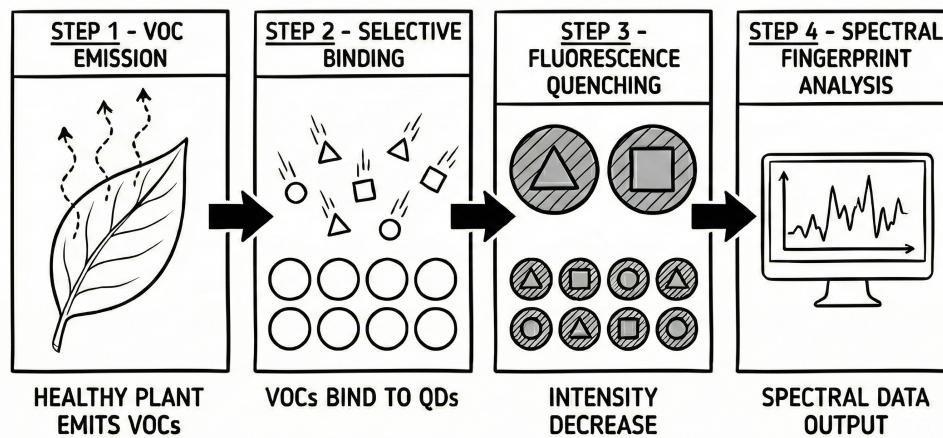


Figure 2: Diagram illustrating the four-step working principle of the QDSS. (1) A stressed plant emits specific VOCs. (2) VOCs selectively bind to the functionalized surface of the quantum dots. (3) This binding process increases non-radiative decay, causing the fluorescence to 'quench'. (4) The specific pattern of quenching across the array creates a unique spectral fingerprint for analysis.

analyte environment is modeled using a modified Stern-Volmer equation (Eq. 1). This provides a mathematical basis for generating a unique "spectral fingerprint" from a mixture of quenching analytes ($[Q_i]$).

$$\frac{F_0}{F} = \left(1 + \sum_{i=1}^N K_{SV,i} [Q_i] \right) \quad (1)$$

It is critical, however, to acknowledge the limitations of this model. While the standard Stern-Volmer equation (Eq. 1) is a linear first-order approximation, our simulation utilizes a hyperbolic tangent function (see Eq. 8 in Section V) to qualitatively capture the non-linear saturation and inner-filter effects often observed in complex VOC mixtures, ensuring the model does not assume infinite sensitivity. Furthermore, creating a robust, field-deployable QDSS presents several challenges:

- **Long-term Stability:** QDs can suffer from degradation and photobleaching when exposed to UV light and harsh environmental conditions, requiring advanced encapsulation techniques to ensure a stable signal over a growing season.
- **Environmental Interference:** The fluorescence of QDs is highly sensitive to temperature and humidity. A physical sensor would require integrated thermal management and calibration algorithms to decorrelate the environmental effects from the target VOC signal.
- **Cross-Sensitivity and Specificity:** While QDs can be functionalized for specific analytes, achieving perfect specificity in a complex blend of dozens of VOCs is a major challenge. Cross-reactivity could lead to ambiguous signals requiring sophisticated deconvolution algorithms.
- **Power Consumption and Scalability:** For wide-scale agricultural deployment, the sensor nodes must be ultra-low-power. Furthermore, mass-producing quantum dots with consistent size, shape, and surface chemistry to ensure uniform performance is a significant manufacturing hurdle.

Table 5 outlines the proposed composition of the sensor array used in the model.

Table 5: Proposed QDSS Array Composition and Target Analytes.

QD Type	Core/Shell Material	Target Analyte (VOC)	Associated Stressor
Dot A	CdSe/ZnS	Jasmonic Acid	Biotic Stress (Pests)
Dot B	InP/ZnS	Salicylic Acid	Biotic Stress (Pathogen)
Dot C	Graphene QD	β -caryophyllene	Abiotic Stress (Drought)
Dot D	Carbon Dot	Ethylene	General Stress / Ripening
Dot E	ZnSe QD	Dissolved Zn^{2+} ions	Soil Nutrient (Micronutrient)
Dot F	CdTe/Cds	Abscisic Acid (ABA)	Abiotic Stress (Drought)
Dot G	Si QD	Methyl Salicylate	Systemic Acquired Resistance
Dot H	AgInS ₂ /ZnS	Dissolved Fe^{3+} ions	Soil Nutrient (Micronutrient)

By creating a theoretical array with dozens of such dots, it is possible to simulate the generation of a unique "spectral fingerprint" for complex mixtures of VOCs and soil ions, as shown in Figure 3. This unique signature provides the rich data needed for the advanced classification models.

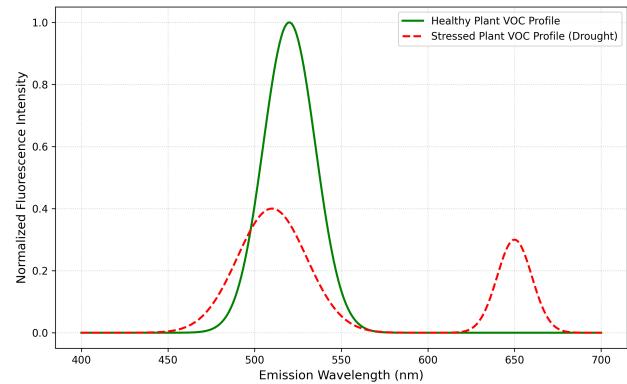


Figure 3: Simulated spectral fingerprint from the QDSS array model. The distinct spectral shift for stressed plants (dashed red line) versus healthy plants (solid green line) highlights the model's ability to generate unique, high-dimensional data for stress classification.

IV. METHODOLOGY

A. DATA ACQUISITION AND CALIBRATION

The first step in the methodology is processing the raw data from the QDSS model into a clean, normalized format. Algorithm 1 outlines this process. To handle anomalous readings from sensor noise, outliers are first identified and replaced using a standard Z-score test. To account for slow sensor drift over time, a dynamic baseline is maintained using an Exponentially Weighted Moving Average (EWMA). The formula for the EWMA, which uses the outlier-filtered sensor data ($S_{valid,t}$), is shown in Eq. 2. This approach ensures that the baseline is not corrupted by transient spikes. The final steps involve calculating the fluorescence quenching and normalizing the vector.

$$B_t = \alpha S_{valid,t} + (1 - \alpha) B_{t-1} \quad (2)$$

B. FEDERATED STRESS-PHENOTYPING MODEL (FSPM)

The core of the system's predictive capability lies in the FSPM, which is trained using FL to protect farmer privacy, as detailed in Algorithm 2. The model is implemented as a feed-forward neural network with two hidden layers, utilizing Batch Normalization and Dropout for regularization. A composite loss function is proposed (Eq. 3), which dynamically balances the objectives of yield prediction and stress classification using a heuristic weighting term, λ_t . During local training on each client device, model weights (W) are iteratively adjusted to reduce error using the Adam optimization algorithm. For this study, the weighting term λ_t was held constant at 0.5 to give equal importance to both the yield prediction and stress classification tasks.

$$\mathcal{L}_{total} = \lambda_t \mathcal{L}_{yield} + (1 - \lambda_t) \mathcal{L}_{stress} + \frac{\beta}{2M} \sum_{l=1}^L \|W_l\|_2^2 \quad (3)$$

Algorithm 1 Enhanced QDSS Data Calibration and Normalization

```

1: Input: Raw fluorescence vector  $S_{raw,t}$ , historical mean
    $\mu_{hist}$ , std dev  $\sigma_{hist}$ , EWMA baseline  $B_{t-1}$ , learning rate
    $\alpha$ 
2: Output: Normalized stress vector  $S_{norm}$ , updated base-
   line  $B_t$ 
3:  $S_{valid} \leftarrow S_{raw,t}$ 
4: for each reading  $s_i \in S_{raw,t}$  do
5:    $Z_i \leftarrow (s_i - \mu_{hist,i}) / \sigma_{hist,i}$ 
6:   if  $|Z_i| > 3$  then
7:      $S_{valid}[i] \leftarrow B_{t-1}[i]$ 
8:    $B_t \leftarrow \alpha S_{valid} + (1 - \alpha) B_{t-1}$ 
9:    $S_{quenched} \leftarrow B_t - S_{valid}$ 
10:  for  $i = 1$  to  $\text{length}(S_{valid})$  do
11:    if  $S_{quenched}[i] < 0$  then
12:       $S_{quenched}[i] \leftarrow 0$ 
13:    $S_{norm} \leftarrow \frac{S_{quenched}}{\|S_{quenched}\|_2}$ 
14:   Update  $\mu_{hist}, \sigma_{hist}$  with new valid readings
15: return  $S_{norm}, B_t$ 

```

Algorithm 2 Federated Learning for Yield Prediction (Ex-
panded)

```

1: Server Executes:
2: Initialize global model weights  $W_G^{(0)}$ 
3: for each round  $t = 0, 1, \dots, T - 1$  do
4:    $m \leftarrow \max(C \cdot K, 1)$ 
5:    $S_t \leftarrow (\text{random subset of } m \text{ clients})$ 
6:   for each client  $k \in S_t$  in parallel do
7:      $W_k^{(t+1)} \leftarrow \text{ClientUpdate}(k, W_G^{(t)})$ 
8:    $W_G^{(t+1)} \leftarrow \sum_{k=1}^m \frac{n_k}{n} W_k^{(t+1)}$ 
9: return  $W_G^{(T)}$ 

10: procedure CLIENTUPDATE( $k, W_G$ )
11:    $W \leftarrow W_G$ 
12:    $\mathcal{B} \leftarrow (\text{split local data } D_k \text{ into batches})$ 
13:   for each local epoch  $i$  from 1 to  $E$  do
14:     for batch  $b \in \mathcal{B}$  do
15:        $W \leftarrow W - \eta \nabla \mathcal{L}_{local}(W; b)$ 
16:   return  $W$  to server

```

C. PROOF-OF-SUSTAINABLE-PRACTICE (PoSP) CONSENSUS

To cryptographically verify and reward sustainable actions, the PoSP consensus mechanism is proposed. The process, detailed in Algorithm 3, begins when a farmer submits a transaction \mathcal{T} , which contains the farmer's signed node ID, N_f . If the claim is rejected, or if a validator is found to be consistently voting against the honest majority, a portion of their staked assets can be "slashed" (i.e., forfeited), creating a strong disincentive for malicious behavior. The validation function, $\mathcal{F}_{validate}$ (Eq. 5), checks for correlation between on-chain and oracle data. The trust score of each validator is

dynamically adjusted after each vote based on their accuracy (Eq. 4), ensuring that honest validators gain influence over time.

$$T_{v,t+1} = (1 - \gamma)T_{v,t} + \gamma A_v \quad (4)$$

$$\mathcal{F}_{validate} = (\delta(\mathcal{D}_{chain}, \mathcal{D}_{oracle}) < \epsilon) \wedge (\text{VerifySignature}(N_f)) \quad (5)$$

The distance metric δ is a crucial component whose implementation depends on the data types being compared, such as Intersection over Union (IoU) for imagery or Mean Squared Error (MSE) for time-series data. Consensus is achieved by a trust-weighted majority, making the system more robust.

Algorithm 3 dPoS-SS Consensus Mechanism with Trust-
Weighted Voting

```

1: Input: Transaction  $\mathcal{T}$  (contains farmer node ID  $N_f$ )
2: Propose block  $\mathcal{B}$  containing  $\mathcal{T}$ 
3: Select random validator set  $\mathcal{V}$ 
4:  $T_{total} \leftarrow \sum_{v \in \mathcal{V}} T_v$ 
5:  $T_{approved} \leftarrow 0$ 
6: for each validator  $v \in \mathcal{V}$  in parallel do
7:   Fetch  $\mathcal{D}_{chain}$  and  $\mathcal{D}_{oracle}$ 
8:    $vote \leftarrow \mathcal{F}_{validate}(\mathcal{T}, \mathcal{D}_{chain}, \mathcal{D}_{oracle})$ 
9:   if  $vote == 1$  then
10:     $T_{approved} \leftarrow T_{approved} + T_v$ 
11: if  $T_{approved} > \frac{2}{3} T_{total}$  then
12:   Append  $\mathcal{B}$  to blockchain; Mint  $R_{token}$  to  $N_f$ 
13: else
14:   Reject  $\mathcal{B}$  and potentially enact slashing penalties.

```

D. ADDITIONAL PROPOSED ALGORITHMS FOR DECISION SUPPORT

To illustrate the practical, downstream applications that the core predictive model is designed to enable, the following decision-support algorithms are also proposed. The QAL framework proposes several other algorithms to form a complete decision support system. While their detailed performance evaluation is outside the scope of this paper, their design is outlined here to illustrate the framework's full potential.

a: Anomaly Detection with Temporal Persistence

Algorithm 4 serves as an intelligent early warning mechanism. It uses the standard Mahalanobis distance to generate an anomaly score. To prevent false positives, the algorithm incorporates temporal persistence, triggering an alert only if the score remains high for a predefined confirmation window.

b: Multi-Variate Dynamic Nutrient Recommendation

To account for complex nutrient interactions, Algorithm 5 proposes a multi-variate optimization model. It aims to compute a cost-effective fertilizer blend that addresses the primary deficiency while preventing secondary imbalances, also integrating weather data to prevent runoff.

Algorithm 4 Anomaly Detection with Temporal Persistence

1: **Input:** Normalized stress vector S_{norm} , historical mean μ_{hist} , covariance matrix Σ_{hist} , dynamic threshold $Threshold_t$
 2: **Persistent State:** Anomaly counter $C_{anomaly}$, Confirmation window $W_{confirm}$
 3: **Output:** Contextual alert (Severity, Duration)
 4: Calculate Mahalanobis distance to quantify deviation from the healthy baseline:

$$\delta \leftarrow \sqrt{(S_{norm} - \mu_{hist})^T \Sigma_{hist}^{-1} (S_{norm} - \mu_{hist})} \quad (6)$$

5: **if** $\delta > Threshold_t$ **then**
 6: $C_{anomaly} \leftarrow C_{anomaly} + 1$
 7: **else**
 8: $C_{anomaly} \leftarrow 0$
 9: **if** $C_{anomaly} \geq W_{confirm}$ **then**
 10: Severity \leftarrow CalculateSeverity(δ , $Threshold_t$)
 11: Trigger alert: "Persistent stress event detected. Severity: **Severity**, Duration: $C_{anomaly}$ hours."
 12: **return** δ , Severity

Algorithm 5 Multi-Variate Dynamic Nutrient Recommendation

1: **Input:** Stress class prediction P_{stress} , soil data vector S_{soil} , Crop Growth Stage G_t , Weather Forecast $F_{weather}$
 2: **Constants:** Optimal nutrient vector $N_{optimal}$, Nutrient Interaction Matrix M_{int} , Fertilizer Cost vector C_{fert} , Fertilizer Nutrient Matrix M_{fert}
 3: **Output:** Optimized fertilizer application plan R_{plan}
 4: $N_{current} \leftarrow \text{decode}(S_{soil})$
 5: $N_{deficit} \leftarrow N_{optimal} - N_{current}$
 6: **if** $\max(N_{deficit}) \leq 0$ **then**
 7: **return** "No nutrient application needed."
 8: **if** $F_{weather}.\text{rainProbability} > 0.8$ **then**
 9: **return** "Recommendation delayed due to impending heavy rain."
 10: Let \mathbf{x} be the vector of fertilizer amounts to apply.
 11: **Minimize:** $C_{fert}^T \mathbf{x}$
 12: **Subject to:**
 1) $N_{current} + M_{fert} \mathbf{x} \geq N_{optimal}$
 2) $(M_{int}(N_{current} + M_{fert} \mathbf{x}))_i \leq \text{toxicity_threshold}_i$
 13: $R_{plan} \leftarrow \text{SolveOptimization}(\mathbf{x})$
 14: **return** R_{plan}

c: Adaptive Anomaly Thresholding

Algorithm 6 proposes a dynamic threshold that adapts to crop growth stage and recent environmental context. By using a rolling statistical window and a volatility index, it aims to maximize detection accuracy.

Algorithm 6 Adaptive Anomaly Thresholding

1: **Input:** Recent history of anomaly scores \mathcal{A}_{window} , crop growth stage G_t , recent environmental data E_{window}
 2: **Output:** Highly contextual, adjusted threshold $Threshold_t$
 3: $\mu_{A,window} \leftarrow \text{mean}(\mathcal{A}_{window})$
 4: $\sigma_{A,window} \leftarrow \text{std}(\mathcal{A}_{window})$
 5: $Threshold_{base} \leftarrow \mu_{A,window} + 3\sigma_{A,window}$
 6: **if** $G_t == \text{"Flowering"}$ or $G_t == \text{"Fruiting"}$ **then**
 7: $\text{Factor}_{stage} \leftarrow 0.85$
 8: **else**
 9: $\text{Factor}_{stage} \leftarrow 1.0$
 10: $\text{VolatilityIndex} \leftarrow \text{std}(E_{window}.\text{temp}) + \text{std}(E_{window}.\text{humidity})$
 11: $\text{Factor}_{volatility} \leftarrow 1 - (0.2 \times \tanh(\text{VolatilityIndex}))$
 12: The final threshold is adjusted by both biological and environmental factors:

$$\text{Threshold}_t \leftarrow \text{Threshold}_{base} \times \text{Factor}_{stage} \times \text{Factor}_{volatility} \quad (7)$$

 13: **return** Threshold_t

V. EXPERIMENTAL SETUP

To evaluate the performance of the QAL framework, a comprehensive simulation was designed and executed. The environment was built in Python 3.9, utilizing PyTorch for the federated neural network models, Scikit-learn for baseline models and data processing, and SimPy for discrete-event simulation of the blockchain network. The evaluation hinges on a synthetic dataset generated to model complex agricultural scenarios under controlled conditions.

A. IOT NETWORK AND SENSOR DATA SIMULATION

To provide realistic inputs for the framework's algorithms, the underlying IoT network and conventional sensor data streams were modeled.

- **Conventional Sensor Data:** The simulation generated time-series data for traditional sensors (temperature, humidity, and soil moisture). To provide richer temporal context for the model, this conventional data stream was expanded to a total of 12 features. This included sinusoidal functions to reflect diurnal cycles (sine/cosine transformations of the hour) and time-lagged values (e.g., from 6 and 12 hours prior) of the primary environmental variables.

- **Network Characteristics:** The communication network between sensor nodes and the edge gateway was also modeled. Based on typical Wi-Fi-based agricultural deployments, the simulation assumed an average data

packet success rate of 98.7% and a network latency modeled by a log-normal distribution with a mean of 52ms.

B. SYNTHETIC DATASET GENERATION FOR QAL

A synthetic dataset representing approximately **20,000 data points** per run was generated to provide a robust basis for model training and validation. The protocol was designed to create a challenging and realistic simulation environment far exceeding simple Gaussian noise.

- **Stress Event Injection:** Biotic (Pest) and abiotic (Drought, Nutrient Deficiency) stress events were introduced following a Bernoulli distribution, with event durations modeled to create a gradual, non-linear onset and recovery profile.
- **Complex VOC Signal Modeling:** To simulate a realistic and challenging scenario, the final signal $S_i(t)$ for each of the 96 QDSS sensor bands was modeled as a superposition of multiple components as described by Eq. 8. This included a common "general stress" signal, $\mathcal{G}(c, t)$, for all non-healthy states, and weaker, class-specific signals, $\mathcal{V}_i(c, t)$, with overlapping spectral features to differentiate between stressors. Signal strength was modeled as a non-linear function of stress duration to reflect biological reality. Crucially, a highly complex and structured noise component, $\mathcal{N}_{complex}(\sigma, t)$, was introduced by layering multiple effects: baseline noise, signal drift, periodic interference, random high-magnitude spikes, and sensor saturation effects simulated via a hyperbolic tangent function. This design forces the model to learn subtle, complex, and co-occurring patterns rather than simple, idealized features.

$$S_i(t) = \tanh(\mathcal{B}_i + \mathcal{G}(c, d_t) + \mathcal{V}_i(c, d_t) + \mathcal{N}_{complex}(\sigma, t)) \quad (8)$$

where \mathcal{B}_i is the baseline signal, c is the stress class, and d_t is the duration of the stress event.

- **Ground-Truth Labels:** Final crop yield (in tons/ha) was calculated using a baseline potential, which was then penalized based on the type and duration of simulated stress events, with added Gaussian noise to simulate natural variability in crop response. It is acknowledged that the current synthetic dataset focuses on physicochemical VOC responses. It does not yet capture complex biological feedback loops, such as evolutionary pest resistance or crop-specific growth stages. Future iterations of the simulator will incorporate these biological dynamics to further validate the model's generalization capabilities.

C. FEDERATED DATA PARTITIONING AND FEATURE ENGINEERING

To simulate a real-world federated scenario with distinct farm datasets, the generated data was partitioned among **20 simulated clients**. The overall dataset was first split into an 80% training set and a 20% hold-out test set using a

standard random split. The 80% training portion was then evenly distributed among the clients.

A crucial step of the methodology was client-side feature engineering. Before local training, each client autonomously processed its time-series data to generate an exhaustive set of temporal features. This process expanded the feature space significantly by transforming the raw sensor readings into a rich representation of trends and volatility. The final input vector size of 108 features is derived from the 96 discrete spectral bands of the QDSS array combined with the 12 engineered features from conventional sensors. The specific features generated are detailed in Table 6. All feature scaling (StandardScaler) was performed based on parameters derived only from the training data partitions to prevent data leakage.

Table 6: Details of Client-Side Temporal Feature Engineering.

Feature Type	Parameters	Applied To
Rolling Mean	Time Windows: 6, 12, 24h	Base Sensor Channels
Rolling Std Dev	Time Windows: 6, 12, 24h	Base Sensor Channels

D. MODEL ARCHITECTURE AND FEDERATED TRAINING PROTOCOL

The FSPM utilized an advanced feed-forward neural network designed for robustness, implemented in PyTorch. The architecture included an input layer sized to the expanded feature set, followed by two hidden layers augmented with **Batch Normalization** and **Dropout** to prevent overfitting and improve generalization. The final model had two output heads: one for yield prediction (regression) and one for multi-class stress classification (softmax).

To ensure optimal model configuration, automated hyperparameter tuning was employed within the federated context using the Optuna framework. This process involved systematically searching for the best-performing combination of learning rates, dropout rates, and layer sizes by running multiple, short-duration federated simulations. The hyperparameter set that yielded the best performance on a validation set was selected for the final, full-duration training runs. The search space for this optimization is defined in Table 7.

Table 7: Hyperparameter Search Space for FSPM Tuning.

Hyperparameter	Distribution	Range
Learning Rate	Log-Uniform	$[1 \times 10^{-4}, 1 \times 10^{-2}]$
Dropout Rate	Uniform	$[0.2, 0.5]$
Layer 1 Neurons	Integer	$[64, 256]$
Layer 2 Neurons	Integer	$[32, 128]$

The main simulation consisted of **10 independent runs** to ensure statistical validity of the results. In each run, the best-found hyperparameters were used to train a global model over **40 communication rounds**. In each round, a subset of clients performed local training for 5 epochs using the Adam

optimizer before their model updates were securely aggregated using the Federated Averaging (FedAvg) algorithm. All reported results are the mean and standard deviation across these 10 runs.

E. COMPREHENSIVE ABLATION STUDY

To rigorously evaluate the contributions of both the advanced data source (QDSS) and the neural network architecture (FSPM), a comprehensive ablation study was conducted. This study was organized into two groups, with three distinct models evaluated in each:

- 1) **Group 1 (Conventional Data Only):** This group establishes performance baselines using only the 12-dimensional conventional environmental sensor data. The following models were trained on this limited dataset:
 - Random Forest (RF)
 - Support Vector Machine (SVM)
 - FSPM (Neural Network)
- 2) **Group 2 (Full Data):** This group evaluates the same three models using the complete, high-dimensional dataset combining both conventional and QDSS sensor streams. This setup is designed to isolate the performance gains attributable to the richer data.

This two-group approach allows for a robust comparison, highlighting not only the value of the QDSS data but also the effectiveness of different machine learning architectures in interpreting that data. The RF and SVM models, while implemented with standard parameters for this study, were included as they represent strong and commonly used benchmarks for tabular data analysis.

F. BLOCKCHAIN NETWORK SIMULATION

The dPoS-SS consensus mechanism was simulated with a network of **100 validator nodes** and benchmarked against a PoW network model.

- **Network Parameters:** The simulation assumed a block time of 15 seconds, an average transaction size of 250 bytes, and network latency modeled using a log-normal distribution.
- **PoW Benchmark Details:** The PoW benchmark was modeled on the parameters of the pre-merge Ethereum network (Ethash). Its energy consumption was calculated based on the model provided by De Vries [18], which correlates energy use with network hash rate and revenue.

G. NOISE SENSITIVITY ANALYSIS

To evaluate the robustness of the trained FSPM against real-world sensor imperfections, a noise sensitivity analysis was conducted. This analysis began by establishing a baseline performance on the clean test set (noise $\sigma = 0.0$). Subsequently, the entire multi-run simulation was repeated at several noise levels, where Gaussian noise with a systematically increasing standard deviation (σ) was added to the

raw sensor data streams before the feature engineering step. The final aggregated model performance (Mean Accuracy and Mean RMSE) was recorded at each noise level to assess the degradation threshold.

VI. RESULTS AND DISCUSSION

The performance of the QAL framework was evaluated through the comprehensive simulation detailed in Section V. The analysis focused on the predictive accuracy and robustness of the FSPM, the efficiency of the blockchain mechanism, and the potential agricultural impact. All key performance metrics are reported as mean \pm standard deviation over 10 independent simulation runs to ensure statistical robustness.

Following the federated hyperparameter tuning process, a set of optimal parameters was identified for the FSPM. A representative set of the best-performing hyperparameters is presented in Table 8. These values were used to configure the FSPM for the final training and evaluation.

Table 8: Representative Optimal Hyperparameters for the FSPM.

Hyperparameter	Optimal Value
Learning Rate	3.7×10^{-4}
Dropout Rate	0.35
Layer 1 Neurons	196
Layer 2 Neurons	115

A. FSPM PERFORMANCE ANALYSIS

1) Yield Prediction and Stress Classification

The FSPM, leveraging the combined conventional and QDSS data streams, demonstrated strong predictive performance on the highly complex synthetic dataset. Figure 4 visualizes the yield prediction performance from a representative simulation run. It contrasts the baseline SVM model trained only on conventional data—which shows a wide prediction variance (RMSE = 1.45 tons/ha for this specific run)—with the full QAL model. The QAL model’s predictions cluster much more tightly around the line of perfect prediction, achieving a mean RMSE of **1.31 \pm 0.20 tons/ha** across all runs.

In the domain of stress classification, the model showed excellent diagnostic precision. The mean confusion matrix, aggregated over all 10 runs, is presented in Figure 5. The high values along the diagonal—for instance, ‘1203.6 \pm 106.6’ correct classifications for the ‘Healthy’ state—confirm the model’s ability to consistently identify the correct condition. The overall mean accuracy was a notable **96.71% \pm 0.47%**. The matrix shows that the most common errors are confusions between the different stress types (e.g., Drought being misclassified as Healthy ‘30.9 \pm 6.4’ times), an expected outcome given the intentionally overlapping signal features in the data model. To further validate this, the mean Receiver Operating Characteristic (ROC) curves are shown in Figure 6. The high Area Under the Curve (AUC) for all classes

(mean AUC > 0.996) demonstrates the model's excellent and stable ability to distinguish between all states.

2) Comprehensive Ablation Study and Feature Importance

To rigorously evaluate the contributions of both the advanced data source and the neural network architecture, a comprehensive two-group ablation study was conducted. The full results are presented in Table 9 and visualized in Figure 7.

The study evaluated three different model architectures (RF, SVM, and FSPM) on both the limited conventional dataset and the full dataset. As seen in Table 9, when trained only on conventional data, the traditional machine learning models (RF and SVM) outperform the FSPM neural network, achieving classification accuracies of 32.5% and 33.7%, respectively, compared to the FSPM's 27.3%. This suggests that for lower-dimensional data, traditional models are highly effective.

However, when the models are given access to the full, high-dimensional dataset including QDSS streams, the FSPM demonstrates its superior capability. Its accuracy dramatically increases to 96.7%, surpassing both the RF (66.6%) and SVM (95.6%) models. This two-group analysis provides a crucial insight: the high performance of the QAL system is a result of the synergy between the rich, molecular-level data from the QDSS and an advanced neural network architecture adept at interpreting its complexity.

The feature importance analysis, shown in Figure 8, confirms the dominant contribution of the novel data source. The top 20 most predictive features are entirely derived from the QDSS vector. This result provides strong evidence that the molecular-level data captured by the QDSS model is the primary driver of the system's high diagnostic accuracy.

B. MODEL ROBUSTNESS AND SENSITIVITY ANALYSIS

1) Federated Learning Convergence

The convergence of the FL model, shown in Figure 9, demonstrates stable and efficient learning. The plot displays the mean validation accuracy over 10 runs, which shows a rapid increase during the initial 25-30 communication rounds before reaching a stable plateau around 96.7%. The narrow shaded region, representing ± 1 standard deviation, indicates low variance between runs and suggests consistent convergence under the simulated ideal network conditions.

2) Noise Sensitivity Analysis

To evaluate the FSPM's resilience to real-world sensor imperfections, a noise sensitivity analysis was conducted. As illustrated in Figure 10, the model's performance shows a clear trend of graceful degradation as sensor noise increases. Starting from its baseline on clean data ($\sigma = 0.0$), the stress classification accuracy (orange dashed line) remains high, at approximately 96%, up to a noise level of $\sigma = 0.10$. It then declines more steeply to approximately 62% at the maximum tested noise level of $\sigma = 0.30$. The Yield Prediction RMSE (blue solid line) shows a corresponding increase in error,

rising from its baseline of 1.31 tons/ha to a peak of approximately 1.95 tons/ha around a noise level of $\sigma = 0.20$. The analysis indicates a notable tolerance for sensor imprecision, although it underscores the importance of high-quality sensor fabrication for optimal real-world performance.

C. BLOCKCHAIN PERFORMANCE AND AGRICULTURAL IMPACT

While the anomaly detection and adaptive thresholding algorithms (Algorithms 4 and 6) performed their auxiliary roles effectively in preliminary tests, the core viability of the verification layer depends on the efficiency of its consensus mechanism. Figure 11 shows the mean performance with error bars representing variance across 10 runs. The dPoS-SS mechanism achieves significantly higher transaction throughput, scaling to 187 ± 11 TPS with 100 validators, compared to the flat performance of PoW. More critically, its energy consumption is orders of magnitude lower, representing a 98.8% reduction per transaction.

These technical improvements are designed to translate into tangible agricultural benefits. The illustrative simulation of nutrient management in Figure 12, now driven by the aggregated model outputs, suggests that the dynamic recommendations can maintain soil nitrogen in the optimal zone. While the model without QAL recommendations sees a steady decline into a stress state, the QAL-driven approach successfully prevents deficiency and reduces potential fertilizer waste over the simulated 60-day period.

D. SECURITY AND ATTACK VECTOR ANALYSIS

Given the decentralized nature of the framework, we analyze its resilience against key attack vectors:

1) Validator Collusion

The dPoS-SS mechanism increases the economic cost of collusion. By implementing a configurable “Slashing” fraction, the protocol ensures that validators voting against the oracle-verified consensus risk forfeiting a portion of their stake.

2) Adversarial Economic Incentives

To counter adversarial economic incentives, the staking and slashing rules are designed to ensure that the cost of short-term malicious strategies (e.g., profitable misreporting followed by exit) exceeds the potential gains, making long-term honest participation the rational equilibrium, although a full game-theoretic analysis is left to future work.

3) Edge Node Forgery

Architecturally, the framework proposes hardware-level signing (e.g., ECDSA) at the Perception Layer to ensure data non-repudiation, though this cryptographic primitive is not explicitly modeled in the current throughput simulation.

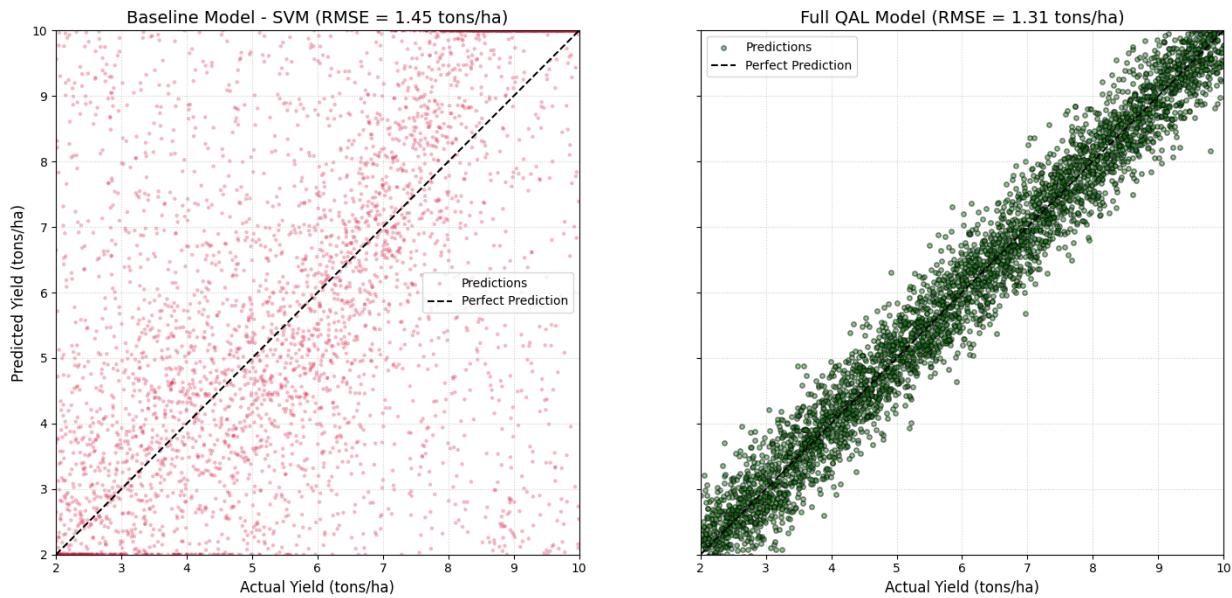


Figure 4: Comparison of simulated yield prediction accuracy from a representative run, plotting predicted vs. actual yield. The dashed line represents a perfect prediction. (Left) The baseline SVM model trained on conventional data shows significant prediction error (RMSE = 1.45 tons/ha). (Right) The Full QAL model's predictions cluster tightly around the line of perfect prediction, demonstrating a much stronger correlation with the actual yield (RMSE = 1.31 tons/ha). The mean RMSE for the QAL model across 10 runs was 1.31 ± 0.20 tons/ha.

Table 9: Comprehensive Ablation Study Results (Mean \pm Std Dev over 10 runs).

Model Configuration	Yield Prediction RMSE (tons/ha)	Stress Classification Accuracy (%)
<i>Group 1: Trained on Conventional Data Only</i>		
Random Forest (RF)	1.35 ± 0.22	32.5 ± 1.6
Support Vector Machine (SVM)	1.45 ± 0.22	33.7 ± 2.5
FSPM (Neural Network)	2.44 ± 0.24	27.3 ± 1.6
<i>Group 2: Trained on Full Data (Conventional + QDSS)</i>		
Random Forest (RF)	1.34 ± 0.22	66.6 ± 11.2
Support Vector Machine (SVM)	1.55 ± 0.23	95.6 ± 0.5
FSPM (QAL Model)	1.31 ± 0.21	96.7 ± 0.5

4) Oracle Manipulation

Reliance on a decentralized oracle network is proposed to mitigate single-point-of-failure risks regarding weather and satellite verification data.

E. DISCUSSION OF LIMITATIONS

While the simulation results are promising, a critical discussion of the framework's limitations is essential to contextualize the findings and guide future work.

1) Technical and Methodological Limitations

The primary limitation is the study's reliance on a synthetic dataset generated from simplified biological and physical models. Real-world agricultural environments exhibit complexities and covariances not fully captured in the simulation. The simulation assumes Independent and Identically Distributed (IID) data partitions. Real-world agricultural data is highly non-IID. Consequently, the reported accuracy represents an upper bound under ideal federated conditions.

Future work will integrate clustered Federated Learning and meta-learning-based personalization to mitigate performance degradation across heterogeneous farm environments. Furthermore, the proposed dPoS-SS mechanism, while conceptually sound, requires a more rigorous security analysis against sophisticated economic and collusion attacks before deployment. Finally, this work does not include a direct performance comparison against existing commercial agricultural monitoring systems, which often operate as black boxes, making such a comparison difficult within a simulation context.

2) Potential Economic Viability

A theoretical cost-benefit analysis frames the potential for real-world adoption.

- **Estimated Costs:** The primary cost drivers would be the fabrication and deployment of the QDSS arrays and the maintenance of the on-farm edge devices and network

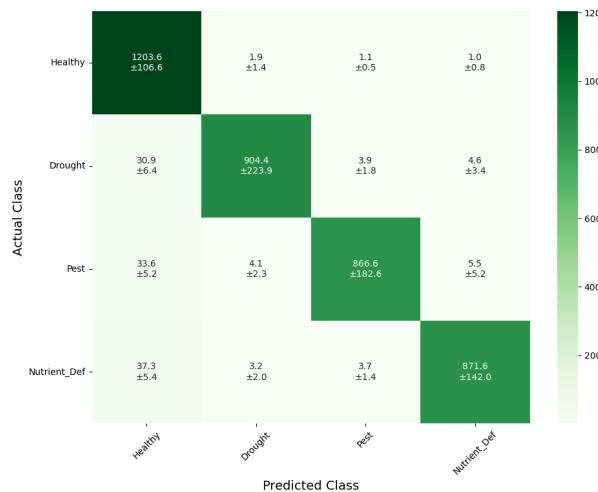


Figure 5: Mean confusion matrix for FSPM stress classification over 10 simulation runs. Each cell displays the mean prediction count \pm one standard deviation. The high values along the diagonal confirm the model's diagnostic precision, achieving a mean accuracy of $96.71\% \pm 0.47\%$.

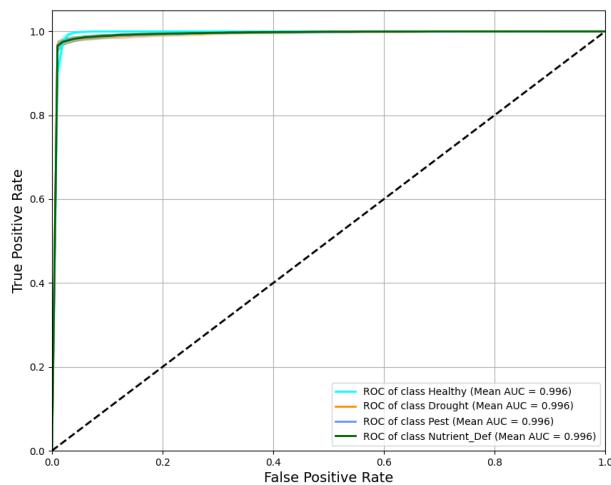


Figure 6: Mean ROC curves over 10 simulation runs. The shaded regions represent ± 1 standard deviation. The high Area Under the Curve (AUC) for all classes demonstrates the model's excellent and stable ability to distinguish between healthy and various stressed states.

infrastructure. Initial costs could be substantial, representing a significant barrier for smallholder farmers.

- **Potential Gains:** The economic benefits could be significant. The improved yield prediction error could help optimize harvesting logistics and forward sales. More importantly, the ability to maintain optimal nutrient levels and preemptively identify stress, as illustrated in the simulations, could prevent catastrophic crop losses and reduce input waste, offering a high potential return on investment.

This preliminary analysis suggests that while the initial capital expenditure is high, the potential for significant operational savings and risk reduction could make the QAL framework economically viable, particularly for high-value crops.

3) Ethical Considerations and Governance

The deployment of such a technology necessitates careful consideration of its ethical implications.

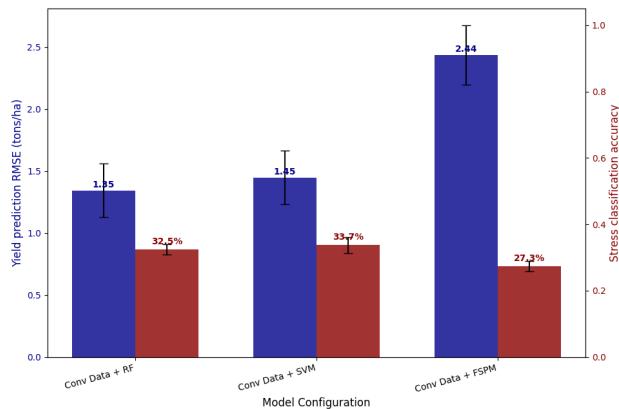
- **Data Sovereignty and Privacy:** The use of FL is a direct response to the critical need for farmer data sovereignty. It ensures that raw, operational farm data remains on-premises, preventing exploitation by large corporations.
- **Equitable Access and the Digital Divide:** The high potential cost of the system risks creating a digital divide, where only large, well-capitalized farms can benefit. To address capital barriers for smallholder farmers, one potential deployment model is “Farming-as-a-Service” (FaaS). In this hypothetical scenario, cooperative-owned QDSS nodes could provide subscription-based insights, removing the burden of individual asset ownership.
- **Decentralized Governance:** The QAL network, with its tokenized rewards, forms a micro-economy. It is crucial to establish a fair, transparent, and decentralized governance model to manage network upgrades, resolve disputes, and ensure that the validator set does not become overly centralized, thereby defeating the purpose of a decentralized system.

VII. CONCLUSION AND FUTURE WORK

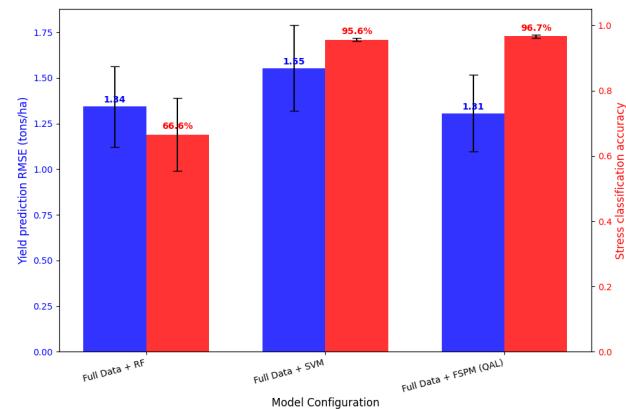
A. CONCLUSION

This paper introduced the Quantum-Enhanced Agri-Ledger (QAL), a framework designed to shift agriculture toward a preemptive management paradigm. While the findings are based on a theoretical framework and a simulated environment with a highly complex dataset, they present a strong blueprint for this paradigm shift. By integrating a novel QDSS sensor model, a privacy-preserving FL architecture, and an energy-efficient decentralized ledger, the QAL framework establishes a blueprint for a system that can enhance productivity, preserve farmer data sovereignty, and directly incentivize sustainability.

The comprehensive simulations yielded promising results. The findings suggest that leveraging molecular-level data from the QDSS model leads to significant improvements in diagnostic accuracy for crop stress, with the FSPM achieving a mean accuracy of 96.71%. Furthermore, the dPoS-SS mechanism was shown to be a viable, energy-efficient alternative to traditional consensus protocols, making a secure verification ledger more feasible for agricultural contexts. While this work is simulation-based, it presents a compelling case for this integrated approach, justifying the future research required to translate this blueprint into a field-ready technology.



(a) Performance on Conventional Data Only.



(b) Performance on Full (Conv + QDSS) Data.

Figure 7: Results of the comprehensive ablation study, visualizing mean performance metrics over 10 runs with error bars representing ± 1 standard deviation. The charts contrast Yield Prediction RMSE (lower is better) with Stress Classification Accuracy (higher is better) for three different models on two different datasets, illustrating the significant accuracy gain achieved by integrating the QDSS data.

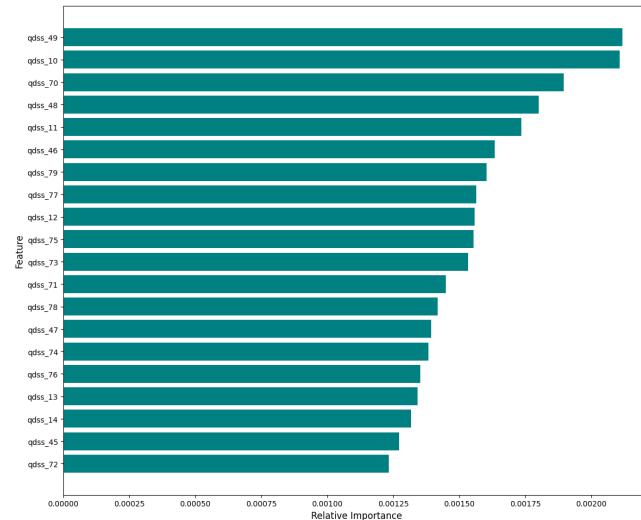


Figure 8: Top 20 relative feature importances from a representative FSPM training run. The analysis shows the overwhelming dominance of QDSS-derived features in the model's predictive power.

B. FUTURE WORK

Building on the simulation results of this work, future research will focus on transitioning the QAL framework from model to practice. The immediate next steps are organized into three key areas:

- 1) **Physical Prototype Development and Field Validation:** The primary goal is to fabricate and deploy physical QDSS prototypes. This involves significant materials science and engineering work to functionalize quantum dots and build a robust, low-power sensor package. Foundational work on a conventional IoT

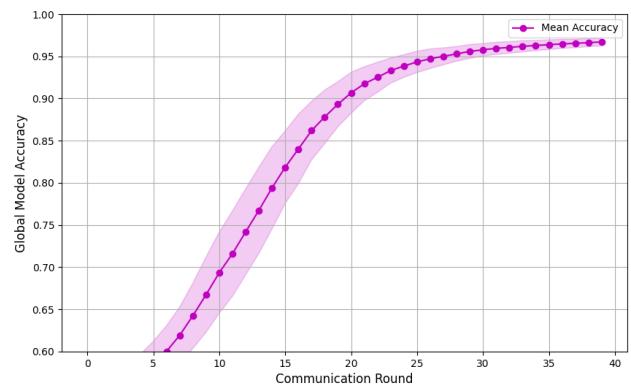


Figure 9: Mean Federated Learning model convergence over 10 simulation runs. The shaded region represents ± 1 standard deviation. The steady increase in validation accuracy demonstrates stable and efficient collaborative training.

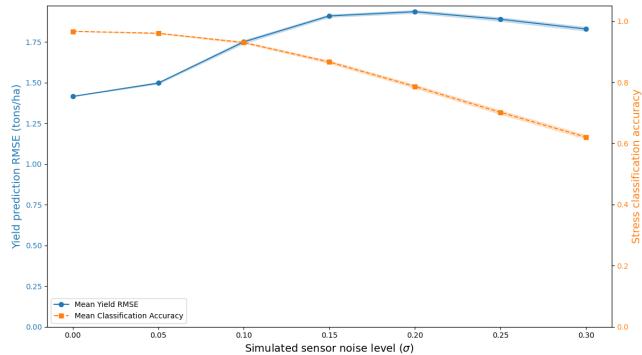


Figure 10: Results of the noise sensitivity analysis. The plot shows the mean degradation of Yield Prediction RMSE and Stress Classification Accuracy as sensor noise increases. Shaded regions represent ± 1 standard deviation over 10 runs.

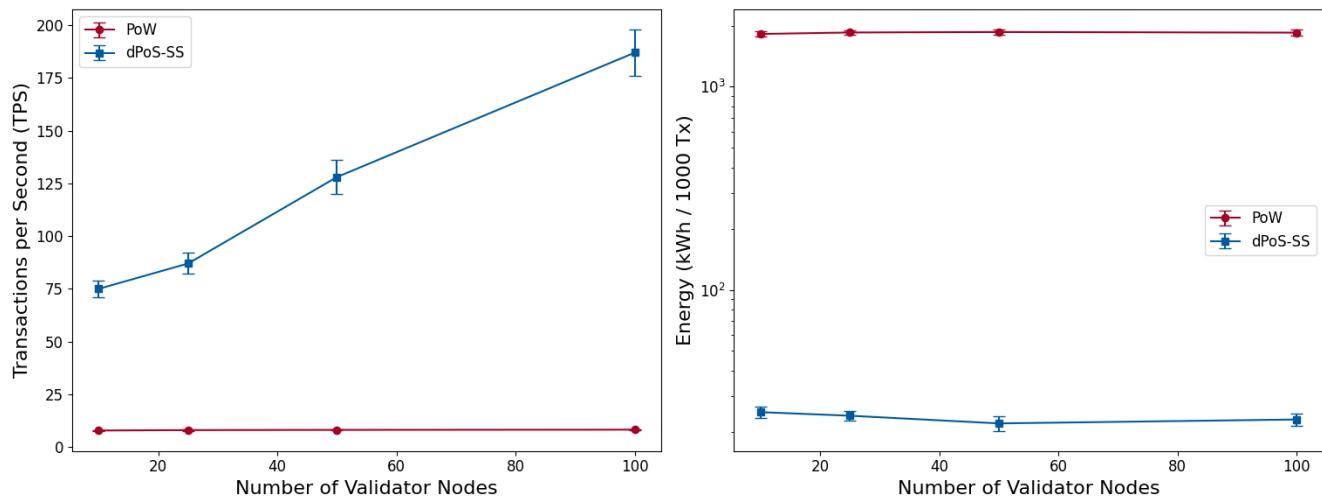


Figure 11: Performance comparison of dPoS-SS and PoW consensus mechanisms. Error bars denote ± 1 standard deviation over 10 runs. (Left) dPoS-SS demonstrates superior scalability, achieving over 180 TPS. (Right) On a logarithmic scale, dPoS-SS shows a >98% reduction in energy consumption per 1000 transactions compared to PoW, highlighting its sustainability.

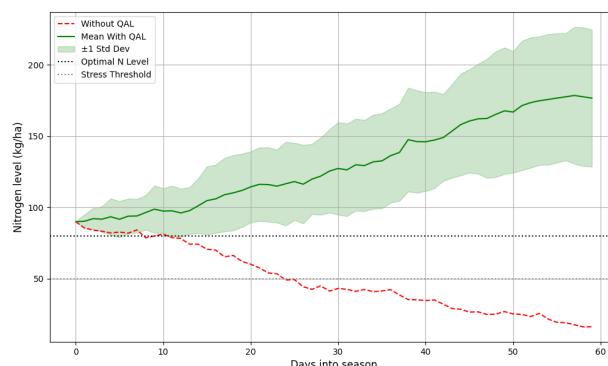


Figure 12: An illustrative simulation of soil nitrogen management over 10 runs. While the baseline scenario (red) shows depletion, the mean trend for the QAL-driven approach (green) maintains the nutrient level within the optimal zone. The shaded region represents ± 1 standard deviation.

node provides a viable platform for the future integration of this novel sensor. This phase will involve calibrating the sensor's sensitivity against gold-standard laboratory instruments (e.g., GC-MS) and assessing its long-term durability.

2) **Model Expansion and Robustness Testing with Real Data:** Once physical prototypes generate real data, the FSPM will be retrained and validated. Research will focus on training the model on a wider range of crops and diverse agro-climatic zones. A key challenge will be training the model to deconvolve complex, co-occurring stress signatures (e.g., simultaneous drought and nutrient deficiency).

3) **Blockchain Testnet Deployment and Economic Modeling:** A plan is in place to deploy the QAL blockchain on a permissioned testnet with pilot farms. This will allow for an evaluation of the economic viability and behavioral impact of the token-based incentive system, stress-testing of the dPoS-SS consensus mechanism, and optimizing the network for real-world latency and scalability.

DATA AVAILABILITY

This study utilizes a synthetic dataset generated through simulation. The source code used to generate the dataset, train the models, and reproduce the results and figures presented in this paper, together with a representative dataset from one of the simulation runs, is available at: <https://github.com/ahkharsha/quantum-agri-ledger>.

References

- [1] Food and Agriculture Organization of the United Nations, "How to feed the world in 2050," FAO, Tech. Rep., 2009, high-Level Expert Forum. [Online]. Available: https://www.fao.org/fileadmin/templates/wsfs/docs/expert_paper/How_to_Feed_the_World_in_2050.pdf
- [2] Ü. Niinemets, A. Kännaste, and L. Copolovici, "Plant volatile organic compound emissions in response to biotic and abiotic stresses: a review," *Trends in Plant Science*, vol. 23, no. 9, pp. 843–856, Sep. 2018. [Online]. Available: <https://doi.org/10.1016/j.tplants.2018.06.001>
- [3] A. Khanna and S. Kaur, "Evolution of Internet of Things (IoT) and its significant impact in the field of Precision Agriculture," *Computers and Electronics in Agriculture*, vol. 157, pp. 218–231, Feb. 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168169918316417>
- [4] M. S. Farooq, S. Riaz, A. Abid, T. Umer, and Y. B. Zikria, "Role of IoT Technology in Agriculture: A Systematic Literature Review," *Electronics*, vol. 9, no. 2, p. 319, Feb. 2020. [Online]. Available: <https://www.mdpi.com/2079-9292/9/2/319>
- [5] T. Adão, J. Hruška, L. Pádua, J. Bessa, E. Peres, R. Morais, and J. Sousa, "Hyperspectral Imaging: A Review on UAV-Based Sensors, Data Processing and Applications for Agriculture and Forestry," *Remote*

Sensing, vol. 9, no. 11, p. 1110, Oct. 2017. [Online]. Available: <https://www.mdpi.com/2072-4292/9/11/110>

[6] A.-K. Mahlein, "Plant Disease Detection by Imaging Sensors – Parallels and Specific Demands for Precision Agriculture and Plant Phenotyping," *Plant Disease*, vol. 100, no. 2, pp. 241–251, Feb. 2016. [Online]. Available: <https://apsjournals.apsnet.org/doi/10.1094/PDIS-03-15-0340-FE>

[7] A. Wilson, "Diverse Applications of Electronic-Nose Technologies in Agriculture and Forestry," *Sensors*, vol. 13, no. 2, pp. 2295–2348, Feb. 2013. [Online]. Available: <https://www.mdpi.com/1424-8220/13/2/2295>

[8] L. Wang, S. Gopalan, and R. Naidu, "Advancements in nanotechnological approaches to volatile organic compound detection and separation," *Current Opinion in Environmental Science & Health*, vol. 37, p. 100528, Feb. 2024. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S2468584423000880>

[9] M. Sliti, A. Elfikky, A. I. Boghdady, S. A. H. Mohsan, and S. Ayouni, "Integrating UAV-Mounted Diffraction Spectroscopy and FBG Sensor Networks for Real-Time Forest Health Monitoring," *IEEE Access*, vol. 13, pp. 196 195–196 205, 2025. [Online]. Available: <https://ieeexplore.ieee.org/document/11224884>

[10] Q. Li, M. El-Hajjar, Y. Sun, and L. Hanzo, "Performance Analysis of Reconfigurable Holographic Surfaces in the Near-Field Scenario of Cell-Free Networks Under Hardware Impairments," *IEEE Transactions on Wireless Communications*, vol. 23, no. 9, pp. 11 972–11 984, Sep. 2024. [Online]. Available: <https://ieeexplore.ieee.org/document/10502274>

[11] A. Kamlaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," *Computers and Electronics in Agriculture*, vol. 147, pp. 70–90, Apr. 2018. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0168169917308803>

[12] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," 2017, arXiv:1602.05629. [Online]. Available: <http://arxiv.org/abs/1602.05629>

[13] J. Wen, Z. Zhang, Y. Lan, Z. Cui, J. Cai, and W. Zhang, "A survey on federated learning: challenges and applications," *International Journal of Machine Learning and Cybernetics*, vol. 14, no. 2, pp. 513–535, Feb. 2023. [Online]. Available: <https://doi.org/10.1007/s13042-022-01647-y>

[14] G. Pramela, "Multi-Modal Deep Learning for Crop Yield Prediction Network: Static and Temporal Feature Space," *Journal of Information Systems Engineering and Management*, vol. 10, no. 35s, pp. 768–782, Apr. 2025. [Online]. Available: <https://jisem-journal.com/index.php/journal/article/view/6055>

[15] A. Kamlaris, A. Fonts, and F. X. Prenafeta-Boldú, "The rise of blockchain technology in agriculture and food supply chains," *Trends in Food Science & Technology*, vol. 91, pp. 640–652, Sep. 2019. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0924224418303686>

[16] Q. Lin, H. Wang, X. Pei, and J. Wang, "Food Safety Traceability System Based on Blockchain and EPCIS," *IEEE Access*, vol. 7, pp. 20 698–20 707, 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8640818/>

[17] Food and Agriculture Organization of the United Nations, "E-agriculture in action: Blockchain for agriculture," FAO, Tech. Rep., 2019. [Online]. Available: <https://www.fao.org/documents/card/en/c/ca5606en/>

[18] A. De Vries, "Bitcoin's Growing Energy Problem," *Joule*, vol. 2, no. 5, pp. 801–805, May 2018. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S2542435118301776>

[19] F. Saleh, "Blockchain without Waste: Proof-of-Stake," *The Review of Financial Studies*, vol. 34, no. 3, pp. 1156–1190, Feb. 2021. [Online]. Available: <https://academic.oup.com/rfs/article/34/3/1156/5868423>

[20] A. Pasdar, Y. C. Lee, and Z. Dong, "Connect API with Blockchain: A Survey on Blockchain Oracle Implementation," *ACM Computing Surveys*, vol. 55, no. 10, pp. 1–39, Oct. 2023. [Online]. Available: <https://dl.acm.org/doi/10.1145/3567582>

[21] X. Tang, X. Lan, L. Li, Y. Zhang, and Z. Han, "Incentivizing Proof-of-Stake Blockchain for Secured Data Collection in UAV-Assisted IoT: A Multi-Agent Reinforcement Learning Approach," *IEEE Journal on Selected Areas in Communications*, vol. 40, no. 12, pp. 3470–3484, Dec. 2022. [Online]. Available: <https://ieeexplore.ieee.org/document/9915381>

[22] Y. Harbi, K. Medani, C. Gherbi, O. Senouci, Z. Aliouat, and S. Harous, "A Systematic Literature Review of Blockchain Technology for Internet of Drones Security," *Arabian Journal for Science and Engineering*, vol. 48, no. 2, pp. 1053–1074, Feb. 2023. [Online]. Available: <https://link.springer.com/10.1007/s13369-022-07380-6>

[23] K. Andrews, L. B. Ngo, and M. Amiruzzaman, "A Detailed Comparative Analysis of Blockchain Consensus Mechanisms," 2025. [Online]. Available: <https://arxiv.org/abs/2511.15730>

[24] Ž. Grbović, B. Ivošević, M. Budjen, R. Waqar, N. Pajević, N. Ljubičić, V. Kandić, M. Pajić, and M. Panić, "Integrating UAV multispectral imaging and proximal sensing for high-precision cereal crop monitoring," *PLOS ONE*, vol. 20, no. 5, p. e0322712, May 2025. [Online]. Available: <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0322712>

[25] Monnit Corp., "New IoT Soil Moisture Sensor by Monnit Corp. for Smart Agriculture," Mar. 2021, accessed: Dec. 2025. [Online]. Available: <https://www.monnit.com/blog/2021/03/17/new-iot-soil-moisture-sensor-by-monnit-for-smart-agriculture/>

[26] M. Hasani, A. M. Coto García, J. M. Costa-Fernández, and A. Sanz-Medel, "Sol-gels doped with polymer-coated ZnS/CdSe quantum dots for the detection of organic vapors," *Sensors and Actuators B: Chemical*, vol. 144, no. 1, pp. 198–202, Jan. 2010. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0925400509008600>

[27] M. M. Rahman, M. R. Karim, M. M. Alam, M. B. Zaman, N. Alharthi, H. Alharbi, and A. M. Asiri, "Facile and efficient 3-chlorophenol sensor development based on photoluminescent core-shell CdSe/ZnS quantum dots," *Scientific Reports*, vol. 10, no. 1, p. 557, Jan. 2020. [Online]. Available: <https://www.nature.com/articles/s41598-019-57091-6>



A. HARSHA KUMAR is currently pursuing the Bachelor of Technology (B.Tech.) degree in Computer Science and Engineering from Vellore Institute of Technology, Chennai, India.

He has participated in research work at the Combat Vehicles Research and Development Establishment (CVRDE), DRDO, Chennai, contributing to research on distributed communication frameworks for in-vehicle networks using OpenDDS. He has also collaborated with Samsung R&D, contributing to the development of hybrid CNN-ViT watermarking systems for secure image authentication. His research contributions include publications on decentralized digital health systems and IoT-enabled maternal health monitoring. His areas of interest include blockchain applications, decentralized systems, machine learning, AI-enhanced healthcare, and the Internet of Things (IoT).

Mr. Kumar was a Grand Finalist at the IEEE YESIST12 2024 for his work on prenatal healthcare technologies.



N. SARAN is currently pursuing the Bachelor of Technology (B.Tech.) degree in Computer Science and Engineering from Vellore Institute of Technology, Chennai, India.

In collaboration with Samsung R&D Institute India, he contributed to research on post-quantum cryptography and secure communication frameworks. He has co-authored publications in the domains of decentralized digital health and IoT-based maternal health monitoring. His research interests include blockchain technology, artificial intelligence, post-quantum cryptography, and secure health informatics.

Mr. Saran was recognized as a Grand Finalist at the IEEE YESIST12 2024 for his work on digital healthcare innovation.



K. KUMARAN received the Ph.D. degree in Computer Science and Engineering from SRM Institute of Science and Technology, Chennai, India.

He is currently an Assistant Professor (Senior Grade) with the School of Computer Science and Engineering, Vellore Institute of Technology (VIT), Chennai, India. He has been in the teaching profession for more than 13 years and has presented over 60 papers in national and international journals, conferences, and symposiums. His research interests include deep learning, cybersecurity, and edge computing.



G. SARANYA received the Ph.D. degree in Computer Science and Engineering from SRM Institute of Science and Technology, Chennai, India.

She is currently an Assistant Professor (Senior Grade) with the School of Computer Science and Engineering, Vellore Institute of Technology (VIT), Chennai, India. She has been in the teaching profession for more than 12 years and has presented over 50 papers in national and international journals, conferences, and symposiums. Her research interests include deep learning and edge computing.



HIMESWAR POTNURU is currently pursuing the Bachelor of Technology (B.Tech.) degree in Electronics and Communication Engineering from Vellore Institute of Technology, Chennai, India.

He has worked with the Combat Vehicles Research and Development Establishment (CVRDE), DRDO, Chennai, on projects involving embedded systems. His academic and research interests include the Internet of Things (IoT), embedded systems, and intelligent hardware design.



DISHA DANIEL is currently pursuing the Bachelor of Technology (B.Tech.) degree in Electronics and Communication Engineering from Vellore Institute of Technology, Chennai, India.

Her research focuses on the intersection of resource-constrained computing, signal processing, and secure decentralized systems. She has contributed to the development of lightweight security frameworks for edge devices and optimized small language models for disaster response applications. Her work also explores multimodal maternal health monitoring systems and privacy-preserving decentralized health record platforms. Her research contributions include publications in high-impact journals and conference proceedings such as *IEEE Access* and *Frontiers in Digital Health*. Her areas of interest include computer architecture, edge AI, biomedical signal processing, and blockchain-enabled security.

• • •