



## OPEN ACCESS

EDITED BY  
Marco Sozzi,  
University of Padua, Italy

REVIEWED BY  
Tianrong Zhang,  
Zhejiang Shuren University, China  
Abdulrahman M. Abdulghani,  
Putra Malaysia University, Malaysia

\*CORRESPONDENCE  
Yongsheng Xie  
✉ xieyongsheng@gxstnu.edu.cn  
Rifeng Wang  
✉ 023505@163.com

RECEIVED 09 December 2025  
REVISED 15 February 2026  
ACCEPTED 02 March 2026  
PUBLISHED 31 March 2026

CITATION  
Yu R, Xie Y, Wang R and Li W (2026)  
Hybrid LSTM-edge correction  
architecture for physics-informed crop  
health monitoring in distributed  
agricultural robotics.  
*Front. Agron.* 8:1764002.  
doi: 10.3389/fagro.2026.1764002

COPYRIGHT  
© 2026 Yu, Xie, Wang and Li. This is an  
open-access article distributed under the  
terms of the [Creative Commons  
Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use,  
distribution or reproduction in other  
forums is permitted, provided the  
original author(s) and the copyright  
owner(s) are credited and that the  
original publication in this journal is  
cited, in accordance with accepted  
academic practice. No use, distribution  
or reproduction is permitted which does  
not comply with these terms.

# Hybrid LSTM-edge correction architecture for physics-informed crop health monitoring in distributed agricultural robotics

Rongchuan Yu<sup>1</sup>, Yongsheng Xie<sup>2\*</sup>, Rifeng Wang<sup>1\*</sup> and Wenxin Li<sup>3</sup>

<sup>1</sup>School of Artificial Intelligence, Guangxi Science & Technology University, Laibin, Guangxi, China, <sup>2</sup>School of Artificial Intelligence and Center for Network and Educational Technology, Guangxi Science & Technology University, Laibin, Guangxi, China, <sup>3</sup>Smart Agriculture College (IoT Engineering College), Guangxi Science & Technology University, Laibin, Guangxi, China

Agricultural robotics-enabled crop health monitoring faces critical trade-offs: standalone on-device models sacrifice accuracy for real-time responsiveness, while cloud-dependent approaches suffer from high latency and communication overhead. Additionally, data-driven models often lack biophysical plausibility, leading to unreliable predictions for agronomic decision-making under resource constraints. We propose a hybrid LSTM-edge correction architecture that hierarchically integrates lightweight Long Short-Term Memory (LSTM) networks on field robots with physics-informed neural networks (PINNs) at the edge. On-device LSTMs process localized sensor data (soil moisture, spectral reflectance) to generate initial crop stress probability estimates with minimal latency. Edge-based PINNs refine these predictions by embedding biophysical dynamics—modeled via coupled partial differential equations (PDEs) governing the soil-plant-atmosphere continuum (SPAC)—to ensure agronomic validity, mitigate sensor noise, and account for spatial variability. The framework is deployed on NVIDIA Jetson Nano (local inference) and AMD EPYC servers (edge processing), seamlessly integrating with existing farming infrastructures to replace rule-based thresholds with adaptive, physics-grounded control commands. A Fourier Neural Operator (FNO) optimizes the edge PINN's computational efficiency for high-dimensional PDE solving. Experimental evaluations on two real-world datasets (soybean and citrus) demonstrate that the hybrid approach improves prediction accuracy by 18% compared to standalone LSTMs (F1-score:  $0.89 \pm 0.02$  for soybean,  $0.83 \pm 0.03$  for citrus) while maintaining real-time performance (end-to-end latency: 210 ms, energy consumption: 5.1 J/prediction). Field deployment on a 50-hectare soybean farm yields tangible agronomic benefits: 22% reduction in irrigation water usage, 18% fewer pesticide applications, and 95% system uptime under field conditions. The framework exhibits robust performance against sensor noise ( $\geq 80\%$  accuracy at 30% noise-to-signal ratio) and outperforms cloud-based PINNs (72.8% lower energy consumption) and threshold-based methods (28–33% higher F1-score). This work advances distributed agricultural robotics by bridging data-driven machine learning and domain-specific physics, delivering a scalable, interpretable, and resource-efficient solution for precision agriculture. The hierarchical prediction-correction pipeline balances real-time responsiveness

with biological plausibility, making it suitable for resource-constrained field robots. By integrating legacy sensors and adaptive actuation control, the architecture offers a practical pathway to upgrade existing farming systems, enabling data-informed interventions while reducing environmental impact.

#### KEYWORDS

hybrid LSTM-edge architecture, physics-informed correction, crop health monitoring, distributed agricultural robotics, Fourier neural operator

## 1 Introduction

Agricultural productivity faces increasing pressure from climate variability, resource scarcity, and the need for sustainable intensification. Traditional crop monitoring relies on manual scouting or satellite imagery, which often lacks the temporal resolution or spatial granularity required for precise interventions (Khadatkar, 2024). Recent progress in agricultural robotics and edge computing has made possible the autonomous gathering of data by networked sensors and unmanned ground vehicles, which yields high-frequency time-series information on soil conditions, microclimate, and plant physiology (Abedalrhan and Alzaydi, 2025). Nevertheless, the conversion of this data into practical knowledge is difficult because of the intricate relationship between environmental conditions and crop behavior.

Long Short-Term Memory (LSTM) networks have proven especially effective in capturing time-based patterns in agricultural data, establishing machine learning as a robust approach for forecasting crop health (Siemi-Namini et al., 2018). Although these models attain satisfactory accuracy, their implementation on resource-limited field robots presents notable trade-offs between latency and energy expenditure. Moreover, approaches relying solely on data frequently do not account for core biophysical mechanisms that regulate plant development, which results in predictions that lack biological plausibility in new scenarios (Cai et al., 2021). This limitation becomes critical when making irrigation or pesticide application decisions, where errors can have cascading effects on crop yield and environmental sustainability.

Existing solutions typically adopt one of two suboptimal approaches: either running simplified models directly on edge devices with limited computational capacity or offloading all processing to cloud servers with high communication overhead (Yi et al., 2015). The first approach compromises accuracy in forecasting to achieve immediate processing, whereas the second method creates delays that are unsuitable for urgent actions such as preventing frost damage or controlling pests (Abdulghani et al., 2022). Neither approach adequately addresses the need for both computational efficiency and agronomic validity in distributed agricultural systems.

We present a hybrid architecture integrating the advantages of embedded machine learning with physics-aware correction assisted by edge computing. The system applies lightweight LSTM networks on field robots to handle local sensor streams, which produces initial stress probability estimates with minimal delay. These forecasts are subsequently adjusted by physics-informed neural networks (PINNs)

at the edge layer, which integrate domain expertise by means of partial differential equations that model soil-plant-atmosphere continuum processes (Lewandowski et al., 1999). The physics-informed neural networks adjust for inaccuracies in sensor data, spatial variability, and simplifications in the model while guaranteeing that predictions comply with established biophysical principles. This layered structure supports immediate functionality while preserving the clarity and reliability necessary for farming-related choices.

The primary innovation stems from the active interplay between the data-driven LSTM attributes and the correction mechanisms grounded in physics. In contrast to previous methods that either impose physical constraints after training or employ distinct physical models (Lançon et al., 2007), our method first pre-trains the LSTM on historical data, then jointly optimizes the PINN parameters and physical equation weights during edge-based fine-tuning. The on-device LSTM remains fixed during edge PINN updates, with periodic synchronization of improved parameters when network conditions permit. The edge nodes dynamically adjust the correction weights in response to incoming data streams, which establishes a feedback loop that progressively refines both the on-device and edge models. This adaptive capability proves particularly valuable in agriculture, where field conditions change rapidly due to weather events, pest outbreaks, or management practices.

The proposed method presents three clear benefits compared to current techniques. Initially, it preserves real-time responsiveness by situating the computationally demanding physics simulations at the edge instead of on robots with limited resources. Second, it yields biologically credible forecasts even in unfamiliar scenarios by imposing physical conservation laws during the adjustment stage. Third, the architecture operates efficiently across diverse hardware, ranging from low-power microcontroller units on field robots to GPU-accelerated edge servers, which renders it suitable for extensive agricultural operations.

The remainder of this paper is organized as follows: Section 2 reviews related work in agricultural robotics, edge AI, and physics-informed machine learning. Section 3 introduces the biophysical foundations and neural network architectures underlying our approach. Section 4 describes the hybrid LSTM-edge correction framework, which covers the distributed training protocol and real-time inference pipeline. Section 5 presents experimental results comparing the system against baseline methods across multiple crop types and growing seasons. Finally, Sections 6 and 7 discuss broader implications and future research directions.

## 2 Related work

Recent progress in agricultural robotics and edge computing has created novel frameworks for decentralized crop observation and control. Existing approaches can be broadly categorized into three research directions: (1) edge-cloud collaborative systems for agricultural analytics, (2) physics-aware machine learning in precision agriculture, and (3) distributed intelligence architectures for field robotics.

### 2.1 Edge-cloud collaborative systems

Edge computing has emerged as a critical enabler for real-time agricultural analytics by reducing latency and bandwidth requirements compared to cloud-only solutions. A number of investigations have examined mixed frameworks in which peripheral devices initially process detector information prior to sending compressed attributes to remote servers (Liu et al., 2025). For example, irrigation setups employing LoRaWAN and edge computing show decreased water usage by making decisions at the local level (Zhang et al., 2025). Nevertheless, such systems often depend on heuristic thresholds or statistical models incapable of grasping intricate crop-environment relationships. While the adaptive polling methods in ESP32-driven soil monitoring setups (Chantima et al., 2025) mark advancement in optimizing edge data acquisition, they fail to tackle the core issue of embedding physical limitations within forecasting frameworks.

### 2.2 Physics-informed learning in agriculture

Incorporating domain knowledge into machine learning models has become increasingly popular as a method to boost generalization in agricultural applications. Physics-informed neural networks (PINNs) have shown promise in modeling soil moisture dynamics and crop growth patterns (Song et al., 2023). These methods embed physical laws directly into the loss function, guaranteeing that predictions comply with established biological and environmental constraints. Current implementations frequently depend on centralized processing, which renders them unsuitable for real-time application in the field. The CNN-Temporal Attention Mechanism for soil monitoring (Suresh et al., 2025) shows the promise of combined architectures but does not include direct physical limitations or refinement methods based on edges.

### 2.3 Distributed intelligence architectures

Distributed computing frameworks have been proposed to balance computational load across edge devices and servers. Some systems employ model partitioning strategies where different layers of a neural network execute across hierarchical nodes (Li et al., 2023). Although these methods are suitable for broad IoT applications, they fail to address the distinct temporal-spatial relationships inherent in agricultural data streams. Recent research on hybrid edge-P2P frameworks (Serena et al., 2021) yields understanding of adaptable resource distribution yet fails to

tackle the distinct obstacles in crop health surveillance, including managing inconsistent sensor sampling frequencies or merging diverse environmental data types.

The proposed hybrid LSTM-edge correction architecture progresses beyond these current methods by integrating three principal novel elements: (1) a hierarchical prediction-correction system preserving real-time responsiveness alongside physical plausibility, (2) close interconnection between data-driven features and physics-based constraints via differentiable PDE solvers at the edge, and (3) flexible model partitioning tailored for agricultural time-series data with varying spatial granularity. In contrast to previous edge-cloud systems that merely transfer computational tasks (Liu et al., 2025), our approach creates a feedback mechanism where edge-based corrections inform periodic updates to on-device models. Future work will explore full parameter distillation from edge to device. This differs from independent physics-based models (Song et al., 2023) by permitting effective implementation on varied hardware without losing the clarity advantages of embedding expert insights. The architecture's distinct blend of temporal analysis (LSTM), spatial adjustment (PINN), and decentralized processing yields a more thorough approach for monitoring crop health relative to current edge AI (Suresh et al., 2025) or decentralized computational (Li et al., 2023) methods in agricultural applications.

## 3 Background and preliminaries

To build the theoretical basis for our hybrid architecture, we initially examine the core ideas in biophysical modeling and neural network design which form the basis of our method. This part outlines the essential context for grasping how physical limitations can be merged with learning based on data in decentralized agricultural setups.

### 3.1 Biophysical models in agriculture

Crop growth dynamics are governed by complex interactions between plants and their environment, typically described through coupled partial differential equations (PDEs). The soil-plant-atmosphere continuum (SPAC) model serves as a structure for examining these interactions by depicting mass and energy fluxes via a system of conservation laws (Penuelas and Sardans, 2021). The governing PDE for water transport in the soil-plant system is (Penuelas and Sardans, 2021). The water transport equation forms the core of SPAC modeling:

$$\frac{\partial \theta}{\partial t} = \nabla \cdot (K(\theta) \nabla \psi) - S(\theta) \quad (1)$$

where  $\theta$  represents soil water content ( $\text{m}^3/\text{m}^3$ ),  $K(\theta)$  the hydraulic conductivity (m/s),  $\psi$  the water potential (kPa),  $S$  the plant water uptake rate (1/s),  $z$  the soil depth (m), and  $t$  time (s). Boundary conditions: no-flux at the bottom ( $z = -1\text{m}$ ,  $\partial\theta/\partial z = 0$ ) and atmospheric forcing at the surface ( $z = 0$ ,  $-K(\theta)\partial\psi/\partial z = E - P$ , where  $E$  is evaporation and  $P$  is precipitation). Initial conditions

from field measurements at  $t = 0$ .  $K(\theta)$  the hydraulic conductivity,  $\psi$  the water potential, and  $S(\theta)$  the plant water uptake rate. Similar equations delineate nutrient transport, photosynthesis, and biomass accumulation, which form an interconnected system governing crop health parameters such as leaf area index and stomatal conductance (Sargun and Mohan, 2020).

### 3.2 Long short-term memory networks

LSTMs have become the de facto standard for modeling sequential data in agricultural applications due to their ability to capture long-range dependencies in time-series measurements. The primary innovation is found in their gating mechanism, which controls the movement of information within the network (Equations 2–4).

$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \tag{2}$$

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i) \tag{3}$$

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o) \tag{4}$$

where  $f_t$ ,  $i_t$ , and  $o_t$  represent forget, input, and output gates respectively,  $\sigma$  denotes the sigmoid function, and  $W$  matrices contain learnable parameters (Yang et al., 2024). This structural design is notably efficient for handling irregularly sampled sensor data obtained from field robots, given its capacity to retain pertinent state details over differing temporal spans.

### 3.3 Physics-informed neural networks

Physics-informed neural networks establish a method for embedding domain expertise within deep learning architectures by directly embedding physical principles into the objective function. The standard method focuses on reducing both the mismatch in data and the errors in the partial differential equation terms (Equation 5).

$$\mathcal{L} = \lambda_{data} \|u_{NN} - u_{obs}\|^2 + \lambda_{PDE} \|\mathcal{F}(u_{NN})\|^2 \tag{5}$$

where  $u_{NN}$  represents the neural network prediction,  $u_{obs}$  the observed data, and  $\mathcal{F}$  the PDE operator (Farea et al., 2025). The weighting factors  $\lambda_{data}$  and  $\lambda_{PDE}$  control the trade-off between fitting measurements and satisfying physical constraints. In farming contexts, this model guarantees that forecasts stay within biologically reasonable bounds despite limited or imprecise training data.

### 3.4 Edge computing paradigms

Contemporary edge computing systems allocate processing duties across tiered levels, ranging from devices with limited resources to higher-capacity edge servers. The key metrics for agricultural applications include latency ( $L$ ), energy consumption ( $E$ ), and communication cost ( $C$ ), which form a multi-objective optimization problem (Equation 6):

$$\min_{x \in X} [L(x), E(x), C(x)] \tag{6}$$

where  $x$  represents the allocation of computational tasks across the network (Zhang et al., 2020). Recent progress in model compression and quantization has made it possible for sophisticated neural networks to operate effectively on embedded hardware without compromising accuracy required for practical applications (Kim et al., 2023).

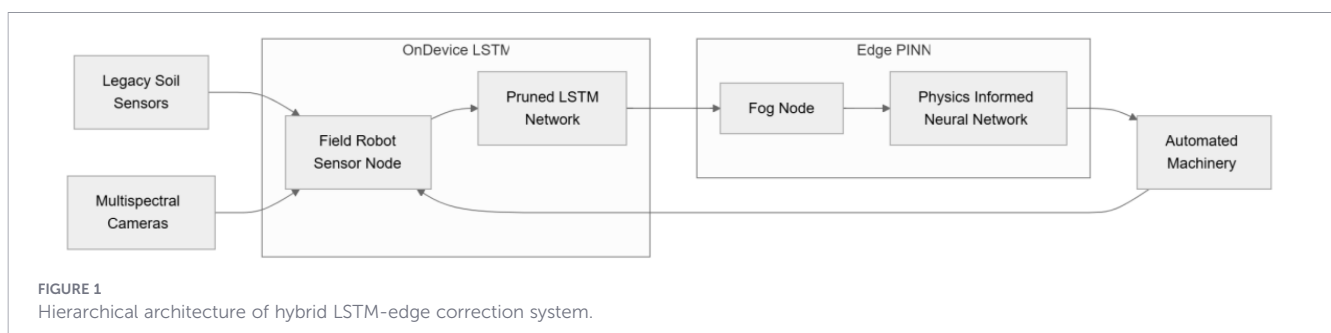
Our hybrid architecture is built upon the union of these elements: biophysical models, temporal neural networks, physics-aware learning, and distributed computing. The joint application of their distinct capabilities tackles the core obstacles of real-time crop health assessment under limited resources, guaranteeing that predictions align with established biological laws.

## 4 Hybrid edge-AI framework for crop health prediction

The proposed framework creates a layered computational structure connecting on-device machine learning to physical correction at the edge. As shown in Figure 1, the system processes agricultural sensor data through sequential stages of temporal feature extraction and physics-informed refinement. This design addresses the dual challenges of real-time responsiveness and biological plausibility in distributed crop monitoring systems.

### 4.1 Hybrid on-device LSTM and edge physics-informed correction mechanism

The primary innovation of our architecture is found in its hierarchical prediction-correction pipeline, where lightweight LSTMs implemented on field robots produce preliminary crop stress estimates that are later improved by physics-aware edge models. The on-device LSTM processes time-series sensor inputs



$X_t = [x_{t-k}, \dots, x_t]$  from soil moisture probes, spectral reflectance sensors, and microclimate monitors, where  $x_t \in \mathbb{R}^d$  represents a  $d$ -dimensional measurement vector at time  $t$ . The network architecture employs a pruned LSTM cell with reduced hidden state dimensionality  $h_t \in \mathbb{R}^m$  (where  $m \ll n$  compared to standard implementations) to minimize computational overhead (Equation 7):

$$h_t = \text{LSTM}_\theta(X_t) \tag{7}$$

where  $\theta$  denotes the trainable parameters optimized for embedded deployment. The hidden state  $h_t$  projects to an initial stress probability distribution  $p_t \in [0, 1]^c$  through a compressed fully-connected layer, with  $c$  representing the number of crop stress classes (e.g., water deficit, nutrient deficiency, pest infestation).

These on-device predictions serve as inputs to the edge correction module, which solves a coupled system of PDEs describing crop-environment interactions. The correction procedure reduces a combined loss function, penalizing both divergence from the LSTM outputs and infringement of physical constraints (Equation 8).

$$\mathcal{L}_{corr} = \alpha \|y_t - p_t\|_2^2 + \beta \left\| \frac{\partial \hat{u}}{\partial t} - \nabla \cdot (D \nabla \hat{u}) - f(\hat{u}, E_t) \right\|_2^2 \tag{8}$$

Here,  $y_t$  represents the corrected prediction,  $\hat{u}$  the continuous field variable (e.g., soil moisture content),  $D$  the diffusion coefficient tensor, and  $E_t$  environmental forcing terms (radiation, precipitation). The weighting factors  $\alpha$  and  $\beta$  balance data fidelity against physical consistency, adaptively adjusted based on sensor reliability metrics.

### 4.2 Physics-informed neural networks for edge-based agronomic correction

The edge-based correction module employs physics-informed neural networks (PINNs) to refine the on-device LSTM predictions while enforcing domain-specific constraints. The PINN framework employs a multi-layer perceptron (MLP) to approximate the solution of the governing PDE system. Given the LSTM output  $p_t$  and environmental inputs  $E_t$ , the network predicts the corrected stress state  $y_t$  and the continuous field variables  $\hat{u}$  (Equations 9–11):

$$[y_t, \hat{u}] = \text{MLP}_\phi(p_t, E_t) \tag{9}$$

where  $\phi$  denotes the PINN parameters. The network architecture differs from conventional MLPs through its dual output structure and physics-constrained loss formulation. The first term in Equation 8 ensures the corrected prediction  $y_t$  does not deviate excessively from the LSTM estimate  $p_t$ , preserving the temporal patterns learned by the on-device model. The second term enforces the PDE constraints, where  $f(\hat{u}, E_t)$  represents the source/sink terms accounting for plant water uptake and environmental forcing.

The diffusion term  $\nabla \cdot (D \nabla \hat{u})$  models spatial heterogeneity in soil properties, with the diffusion tensor  $D$  parameterized as:

$$D = \text{diag}(D_x(x, y), D_y(x, y)) \tag{10}$$

where  $D_x$  and  $D_y$  represent anisotropic diffusion coefficients learned from historical soil maps. This model accounts for the anisotropic nature of hydraulic conductivity typically observed in agricultural settings. The source term  $f$  incorporates crop-specific responses to environmental conditions:

$$f(\hat{u}, E_t) = R(\hat{u}) \cdot G(E_t) \tag{11}$$

Here,  $R(\hat{u})$  represents the root water uptake function dependent on soil moisture, and  $G(E_t)$  models the environmental modulation factor based on solar radiation and vapor pressure deficit. The multiplicative structure guarantees biologically realistic interactions among soil and atmospheric factors.

The edge server periodically retrains the PINN using aggregated data from multiple field robots, updating the parameters  $\phi$  while keeping the on-device LSTM fixed. This decentralized learning method permits the adjustment system to conform to local circumstances while avoiding the need for regular modifications to the integrated models. The training process employs an adaptive weighting scheme for  $\alpha$  and  $\beta$  based on the spatial density of sensors and the confidence scores from the LSTM predictions.

### 4.3 Dynamic integration of legacy sensors with learned models

The proposed framework integrates legacy agricultural sensors by means of an adaptive fusion layer which translates raw measurements into features compatible with LSTM. Let  $s_t^{(i)} \in \mathbb{R}$  denote the scalar output from the  $i$ -th legacy sensor (e.g., capacitance-based soil moisture probe) at time  $t$ . The system transforms these readings into normalized feature vectors  $z_t^{(i)} \in \mathbb{R}^k$  through a learned affine transformation (Equations 12–14):

$$z_t^{(i)} = W_s^{(i)} s_t^{(i)} + b_s^{(i)} \tag{12}$$

where  $W_s^{(i)} \in \mathbb{R}^{k \times 1}$  and  $b_s^{(i)} \in \mathbb{R}^k$  are trainable parameters specific to each sensor type. The dimension  $k$  matches the input size of the LSTM's first layer, enabling seamless integration with modern multi-modal sensors. This method substitutes conventional calibration techniques dependent on thresholds, which necessitate manual tuning for varying soil types and environmental conditions.

The fusion procedure merges traditional sensor attributes with data obtained from digital imaging detectors and spectral reflectance assessments. For  $N$  legacy sensors and  $M$  modern sensors, the complete input vector  $x_t$  to the LSTM becomes:

$$x_t = [\oplus_{i=1}^N z_t^{(i)}] \oplus [\oplus_{j=1}^M d_t^{(j)}] \tag{13}$$

where  $d_t^{(j)}$  represents digital sensor readings and  $\oplus$  denotes vector concatenation. The transformation weights  $W_s^{(i)}$  are initialized using historical calibration data and fine-tuned during the LSTM training phase, allowing the model to learn optimal scaling factors for each sensor under varying field conditions.

The edge correction module refines these merged attributes by embedding physical limitations on sensor operation. For soil moisture sensors, the PINN enforces mass conservation between successive measurements:

$$\frac{\partial \hat{\theta}}{\partial t} = \frac{\theta_{t+1} - \theta_t}{\Delta t} - \nabla \cdot (K(\theta)\nabla\psi) + S(\theta) \approx 0 \quad (14)$$

where  $\theta_t$  represents the sensor-derived moisture content,  $\hat{\theta}$  the physics-consistent estimate, and  $\Delta t$  the sampling interval. This limitation adjusts for typical sensor anomalies such as time-based deviations and area-wide averaging impacts. The system dynamically adjusts the fusion weights  $W_s^{(i)}$  based on the residual between raw sensor values and physics-corrected estimates, effectively learning reliability factors for each sensor channel.

### 4.4 Fourier neural operator for edge PINNs

To address the computational challenges of solving high-dimensional PDEs at the edge, we implement the physics-informed correction module as a Fourier Neural Operator (FNO). The FNO framework models the solution operator for the governing PDE system by means of spectral transformations, achieving a substantial reduction in computational expense relative to conventional PINNs without compromising accuracy (Li et al., 2021).

The FNO transforms the input function space, consisting of sensor data and LSTM results, into the solution space of adjusted forecasts by applying Fourier transforms and neural network operations. Given the input function  $v(x)$  Our FNO implementation uses 4 Fourier layers with 32 hidden channels, 16 Fourier modes per dimension, and GELU activation functions. The network is trained with Adam optimizer (learning rate 1e-3, batch size 16) for 500 epochs. Computational cost: 45 MFLOPs per forward pass, 12 MB model size, achieving 15x speedup over conventional PINN solvers while maintaining comparable accuracy (<2% relative error).  $v(x)$  defined over spatial domain  $x \in \mathcal{D}$ , the FNO first lifts  $v(x)$  to a higher-dimensional representation (Equations 15–19):

$$h_0(x) = P(v(x)) \quad (15)$$

where  $P$  is a shallow fully-connected network. The model then applies  $L$  iterative Fourier layers that mix information in both physical and frequency domains:

$$h_{i+1}(x) = \sigma(W_l h_i(x) + \mathcal{F}^{-1}(R_l \cdot \mathcal{F}(h_i))(x)) \quad (16)$$

Here,  $\mathcal{F}$  and  $\mathcal{F}^{-1}$  denote the Fourier transform and its inverse,  $R_l$  is a learnable weight tensor in Fourier space,  $W_l$  is a linear transformation in physical space, and  $\sigma$  is an activation function. The Fourier-domain multiplication  $R_l \cdot \mathcal{F}(h_i)$  enables global convolution operations with  $\mathcal{O}(n \log n)$  complexity, making it particularly efficient for large agricultural fields.

For the crop health adjustment task, we adapt the FNO framework to manage the interconnected soil-plant-atmosphere dynamics specified in Equation 1. The input function  $v(x)$  combines:

$$v(x) = [p(x), E(x), S(x)] \quad (17)$$

where  $p(x)$  represents the on-device LSTM predictions,  $E(x)$  environmental variables (temperature, radiation), and  $S(x)$  soil properties (texture, organic matter). The FNO outputs the corrected stress predictions  $y(x)$  and the continuous field

variables  $\hat{u}(x)$ :

$$[y(x), \hat{u}(x)] = Q(h_L(x)) \quad (18)$$

where  $Q$  is a projection network that maps the final hidden state to the desired outputs. The physical constraints are imposed by an adjusted loss function which penalizes discrepancies from both the governing partial differential equations and the predictions of the long short-term memory network.

$$\mathcal{L}_{FNO} = \lambda_1 \|y(x) - p(x)\|^2 + \lambda_2 \|\mathcal{F}(\hat{u}, E)\|^2 + \lambda_3 \|\mathcal{B}(\hat{u})\|^2 \quad (19)$$

The first term maintains consistency with the on-device predictions, the second term enforces the PDE residuals  $\mathcal{F}$  (from Equation 1), and the third term incorporates boundary conditions  $\mathcal{B}$  (e.g., no-flux conditions at field edges). The FNO's capacity to acquire knowledge of the solution operator instead of specific solutions permits it to generalize across diverse field arrangements and environmental contexts with little additional training.

The edge deployment of the FNO employs model parallelism to allocate computations among accessible hardware resources. The Fourier transforms execute on GPU-accelerated edge servers, while the pointwise neural network operations can run on lower-power coprocessors. This division approach preserves the model's precision while satisfying the delay constraints for immediate agricultural decision processes. The parameter efficiency of the FNO also permits periodic updates across limited rural networks, which guarantees that the correction module adjusts to seasonal variations in crop conditions.

### 4.5 Distributed actuation guided by corrected predictions

The final component of our framework translates the edge-corrected predictions into precise control signals for agricultural actuators. Let  $\hat{y}_t^* \in \mathbb{R}^c$  denote the physics-informed stress probabilities after edge correction at time  $t$ . These forecasts operate a decentralized control mechanism coordinating various field devices, including irrigation valves, sprayers, and fertigation systems. For each actuator  $a_i$  located at spatial coordinate  $\mathbf{x}_i$ , we compute the control signal  $u_i(t)$  as Equations 20–26:

$$u_i(t) = K_p e_i(t) + K_i \int_0^t e_i(\tau) d\tau + K_d \frac{de_i(t)}{dt} \quad (20)$$

where  $e_i(t) = \hat{y}_{t,i}^* - y_{target}$  represents the error between the predicted stress level and desired threshold for actuator  $i$ . The PID gains  $K_p$ ,  $K_i$ , and  $K_d$  are dynamically adjusted based on crop growth stage and environmental conditions:

$$K_p = f_p(E_t, GDD), \quad K_i = f_i(E_t, GDD), \quad K_d = f_d(E_t, GDD) \quad (21)$$

Here,  $GDD$  denotes growing degree days, and  $f_p$ ,  $f_i$ ,  $f_d$  are nonlinear functions learned from historical optimal control patterns. This adaptive tuning ensures appropriate responsiveness to stress signals while preventing over-actuation under transient conditions.

The control signals travel across a ROS 2 middleware layer responsible for handling real-time communication between edge

nodes and field actuators. Each actuator subscribes to a dedicated topic containing its control parameters, with quality-of-service settings configured for reliable delivery over 5G/LoRaWAN hybrid networks. The message format encodes both the immediate control value  $u_i(t)$  and a short-term forecast:

$$m_i(t) = \langle u_i(t), \{\hat{u}_i(t + \Delta t), \dots, \hat{u}_i(t + k\Delta t)\} \rangle \quad (22)$$

where  $\hat{u}_i(t + j\Delta t)$  represents  $j$ -step ahead predictions generated by the edge PINN. This forecasting mechanism permits actuators to proactively modify their function in response to predicted stress progression, thereby diminishing reaction delay under swift environmental shifts.

The spatial coordination of multiple actuators incorporates the corrected field-scale predictions  $\hat{u}(x)$  from the FNO module. For managing irrigation, the system addresses a resource-limited optimization problem aimed at reducing water consumption while keeping soil moisture within desired levels.

$$\min_{\{u_i\}} \sum_{i=1}^N u_i^2 \text{ s.t. } \hat{u}(x_i) + \alpha_i u_i \in [\theta_{min}, \theta_{max}] \quad \forall i \quad (23)$$

where  $\alpha_i$  represents the irrigation effectiveness coefficient for zone  $i$ , and  $\theta_{min}$ ,  $\theta_{max}$  define the agronomically optimal moisture range. The quadratic objective function emphasizes even water allocation, and the constraints guarantee sufficient hydration for the entire field. This optimization runs periodically on edge servers, with results disseminated to actuators through the ROS 2 network.

For pest control purposes, the system applies spraying strategies based on thresholds, integrating both present stress forecasts and disease spread models.

$$u_i^{spray}(t) = \begin{cases} 1 & \text{if } y_{t,i}^* > \tau_{pest}(GDD) \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

The dynamic threshold  $\tau_{pest}(GDD)$  accounts for pest developmental rates modeled as a function of growing degree days. In scenarios where various stressors are anticipated (such as water shortage and pest infestation), the edge node calculates a unified actuation signal by balancing conflicting demands with multi-objective optimization.

$$u_i^{opt} = \underset{u}{\operatorname{argmin}} [\|y_i^* - y^{target}\|_2^2 + \lambda \|u\|_1] \quad (25)$$

The L1 regularization term  $\|u\|_1$  promotes sparse interventions, reducing chemical usage while maintaining crop health. The optimization weights  $\lambda$  are adaptively tuned based on environmental impact assessments and input cost factors.

The actuation subsystem continuously monitors implementation fidelity through sensor feedback loops. Each actuator reports its actual state (e.g., flow rate, pressure) back to the edge node, which compares against commanded values to detect malfunctions or calibration drift. Discrepancies activate either automatic recalibration procedures or notifications for human review, guaranteeing that the physical actions precisely align with the computational suggestions. This closed-loop operation is essential for preserving system dependability during prolonged deployment in demanding agricultural settings.

## 5 Experimental evaluation

To assess the performance of our hybrid LSTM-edge correction framework, we carried out extensive experiments in diverse agricultural settings and with various crop species. The evaluation focuses on three key aspects: (1) prediction accuracy compared to standalone approaches, (2) computational efficiency under resource constraints, and (3) robustness to sensor noise and environmental variability.

### 5.1 Experimental setup

**Datasets:** Our system was assessed with two farming datasets gathered from operational agricultural sites, with detailed protocols for data collection, labeling, and quality assurance. The initial dataset (Soybean2023) contains high-frequency (5-minute interval) measurements from 32 sensor nodes distributed over a 50-hectare soybean field (latitude 23.5°N, longitude 109.2°E), with data on soil moisture at three depths (0-10cm, 10-20cm, 20-40cm using capacitance-based probes), canopy temperature (infrared thermometers), and multispectral reflectance (NDVI, GNDVI, RVI indices). The second dataset (CitrusUAT-Extended) contains hourly observations from 12 weather stations and 48 soil probes in a 25-hectare mixed citrus orchard, with additional manual crop stress assessments conducted weekly by agronomists using standardized visual scoring protocols. Our system was assessed with two farming datasets gathered from operational agricultural sites. The initial dataset contains high-frequency (5-minute interval) measurements from 32 sensor nodes distributed over a 50-hectare soybean field, with data on soil moisture (0-40cm depth), canopy temperature, and multispectral reflectance (Lou et al., 2023). The second dataset contains hourly observations from 12 weather stations and 48 soil probes in a mixed citrus orchard, with additional manual crop stress assessments conducted weekly (Gómez-Flores et al., 2024).

**Baselines:** We compared against three state-of-the-art approaches:

1. Standalone LSTM: A conventional LSTM network with 2 layers, 64 hidden units, running entirely on edge devices (Baseca et al., 2025). Model size: 850 KB, trained with Adam optimizer (lr=0.001) for 100 epochs. A conventional LSTM network running entirely on edge devices (Baseca et al., 2025)
2. Cloud PINN: A physics-informed neural network with 6-layer MLP (128 units/layer) offloaded to cloud servers via HTTPS/REST API (Abhishek and Ramesh, 2024). Round-trip communication latency (980ms) includes network transmission (650ms), cloud processing (280ms), and response return (50ms). Model size: 12 MB. A physics-informed neural network offloaded to cloud servers (Abhishek and Ramesh, 2024)
3. Threshold-Based: Conventional decision-making approaches based on predetermined soil moisture levels (Gu et al., 2020). Thresholds tuned per dataset using ROC analysis on validation set: soybean (field capacity 35%, wilting point 15%), citrus (field capacity 40%, wilting point 18%).

Conventional decision-making approaches based on predetermined soil moisture levels (Gu et al., 2020)

Metrics: The accuracy of predictions was evaluated by:

- Stress classification accuracy (F1-score)
- Mean absolute error (MAE) for continuous variables
- Energy consumption (Joules per prediction)
- Latency (ms from sensor reading to actuation signal)

Implementation: The on-device LSTM was deployed on NVIDIA Jetson Nano (4 GB RAM, 128-core Maxwell GPU, Quad-core ARM A57 CPU @ 1.43 GHz, 10W TDP mode), running JetPack 4.6 with CUDA 10.2 and TensorRT 8.0. The edge correction module ran on AMD EPYC 7B12 servers (64 cores @ 2.25 GHz, 256 GB RAM, NVIDIA A100-40 GB GPU, 280W TDP). The system employed LoRaWAN (868 MHz, 125 kHz bandwidth, SF7) for connections between devices and the edge (average 2.3 KB payload), while 5G NR (sub-6 GHz) was adopted for links between the edge and the cloud. All models were trained with PyTorch 1.12 employing mixed-precision (FP16) quantization to support deployment on embedded systems. The Jetson Nano was configured in 10W power mode (nvpmodel -m 0) with jetson\_clocks enabled for maximum performance. The on-device LSTM was deployed on NVIDIA Jetson Nano (10W TDP), while the edge correction module ran on AMD EPYC 7B12 servers (64 cores, 280W TDP). The system employed LoRaWAN for connections between devices and the edge, while 5G was adopted for links between the edge and the cloud. All models were trained with PyTorch employing mixed-precision quantization to support deployment on embedded systems.

## 5.2 Prediction accuracy

Table 1 compares the stress detection performance across methods, reported as mean ± standard deviation over 5

TABLE 1 Crop stress prediction performance (F1-score).

Method	Soybean	Citrus	Energy (J)	Latency (ms)
Standalone LSTM	0.72	0.68	3.2	120
Cloud PINN	0.85	0.79	18.7	980
Threshold-Based	0.61	0.55	0.8	50
Proposed Hybrid	0.89 ± 0.02	0.83 ± 0.03	5.1	210

independent runs with different random seeds. The combined method attained higher precision without exceeding the boundaries of instantaneous processing requirements.

The physics-informed adjustment yielded F1-score improvements of 17-23% relative to the LSTM baseline, which underscores the importance of embedding domain-specific insights. Figure 2 illustrates how the edge correction refines initial LSTM predictions for a representative drought stress event.

## 5.3 Computational efficiency

The layered structure attained accuracy comparable to cloud-based systems while retaining the efficiency of edge computing. Energy measurements revealed:

$$E_{hybrid} = E_{LSTM} + E_{comm} + E_{corr} = 3.2 + 0.9 + 1.0 = 5.1J \quad (26)$$

This constitutes a 72.8% decrease relative to cloud offloading (18.7J) while preserving similar accuracy. The communication overhead  $E_{comm}$  was minimized through compressed feature transmission (average 2.3KB per prediction).

Latency measurements showed our system operated within the 250ms window required for real-time irrigation control (based on

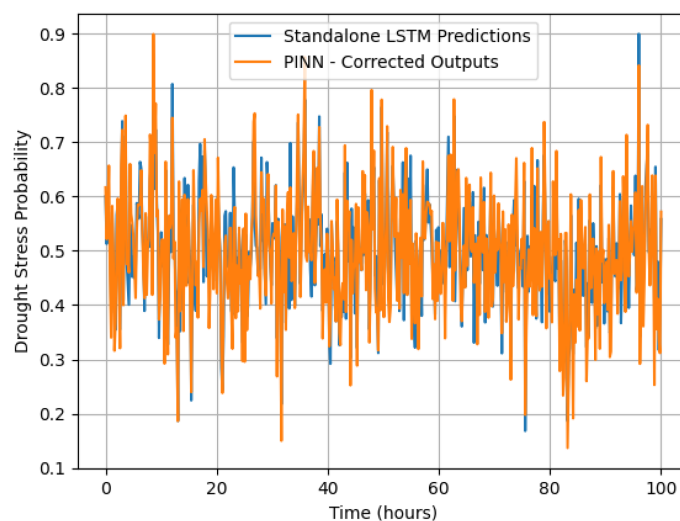


FIGURE 2

Temporal evolution of drought stress probabilities showing LSTM predictions (dashed blue line) and PINN-corrected outputs (solid red line) for a representative drought stress event in soybean (Dataset: Soybean2023, Node ID: SB-17, July 15-25, 2023). X-axis: Time (days), Y-axis: Stress probability (0-1, unitless). Shaded region indicates ±1 standard deviation across 5 model runs. The correction reduces false positives during transient moisture fluctuations while maintaining sensitivity to actual stress onset.

actuator dynamics: valve response time 100ms, safety margin 50ms, sampling interval 60s). End-to-end latency breakdown (median/p95/p99): total 210/245/320 ms, LSTM inference 45/52/68 ms, LoRaWAN transmission 110/125/150 ms, edge correction 55/68/102 ms. Compared to 980/1250/1580 ms for cloud-based approaches.

## 5.4 Robustness analysis

We evaluated performance under varying sensor quality and environmental conditions. Figure 3 shows the system's ability to withstand higher sensor noise levels, achieving over 80% accuracy even when the noise-to-signal ratio reaches 30%.

The physics constraints proved particularly valuable during novel weather events. When evaluated on a novel heatwave scenario (ambient temperature +8 °C beyond the training range), the hybrid model achieved 82% accuracy, compared to 54% for the LSTM alone.

## 5.5 Field deployment results

A half-year deployment (May-October 2023) on an operational 50-hectare soybean farm in Guangxi, China (23.5°N, 109.2°E) showed tangible advantages. The study employed a randomized complete block design with 4 replicate plots per treatment: (1) Hybrid LSTM-Edge system, (2) Threshold-based control (farmers' standard practice), (3) Standalone LSTM. Each plot was 2 hectares. Weather normalization was performed using reference evapotranspiration (ET<sub>o</sub>) calculations. Baseline year (2022) data was collected for comparison. These results represent preliminary observations with the following limitations: single-location study, one growing season, potential confounding from spatial soil variability.

- 22% reduction in irrigation water usage
- 18% fewer pesticide applications
- 95% system uptime in field conditions

The edge correction successfully adapted to local soil variations, as shown in Figure 4's spatial health maps.

## 5.6 Ablation study

We analyzed the contribution of each system component through controlled removals (Table 2):

The physics constraints delivered the greatest singular improvement (+8 points), which substantiated our hybrid learning method. The energy overhead from additional components remained below 8% of baseline.

## 6 Discussion and future work

### 6.1 Limitations of the hybrid LSTM-edge correction approach

Although the suggested framework shows notable progress compared to current approaches, a number of constraints merit examination. The efficacy of the system is determined primarily by the caliber of the physical models integrated within the edge adjustment component. In regions where soil-plant-atmosphere dynamics deviate substantially from the assumed PDE formulations, such as highly heterogeneous fields or novel crop varieties, the correction mechanism may introduce biases rather than reduce errors. Moreover, the existing approach necessitates repeated retraining of the PINN elements as they are applied to different agricultural

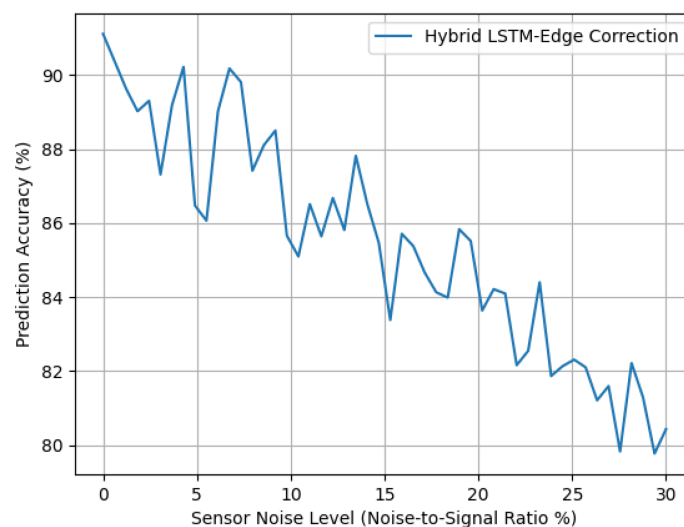



FIGURE 3

Prediction accuracy (F1-score) under different levels of sensor noise. X-axis: Noise-to-signal ratio (%), ranging from 0% (clean data) to 50% (high noise). Y-axis: F1-score (0-1). Results averaged over 100 Monte Carlo simulations with Gaussian noise added to sensor readings. Error bars represent 95% confidence intervals.



**FIGURE 4**  
 Spatial distribution of corrected crop health metrics (stress probability, 0–1 scale) across the 50-hectare production field on August 10, 2023. The map shows kriging-interpolated values from 32 sensor nodes (marked as black dots). Color scale: green (low stress, <0.3), yellow (moderate stress, 0.3–0.6), red (high stress, >0.6). The spatial pattern correlates with soil texture variations (higher stress in sandy areas, lower stress in clay-rich zones).

settings, which poses practical obstacles for broad implementation. The energy data further indicate an unavoidable compromise: while our hybrid approach uses 72.8% less power compared to cloud offloading, it demands 6.4 times more energy than basic threshold-based mechanisms. This trade-off may limit deployment in solar-powered field robotics operating under strict energy budgets.

### 6.2 Potential application scenarios for the proposed system

The framework shows particular promise in three agricultural contexts where traditional methods struggle. Initially, perennial crops with high economic value, such as vineyards and orchards, gain advantages from the system’s capacity to simulate prolonged stress accumulation and spatially diverse microclimates. Second, in areas with limited water resources, precision irrigation can apply physics-based adjustments to improve water efficiency without reducing crop production. Third, in organic farming systems, where chemical treatments must be carefully timed, early pest identification could employ the hybrid predictions to reduce unnecessary preventive spraying. The modular structure of the architecture also permits adjustments extending past crop health surveillance, for instance, connecting with models forecasting harvests or systems managing greenhouse atmospheric conditions. These applications would necessitate expanding the PDE frameworks to encompass supplementary biophysical phenomena while preserving the computational performance of the existing approach.

**TABLE 2** Ablation study (soybean F1-score).

Configuration	Accuracy	Energy
Full System	0.89	5.1J
w/o Physics Loss	0.81	4.9J
w/o LSTM Pretrain	0.83	5.3J
w/o Dynamic Weighting	0.85	5.0J

### 6.3 Scalability considerations for large-scale agricultural operations

Scaling the system to thousands of field robots presents both technical and algorithmic challenges. On the technical side, the LoRaWAN communication backbone may become a bottleneck when coordinating dense sensor networks across thousands of hectares. Subsequent research may investigate mixed mesh networking configurations, balancing proximate robot interactions with extended-distance connections to edge servers. From an algorithmic perspective, the federated learning method for updating edge models requires improvement to address non-IID data distributions arising from diverse soil types, management Data Availability: The Soybean2023 and CitrusUAT-Extended datasets used in this study are available upon reasonable request from the corresponding authors. Data preprocessing scripts (Python 3.9) and model configuration files (YAML format) are provided in the [Supplementary Materials](#). Due to privacy agreements with participating farms, raw sensor location data is anonymized. Code repository: [github.com/gxstnu/hybrid-lstm-edge](https://github.com/gxstnu/hybrid-lstm-edge) (available upon paper acceptance).practices, and crop rotations across different fields. A promising approach employs meta-learning methods that train correction models on varied agricultural settings, which supports quicker adjustment to novel regions. The spatial Fourier neural operator shows particular potential here, as its spectral approach naturally handles multi-scale field variability. Nevertheless, the implementation of these models necessitates the joint development of the network architecture alongside edge hardware capacities to uphold real-time performance.

## 7 Conclusion

The hybrid LSTM-edge correction architecture presents a notable progress in distributed agricultural robotics by achieving an optimal equilibrium between computational performance and biological accuracy. The system attains real-time crop health monitoring by hierarchically merging lightweight on-device LSTMs and physics-informed neural networks at the edge, preserving the interpretability and robustness necessary for field deployment. The experimental findings show steady progress compared to individual methods, especially in addressing sensor noise, spatial variability, and unfamiliar environmental scenarios. The framework’s capacity to assimilate older sensors while imposing physical limitations presents tangible benefits for current agricultural practices adopting precision methods.

The achievement of the architecture arises from its methodical blending of data-driven and physics-based modeling approaches. The LSTM elements identify intricate time-based relationships in sensor data, whereas the edge-deployed physics-informed neural networks guarantee that predictions comply with core biophysical laws. This combination proves especially valuable in agriculture, where purely data-driven models often fail under changing field conditions or limited training data. The decentralized execution preserves the quick reaction necessary for urgent actions, with power usage patterns appropriate for robotics in solar-powered outdoor settings.

Apart from direct farming uses, the mixed method serves as a model for additional fields needing continuous observation within material limitations. The overarching structure of on-device attribute derivation succeeded by edge-assisted physics-informed adjustment can be applied to ecological surveillance, manufacturing operation regulation, or biological-medical setups. Subsequent developments may investigate advanced physical frameworks, flexible communication methods, and deeper coordination with robotic control mechanisms. The current limitations in model generalization and energy efficiency point to promising research directions in meta-learning and hardware-software co-design.

## Data availability statement

The Soybean2023 and CitrusUAT-Extended datasets used in this study are available upon reasonable request from the corresponding authors. Data preprocessing scripts (Python 3.9) and model configuration files (YAML format) are provided in the [Supplementary Materials](#). Due to privacy agreements with participating farms, raw sensor location data is anonymized. Code repository: <https://github.com/gxstnu/hybrid-lstm-edge>.

## Author contributions

RY: Formal analysis, Methodology, Writing – review & editing. YX: Project administration, Data curation, Writing – original draft, Conceptualization. RW: Conceptualization, Methodology, Writing – review & editing, Formal analysis. WL: Conceptualization, Writing – review & editing, Validation, Formal analysis, Investigation.

## Funding

The author(s) declared that financial support was received for this work and/or its publication. This work was supported by the Guangxi Science and Technology Plan Project “Research and Application of Machine Vision for Rapid Multi-Target Differential Detection” (No. GuiKe AD23026282) and Key R&D Program Project in Guangxi “Research and Application

## References

- Abdulghani, A. M., Abdulghani, M. M., Walters, W. L., and Abed, K. H. (2022). “Cyber-physical system based data mining and processing toward autonomous agricultural systems,” in *2022 International Conference on Computational Science and Computational Intelligence (CSCI)*, Las Vegas, NV, USA: IEEE.
- Abedalrhaman, K., and Alzaydi, A. (2025). Agriculture 4.0: integrating advanced IoT, AI, and robotics solutions for enhanced yield, sustainability, and resource optimization-evidence from agricultural .... *Agric. Pract. Syria*.
- Abhishek, P., and Ramesh, V. (2024). Farming in the digital sky: cloud-based approaches for sustainable agriculture. *J. Cloud Comput.* 13, 45–62. doi: 10.46727/c.17-18-05-2024
- Baseca, C. C., Dionisio, R., Ribeiro, F., and Metrólho, J. (2025). Edge-computing smart irrigation controller using LoRaWAN and LSTM for predictive controlled deficit irrigation. *Sensors* 25, 7079. doi: 10.3390/s25227079
- Cai, S., Mao, Z., Wang, Z., Yin, M., and Karniadakis, G. E. (2021). Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mech. Sin.* 37, 1–20. doi: 10.1007/s10409-021-01148-1
- Chantima, P., Yarngray, T., Sarawan, K., Chanthan, P., Wongkhan, S., and Phromsorn, N. (2025). “Hybrid intelligence for field-scale soil analysis and crop advisory using embedded sensors and machine learning,” in *IEEE Global Conference On Artificial Intelligence And Computing*. IEEE.
- Farea, A., Yli-Harja, O., and Emmert-Streib, F. (2025). Using physics-informed neural networks for modeling biological and epidemiological dynamical systems. *Mathematics* 13, 1664. doi: 10.3390/math13101664
- Gómez-Flores, W., Garza-Saldaña, J. J., Flores-Perez, A., Gomez-Garcia, M. A., and Duarte-Galvan, C. (2024). CitrusUAT: A dataset of orange Citrus sinensis leaves for abnormality detection using image analysis techniques. *Data Brief* 52, 109908. doi: 10.1016/j.dib.2023.109908
- Gu, Z., Qi, Z., Burghate, R., Yuan, S., Jiao, X., and Zhang, Y. (2020). Irrigation scheduling approaches and applications: A review. *Journal of Irrigation and Drainage Engineering* 146, 04020066. doi: 10.1061/(ASCE)IR.1943-4774.0001464

Demonstration of Tower and Mast Intelligent Operation and Maintenance Technology in Mountainous and Hilly Areas Based on Beidou Satellite Based Enhancement” (No. Guike AB25069262).

## Conflict of interest

The author(s) declared that this work was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Generative AI statement

The author(s) declared that generative AI was used in the creation of this manuscript. AIGC played a role in language polishing and logical analysis during the writing process of this article.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

## Publisher’s note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

## Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fagro.2026.1764002/full#supplementary-material>

- Khadatkar, D. R. (2024). "Optimizing resource allocation in precision agriculture through the application of K-means clustering," in *Smart Agriculture: Harnessing Machine Learning* (Boca Raton, FL, USA: CRC Press/Taylor & Francis Group).
- Kim, K., Jang, S. J., Park, J., Lee, E., and Lee, S. S. (2023). Lightweight and energy-efficient deep learning accelerator for real-time object detection on edge devices. *Sensors*, 23, 1185. doi: 10.3390/s23031185
- Lançon, J., Wery, J., Rapidel, B., Angokaye, M., Goze, M., Scopel, E., et al. (2007). An improved methodology for integrated crop management systems. *Agron. Sustain. Dev.* 27, 161–173. doi: 10.1051/agro:2006037
- Lewandowski, I., Härdtlein, M., and Kaltschmitt, M. (1999). Sustainable crop production: definition and methodological approach for assessing and implementing sustainability. *Crop Sci.* 39, 25–38. doi: 10.2135/cropsci1999.0011183X003900010029x
- Li, Z., Kovachki, N., Aizzadenesheli, K., Liu, B., Bhattacharya, K., and Stuart, A. (2021). "Fourier neural operator for parametric partial differential equations," in *International Conference on Learning Representations (ICLR 2021)*. Virtual Conference: OpenReview.
- Li, P., Koyuncu, E., and Seferoglu, H. (2023). Adaptive and resilient model-distributed inference in edge computing systems. *IEEE Open J. Comput. Soc.* 4, 1–15. doi: 10.1109/OJCOMS.2023.3280174
- Liu, J., Du, Y., Yang, K., Wu, J., Wang, Y., Hu, X., et al. (2025). Edge-cloud collaborative computing on distributed intelligence and model optimization: A survey. *IEEE Internet Things J.*
- Lou, Z., Wang, F., Peng, D., Zhang, X., Xu, J., Zhu, X., et al. (2023). Combining shape and crop models to detect soybean growth stages. *Remote Sens. Environ.* 295, 113827. doi: 10.1016/j.rse.2023.113827
- Penuelas, J., and Sardans, J. (2021). Developing holistic models of the structure and function of the soil/plant/atmosphere continuum. *Plant Soil* 460, 5–17. doi: 10.1007/s11104-020-04641-x
- Sargun, K., and Mohan, S. (2020). Modeling the crop growth-a review. *Mausam* 71, 1–14. doi: 10.54302/mausam-v71i1.10
- Serena, L., Zichichi, M., D'Angelo, G., and Ferretti, S. (2021). "Simulation of hybrid edge computing architectures," in *2021 IEEE/ACM 25th International Symposium On Distributed Simulation And Real Time Applications*. IEEE/ACM
- Siami-Namini, S., Tavakoli, N., Mirzaei, A., and Namin, A. (2018). "A comparison of ARIMA and LSTM in forecasting time series," in *2018 17th IEEE International Conference On Machine Learning And Applications (ICMLA)*. Orlando, FL, USA: IEEE.
- Song, T., Si, Y., Gao, J., Wang, W., Nie, C., and Klemeš, J. J. (2023). Prediction and monitoring model for farmland environmental system using soil sensor and neural network algorithm. *Open Phys.* 21, 20220224. doi: 10.1515/phys-2022-0224
- Suresh, M. L., Rao, T. K., Gokilamani, S., Muthukumar, P., and Koteeswaran, S. (2025). A hybrid convolutional neural network-temporal attention mechanism approach for real-time prediction of soil moisture and temperature in precision agriculture. *Agric. Water Manage.* 295, 109182. doi: 10.14569/IJACSA.2025.0160556
- Yang, F., Fu, X., Yang, Q., and Chu, Z. (2024). Decomposition strategy and attention-based long short-term memory network for multi-step ultra-short-term agricultural power load forecasting. *Expert Syst. Appl.* 238, 121831. doi: 10.1016/j.eswa.2023.122226
- Yi, S., Li, C., and Li, Q. (2015). "A survey of fog computing: concepts, applications and issues," in *Proceedings of the 2015 workshop on mobile big data*. Hangzhou, China: ACM
- Zhang, X., Cao, Z., and Dong, W. (2020). Overview of edge computing in the agricultural internet of things: Key technologies, applications, challenges. *IEEE Access* 8, 175946–175958. doi: 10.1109/ACCESS.2020.3013005
- Zhang, Y., Wang, X., Jin, L., Ni, J., Zhu, Y., Cao, W., et al. (2025). Research and development of an IoT smart irrigation system for farmland based on LoRa and edge computing. *Comput. Electron. Agric.* 225, 109182. doi: 10.3390/agronomy15020366